






# Pyramid blending



(d)

(h)

(l)

Create a Laplacian pyramid, blend each level

- Burt, P. J. and Adelson, E. H., [A multiresolution spline with applications to image mosaics](#), ACM Transactions on Graphics, 42(4), October 1983, 217-236.

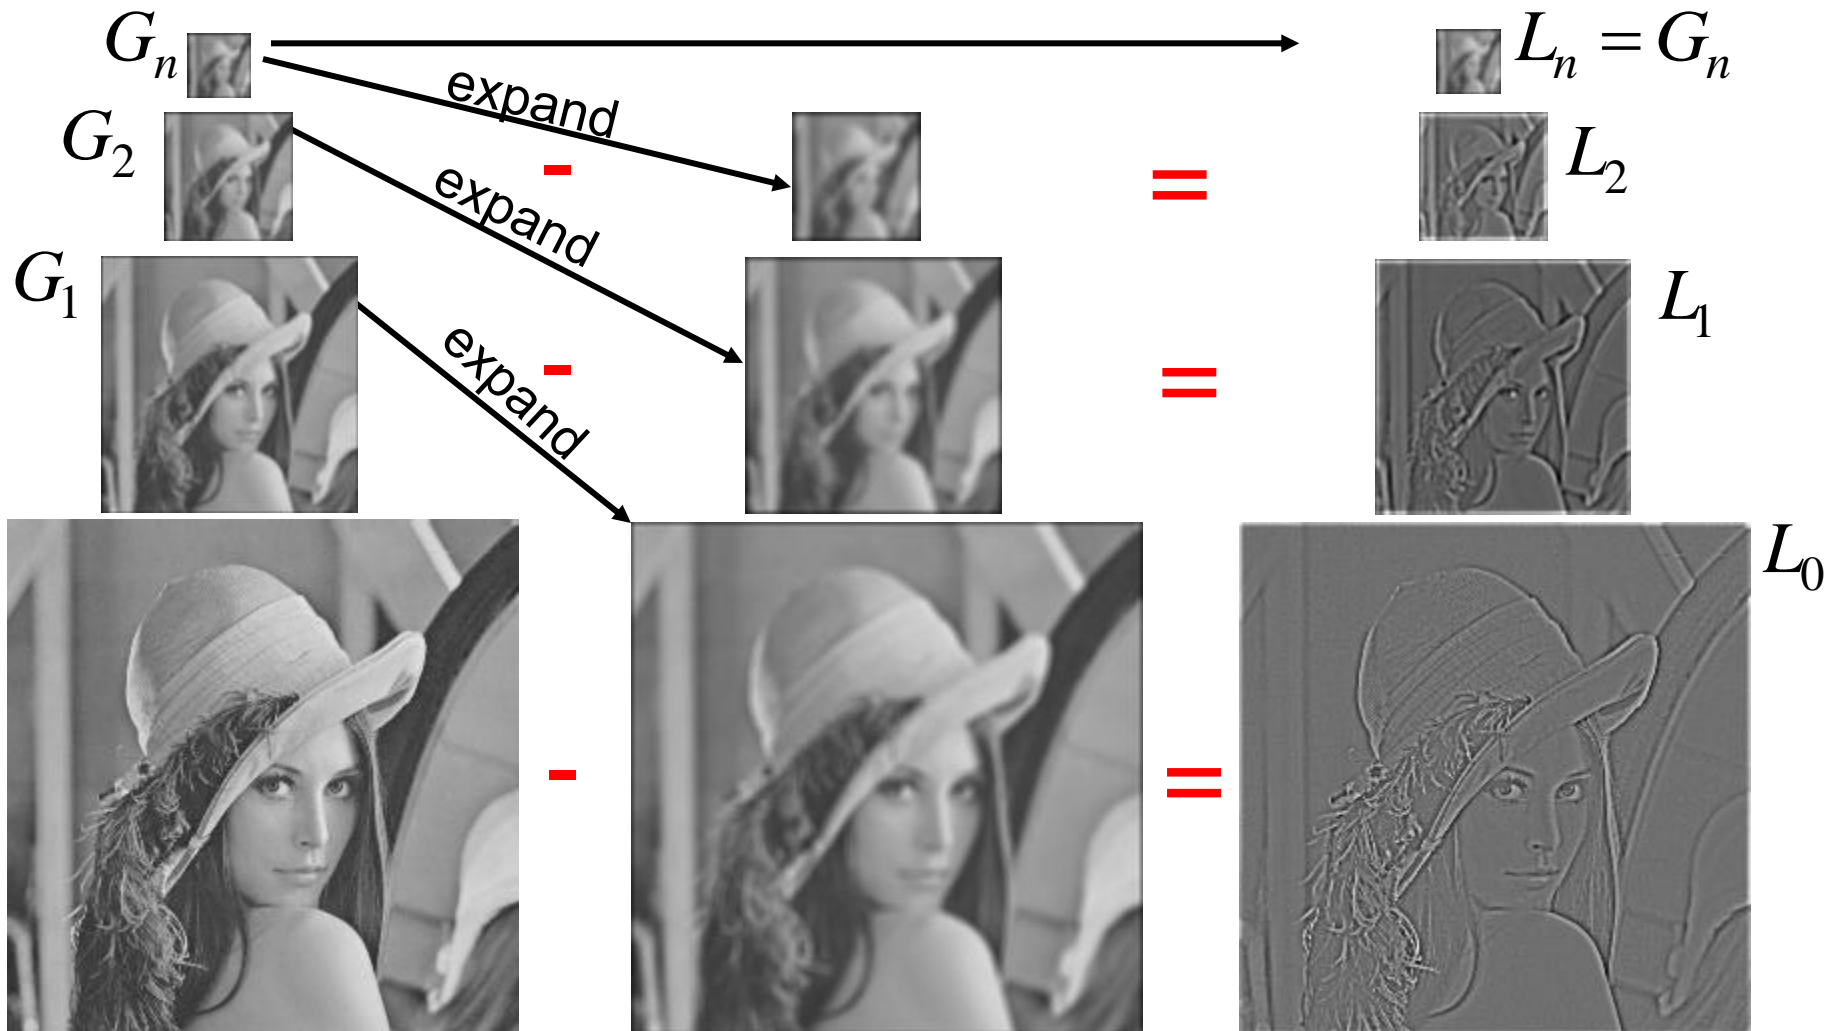
# The Laplacian Pyramid

$$L_i = G_i - \text{expand}(G_{i+1})$$


Gaussian Pyramid

$$G_i = L_i + \text{expand}(G_{i+1})$$

Laplacian Pyramid



# Alpha Blending



Optional: see Blinn (CGA, 1994) for details:

<http://ieeexplore.ieee.org/iel1/38/7531/00310740.pdf?isNumber=7531&prod=JNL&arnumber=310740&arSt=83&ared=87&arAuthor=Blinn%2C+J.F.>

Encoding blend weights:  $I(x,y) = (\alpha R, \alpha G, \alpha B, \alpha)$

color at  $p = \frac{(\alpha_1 R_1, \alpha_1 G_1, \alpha_1 B_1) + (\alpha_2 R_2, \alpha_2 G_2, \alpha_2 B_2) + (\alpha_3 R_3, \alpha_3 G_3, \alpha_3 B_3)}{\alpha_1 + \alpha_2 + \alpha_3}$

Implement this in two steps:

1. accumulate: add up the ( $\alpha$  premultiplied)  $RGB\alpha$  values at each pixel
2. normalize: divide each pixel's accumulated  $RGB$  by its  $\alpha$  value

Q: what if  $\alpha = 0$ ?

# Poisson Image Editing



sources/destinations



cloning



seamless cloning

- For more info: Perez et al, SIGGRAPH 2003

– [http://research.microsoft.com/vision/cambridge/papers/perez\\_siggraph03.pdf](http://research.microsoft.com/vision/cambridge/papers/perez_siggraph03.pdf)

# Some panorama examples



Before Siggraph Deadline:

<http://www.cs.washington.edu/education/courses/cse590ss/01wi/projects/project1/students/doug/siggraph-hires.html>

# Magic: ghost removal



M. Uyttendaele, A. Eden, and R. Szeliski.

*Eliminating ghosting and exposure artifacts in image mosaics.*

In Proceedings of the International Conference on Computer Vision and Pattern Recognition, volume 2, pages 509--516, Kauai, Hawaii, December 2001.

