

CS4670/5670: Computer Vision

Noah Snavely

Single-view modeling, Part 2



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>


Announcements

- Midterm to be handed in Thursday by 5pm
- Please hand it back at my office, Upson 4157

Three point perspective




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** $\mathbf{v} = \mathbf{H}\mathbf{P}_\infty$
 - \mathbf{P}_∞ 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}_∞

Computing vanishing lines





• Properties


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




Fun with vanishing points




Perspective cues



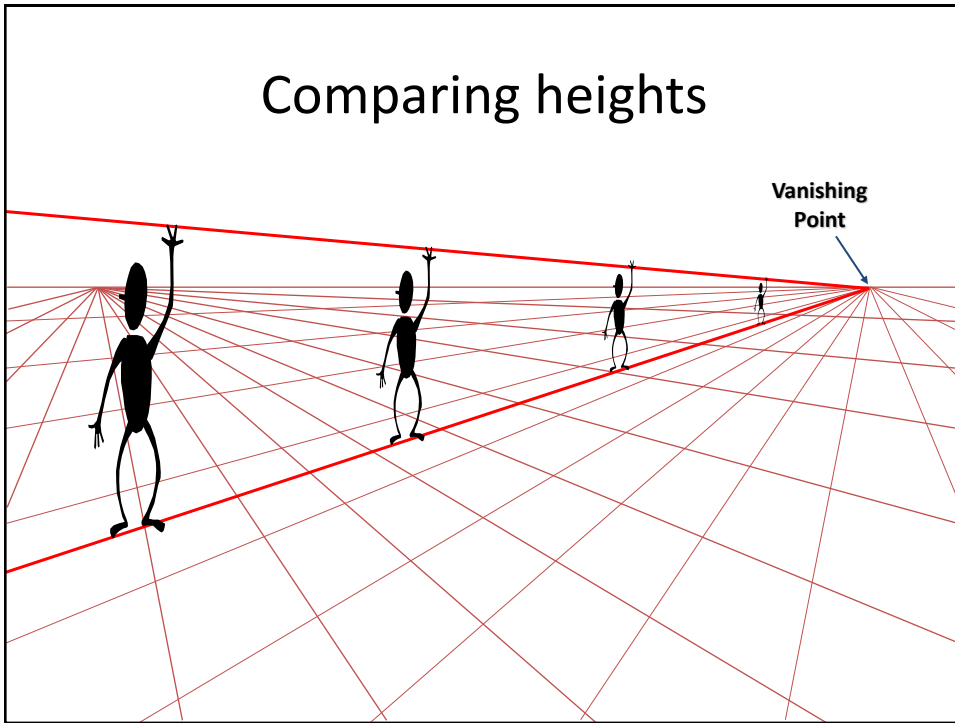
Perspective cues




Perspective cues




Comparing heights



Measuring height



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](http://www-2.cs.cmu.edu/~ph/869/www/notes/vanishing.txt) for one good way of doing this:
 - <http://www-2.cs.cmu.edu/~ph/869/www/notes/vanishing.txt>

Measuring height without a ruler



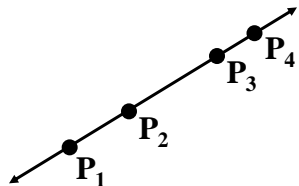
Compute Z from image measurements

- 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




$$\frac{\|P_3 - P_1\| \|P_4 - P_2\|}{\|P_3 - P_2\| \|P_4 - P_1\|}$$

$$P_i = \begin{bmatrix} X_i \\ Y_i \\ Z_i \\ 1 \end{bmatrix}$$

- Can permute the point ordering
- $$\frac{\|P_1 - P_3\| \|P_4 - P_2\|}{\|P_1 - P_2\| \|P_4 - P_3\|}$$
- $4! = 24$ different orders (but only 6 distinct values)
- This is the fundamental invariant of projective geometry

Measuring height



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

scene cross ratio

$$\frac{\|t - b\| \|v_z - r\|}{\|r - b\| \|v_z - t\|} = \frac{H}{R}$$



image cross ratio

scene points represented as

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

image points as

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



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



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

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

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

$$\mathbf{\Pi} = [\mathbf{v}_x \quad \mathbf{v}_y \quad \mathbf{v}_z \quad \mathbf{o}]$$


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

$$\mathbf{\Pi} = [a \mathbf{v}_x \quad b \mathbf{v}_y \quad c \mathbf{v}_z \quad \mathbf{o}]$$

- 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


Chromaglyphs



Courtesy of Bruce Culbertson, HP Labs
http://www.hpl.hp.com/personal/Bruce_Culbertson/ibr98/chromagl.htm

Estimating the projection matrix

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



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

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

$$\begin{bmatrix} X_1 & Y_1 & Z_1 & 1 & 0 & 0 & 0 & 0 & -u_1X_1 & -u_1Y_1 & -u_1Z_1 & -u_1 \\ 0 & 0 & 0 & 0 & X_1 & Y_1 & Z_1 & 1 & -v_1X_1 & -v_1Y_1 & -v_1Z_1 & -v_1 \\ & & & & & & & & & & & \\ X_n & Y_n & Z_n & 1 & 0 & 0 & 0 & 0 & -u_nX_n & -u_nY_n & -u_nZ_n & -u_n \\ 0 & 0 & 0 & 0 & X_n & Y_n & Z_n & 1 & -v_nX_n & -v_nY_