# Magic: ghost removal



M. Uyttendaele, A. Eden, and R. Szeliski.

*Eliminating ghosting and exposure artifacts in image mosaics.*

In Proceedings of the International Conference on Computer Vision and Pattern Recognition, volume 2, pages 509--516, Kauai, Hawaii, December 2001.

# Some panorama examples

- Every image on Google Streetview





Questions?

# CS6670: Computer Vision

Noah Snavely

## Lecture 8: Single-view Modeling



# Projective geometry




[Ames Room](#)

- Readings

- Mundy, J.L. and Zisserman, A., Geometric Invariance in Computer Vision, Appendix: Projective Geometry for Machine Vision, MIT Press, Cambridge, MA, 1992, (read 23.1 - 23.5, 23.10)
  - available online: <http://www.cs.cmu.edu/~ph/869/papers/zisser-mundy.pdf>

# Projective geometry—what's it good for?

- Uses of projective geometry
  - Drawing
  - Measurements
  - Mathematics for projection
  - Undistorting images
  - Camera pose estimation
  - **Object recognition**




[Paolo Uccello](#)

# Applications of projective geometry







Vermeer's *Music Lesson*



Reconstructions by Criminisi et al.


# Measurements on planes



Approach: unwarp then measure

# Point and line duality

- A line  $l$  is a homogeneous 3-vector
- It is  $\perp$  to every point (ray)  $p$  on the line:  $l \cdot p = 0$



What is the line  $l$  spanned by rays  $p_1$  and  $p_2$  ?



- $l$  is  $\perp$  to  $p_1$  and  $p_2 \Rightarrow l = p_1 \times p_2$
- $l$  can be interpreted as a *plane normal*

What is the intersection of two lines  $l_1$  and  $l_2$  ?

- $p$  is  $\perp$  to  $l_1$  and  $l_2 \Rightarrow p = l_1 \times l_2$

Points and lines are *dual* in projective space

# Ideal points and lines



- Ideal point (“point at infinity”)
  - $p \cong (x, y, 0)$  – parallel to image plane
  - It has infinite image coordinates

## Ideal line

- $l \cong (a, b, 0)$  – parallel to image plane
- Corresponds to a line in the image (finite coordinates)
  - goes through image origin (*principle point*)



# 3D projective geometry


- These concepts generalize naturally to 3D
  - Homogeneous coordinates
    - Projective 3D points have four coords:  $\mathbf{P} = (X,Y,Z,W)$
  - Duality
    - A plane  $\mathbf{N}$  is also represented by a 4-vector
    - Points and planes are dual in 3D:  $\mathbf{N} \mathbf{P}=0$
    - Three points define a plane, three planes define a point

# 3D to 2D: perspective projection


Projection:

$$\mathbf{p} = \begin{bmatrix} wx \\ wy \\ w \end{bmatrix} = \begin{bmatrix} * & * & * & * \\ * & * & * & * \\ * & * & * & * \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix} = \mathbf{P}\mathbf{P}$$


# Vanishing points (1D)




- Vanishing point
  - projection of a point at infinity
  - can often (but not always) project to a finite point in the image



# Vanishing points (2D)




# Vanishing points




- Properties
  - Any two parallel lines (in 3D) have the same vanishing point  $\mathbf{v}$
  - The ray from  $\mathbf{C}$  through  $\mathbf{v}$  is parallel to the lines
  - An image may have more than one vanishing point
    - in fact, every image point is a potential vanishing point

# Two point perspective




# Three point perspective



Questions?




# Vanishing lines




- Multiple Vanishing Points
  - Any set of parallel lines on the plane define a vanishing point
  - The union of all of these vanishing points is the *horizon line*
    - also called *vanishing line*
  - Note that different planes (can) define different vanishing lines

# Vanishing lines




- Multiple Vanishing Points
  - Any set of parallel lines on the plane define a vanishing point
  - The union of all of these vanishing points is the *horizon line*
    - also called *vanishing line*
  - Note that different planes (can) define different vanishing lines

# Computing vanishing points





# Computing vanishing points



$$\mathbf{P}_t = \begin{bmatrix} P_X + tD_X \\ P_Y + tD_Y \\ P_Z + tD_Z \\ 1 \end{bmatrix} \cong \begin{bmatrix} P_X / t + D_X \\ P_Y / t + D_Y \\ P_Z / t + D_Z \\ 1/t \end{bmatrix}$$

- **Properties**      $\mathbf{v} = \mathbf{IIP}_\infty$ 
  - $\mathbf{P}_\infty$  is a point at *infinity*,  $\mathbf{v}$  is its projection
  - Depends only on line *direction*
  - Parallel lines  $\mathbf{P}_0 + t\mathbf{D}$ ,  $\mathbf{P}_1 + t\mathbf{D}$  intersect at  $\mathbf{P}_\infty$

# Computing vanishing lines





- **Properties**


- $l$  is intersection of horizontal plane through  $C$  with image plane
- Compute  $l$  from two sets of parallel lines on ground plane
- All points at same height as  $C$  project to  $l$ 
  - points higher than  $C$  project above  $l$
- Provides way of comparing height of objects in the scene



# Fun with vanishing points




# Perspective cues







# Perspective cues



# Perspective cues




# Comparing heights




# Measuring height

How high is the camera?



# Computing vanishing points (from lines)




- Intersect  $p_1q_1$  with  $p_2q_2$

$$v = (p_1 \times q_1) \times (p_2 \times q_2)$$

Least squares version

- Better to use more than two lines and compute the “closest” point of intersection
- See notes by [Bob Collins](#) for one good way of doing this:
  - <http://www-2.cs.cmu.edu/~ph/869/www/notes/vanishing.txt>

# Measuring height without a ruler




Compute  $Z$  from image measurements

- Need more than vanishing points to do this

# The cross ratio

- A Projective Invariant
  - Something that does not change under projective transformations (including perspective projection)

The *cross-ratio* of 4 collinear points



$$\frac{\| \mathbf{P}_3 - \mathbf{P}_1 \| \| \mathbf{P}_4 - \mathbf{P}_2 \|}{\| \mathbf{P}_3 - \mathbf{P}_2 \| \| \mathbf{P}_4 - \mathbf{P}_1 \|}$$

$$\mathbf{P}_i = \begin{bmatrix} X_i \\ Y_i \\ Z_i \\ 1 \end{bmatrix}$$


Can permute the point ordering

$$\frac{\| \mathbf{P}_1 - \mathbf{P}_3 \| \| \mathbf{P}_4 - \mathbf{P}_2 \|}{\| \mathbf{P}_1 - \mathbf{P}_2 \| \| \mathbf{P}_4 - \mathbf{P}_3 \|}$$

- $4! = 24$  different orders (but only 6 distinct values)

This is the fundamental invariant of projective geometry

# Measuring height



scene points represented as

$$\mathbf{P} = \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix}$$

image points as

$$\mathbf{p} = \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

$$\frac{\|\mathbf{T} - \mathbf{B}\| \|\infty - \mathbf{R}\|}{\|\mathbf{R} - \mathbf{B}\| \|\infty - \mathbf{T}\|} = \frac{H}{R}$$


scene cross ratio

$$\frac{\|\mathbf{t} - \mathbf{b}\| \|\mathbf{v}_Z - \mathbf{r}\|}{\|\mathbf{r} - \mathbf{b}\| \|\mathbf{v}_Z - \mathbf{t}\|} = \frac{H}{R}$$

image cross ratio



# Measuring height



# 3D Modeling from a photograph




*St. Jerome in his Study*, H. Steenwick

# 3D Modeling from a photograph



# 3D Modeling from a photograph





*Flagellation*, Piero della Francesca

# 3D Modeling from a photograph



video by Antonio Criminisi

# 3D Modeling from a photograph



# Questions?

- 3-minute break

# Camera calibration

- Goal: estimate the camera parameters
  - Version 1: solve for projection matrix

$$\mathbf{X} = \begin{bmatrix} wx \\ wy \\ w \end{bmatrix} = \begin{bmatrix} * & * & * & * \\ * & * & * & * \\ * & * & * & * \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix} = \mathbf{P}\mathbf{X}$$

- Version 2: solve for camera parameters separately
  - intrinsics (focal length, principle point, pixel size)
  - extrinsics (rotation angles, translation)
  - radial distortion



# Vanishing points and projection matrix

$$\mathbf{\Pi} = \begin{bmatrix} * & * & * & * \\ * & * & * & * \\ * & * & * & * \end{bmatrix} = \begin{bmatrix} \boldsymbol{\pi}_1 & \boldsymbol{\pi}_2 & \boldsymbol{\pi}_3 & \boldsymbol{\pi}_4 \end{bmatrix}$$

- $\boldsymbol{\pi}_1 = \mathbf{\Pi} \begin{bmatrix} 1 & 0 & 0 & 0 \end{bmatrix}^T = \mathbf{v}_x$  (X vanishing point)
- similarly,  $\boldsymbol{\pi}_2 = \mathbf{v}_y$ ,  $\boldsymbol{\pi}_3 = \mathbf{v}_z$
- $\boldsymbol{\pi}_4 = \mathbf{\Pi} \begin{bmatrix} 0 & 0 & 0 & 1 \end{bmatrix}^T =$  projection of world origin

$$\mathbf{\Pi} = \begin{bmatrix} \mathbf{v}_x & \mathbf{v}_y & \mathbf{v}_z & \mathbf{o} \end{bmatrix}$$


Not So Fast! We only know  $\mathbf{v}$ 's up to a scale factor

$$\mathbf{\Pi} = \begin{bmatrix} a \mathbf{v}_x & b \mathbf{v}_y & c \mathbf{v}_z & \mathbf{o} \end{bmatrix}$$

- Can fully specify by providing 3 reference points

# Calibration using a reference object

- Place a known object in the scene
  - identify correspondence between image and scene
  - compute mapping from scene to image




## Issues

- must know geometry very accurately
- must know 3D->2D correspondence

# Estimating the projection matrix

- Place a known object in the scene
  - identify correspondence between image and scene
  - compute mapping from scene to image



$$\begin{bmatrix} u_i \\ v_i \\ 1 \end{bmatrix} \approx \begin{bmatrix} m_{00} & m_{01} & m_{02} & m_{03} \\ m_{10} & m_{11} & m_{12} & m_{13} \\ m_{20} & m_{21} & m_{22} & m_{23} \end{bmatrix} \begin{bmatrix} X_i \\ Y_i \\ Z_i \\ 1 \end{bmatrix}$$

# Direct linear calibration

$$\begin{bmatrix} u_i \\ v_i \\ 1 \end{bmatrix} \cong \begin{bmatrix} m_{00} & m_{01} & m_{02} & m_{03} \\ m_{10} & m_{11} & m_{12} & m_{13} \\ m_{20} & m_{21} & m_{22} & m_{23} \end{bmatrix} \begin{bmatrix} X_i \\ Y_i \\ Z_i \\ 1 \end{bmatrix}$$

$$u_i = \frac{m_{00}X_i + m_{01}Y_i + m_{02}Z_i + m_{03}}{m_{20}X_i + m_{21}Y_i + m_{22}Z_i + m_{23}}$$

$$v_i = \frac{m_{10}X_i + m_{11}Y_i + m_{12}Z_i + m_{13}}{m_{20}X_i + m_{21}Y_i + m_{22}Z_i + m_{23}}$$

$$u_i(m_{20}X_i + m_{21}Y_i + m_{22}Z_i + m_{23}) = m_{00}X_i + m_{01}Y_i + m_{02}Z_i + m_{03}$$

$$v_i(m_{20}X_i + m_{21}Y_i + m_{22}Z_i + m_{23}) = m_{10}X_i + m_{11}Y_i + m_{12}Z_i + m_{13}$$

$$\begin{bmatrix} X_i & Y_i & Z_i & 1 & 0 & 0 & 0 & 0 & -u_iX_i & -u_iY_i & -u_iZ_i & -u_i \\ 0 & 0 & 0 & 0 & X_i & Y_i & Z_i & 1 & -v_iX_i & -v_iY_i & -v_iZ_i & -v_i \end{bmatrix} \begin{bmatrix} m_{00} \\ m_{01} \\ m_{02} \\ m_{03} \\ m_{10} \\ m_{11} \\ m_{12} \\ m_{13} \\ m_{20} \\ m_{21} \\ m_{22} \\ m_{23} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$$

# Direct linear calibration

$$\begin{bmatrix}
 X_1 & Y_1 & Z_1 & 1 & 0 & 0 & 0 & 0 & -u_1 X_1 & -u_1 Y_1 & -u_1 Z_1 & -u_1 \\
 0 & 0 & 0 & 0 & X_1 & Y_1 & Z_1 & 1 & -v_1 X_1 & -v_1 Y_1 & -v_1 Z_1 & -v_1 \\
 & & & & & & & \vdots & & & & \\
 X_n & Y_n & Z_n & 1 & 0 & 0 & 0 & 0 & -u_n X_n & -u_n Y_n & -u_n Z_n & -u_n \\
 0 & 0 & 0 & 0 & X_n & Y_n & Z_n & 1 & -v_n X_n & -v_n Y_n & -v_n Z_n & -v_n
 \end{bmatrix}
 \begin{bmatrix}
 m_{00} \\
 m_{01} \\
 m_{02} \\
 m_{03} \\
 m_{10} \\
 m_{11} \\
 m_{12} \\
 m_{13} \\
 m_{20} \\
 m_{21} \\
 m_{22}
 \end{bmatrix}
 =
 \begin{bmatrix}
 0 \\
 0 \\
 \vdots \\
 0 \\
 0
 \end{bmatrix}$$

Can solve for  $m_{ij}$  by linear least squares

- use eigenvector trick that we used for homographies


# Direct linear calibration

- Advantage:
  - Very simple to formulate and solve
- Disadvantages:
  - Doesn't tell you the camera parameters
  - Doesn't model radial distortion
  - Hard to impose constraints (e.g., known  $f$ )
  - Doesn't minimize the right error function

For these reasons, *nonlinear methods* are preferred

- Define error function  $E$  between projected 3D points and image positions
  - $E$  is nonlinear function of intrinsics, extrinsics, radial distortion
- Minimize  $E$  using nonlinear optimization techniques

# Alternative: multi-plane calibration



Images courtesy Jean-Yves Bouguet, Intel Corp.

## Advantage

- Only requires a plane
- Don't have to know positions/orientations
- Good code available online! (including in OpenCV)
  - Matlab version by Jean-Yves Bouguet:  
[http://www.vision.caltech.edu/bouguetj/calib\\_doc/index.html](http://www.vision.caltech.edu/bouguetj/calib_doc/index.html)
  - Zhengyou Zhang's web site: <http://research.microsoft.com/~zhang/Calib/>

# Some Related Techniques

- Image-Based Modeling and Photo Editing
  - Mok et al., SIGGRAPH 2001
  - <http://graphics.csail.mit.edu/ibedit/>
- Single View Modeling of Free-Form Scenes
  - Zhang et al., CVPR 2001
  - <http://grail.cs.washington.edu/projects/svm/>
- Tour Into The Picture
  - Anjyo et al., SIGGRAPH 1997
  - [http://koigakubo.hitachi.co.jp/little/DL\\_TipE.html](http://koigakubo.hitachi.co.jp/little/DL_TipE.html)



# More than one view?


- W
- W



a

What's the transformation?

HON. ABRAHAM LINCOLN, President of United States.






Public Library, Stereoscopic Looking Room, Chicago, by Phillips, 1923






Mark Twain at Pool Table", no date, UCR Museum of Photography



# Epipolar geometry




Two images captured by a purely horizontal translating camera  
(*rectified* stereo pair)

$x_2 - x_1 =$  the *disparity* of pixel  $(x_1, y_1)$

# Stereo matching algorithms

- Match Pixels in Conjugate Epipolar Lines
  - Assume brightness constancy
  - This is a tough problem
  - Numerous approaches
    - A good survey and evaluation: <http://www.middlebury.edu/stereo/>

# Your basic stereo algorithm



For each epipolar line

For each pixel in the left image

- compare with every pixel on same epipolar line in right image
- pick pixel with minimum match cost


Improvement: match *windows*



# Window size



$W = 3$



$W = 20$

Better results with *adaptive window*

- T. Kanade and M. Okutomi, [A Stereo Matching Algorithm with an Adaptive Window: Theory and Experiment](#), Proc. International Conference on Robotics and Automation, 1991.
- D. Scharstein and R. Szeliski. [Stereo matching with nonlinear diffusion](#). International Journal of Computer Vision, 28(2):155-174, July 1998

# Stereo results

- Data from University of Tsukuba
- Similar results on other images without ground truth




Scene



Ground truth

# Results with window search



Window-based matching  
(best window size)



Ground truth

# Better methods exist...



State of the art method


Boykov et al., [Fast Approximate Energy Minimization via Graph Cuts](#),  
International Conference on Computer Vision, September 1999.



Ground truth

For the latest and greatest: <http://www.middlebury.edu/stereo/>

# Stereo as energy minimization



- What defines a good stereo correspondence?
  1. Match quality
    - Want each pixel to find a good match in the other image
  2. Smoothness
    - If two pixels are adjacent, they should (usually) move about the same amount

# Stereo as energy minimization

- Expressing this mathematically

1. Match quality

- Want each pixel to find a good match in the other image

$$matchCost = \sum_{x,y} \|I(x, y) - J(x + d_{xy}, y)\|$$

2. Smoothness

- If two pixels are adjacent, they should (usually) move about the same amount


$$smoothnessCost = \sum_{neighbor\ pixels\ p,q} |d_p - d_q|$$

- We want to minimize  $Energy = matchCost + smoothnessCost$

- This is a special type of energy function known as an MRF (Markov Random Field)

- Effective and fast algorithms have been recently developed:
  - Graph cuts, belief propagation....
  - for more details (and code): <http://vision.middlebury.edu/MRF/>
  - Great [tutorials](#) available online (including video of talks)

# Depth from disparity



$$\text{disparity} = x - x' = \frac{\text{baseline} * f}{z}$$

# Real-time stereo



[Nomad robot](http://www.frc.ri.cmu.edu/projects/meteorobot/index.html) searches for meteorites in Antarctica  
<http://www.frc.ri.cmu.edu/projects/meteorobot/index.html>

- Used for robot navigation (and other tasks)
  - Several software-based real-time stereo techniques have been developed (most based on simple discrete search)





# Stereo reconstruction pipeline

- Steps
  - Calibrate cameras
  - Rectify images
  - Compute disparity
  - Estimate depth



What will cause errors?

- Camera calibration errors
- Poor image resolution
- Occlusions
- Violations of brightness constancy (specular reflections)
- Large motions
- Low-contrast image regions

# Active stereo with structured light





Li Zhang's one-shot stereo



- Project “structured” light patterns onto the object
  - simplifies the correspondence problem

# Laser scanning



Digital Michelangelo Project  
<http://graphics.stanford.edu/projects/mich/>


- Optical triangulation
  - Project a single stripe of laser light
  - Scan it across the surface of the object
  - This is a very precise version of structured light scanning

# Laser scanned models




*The Digital Michelangelo Project, Levoy et al.*

# Laser scanned models




*The Digital Michelangelo Project, Levoy et al.*

# Laser scanned models



*The Digital Michelangelo Project, Levoy et al.*

# Laser scanned models



*The Digital Michelangelo Project, Levoy et al.*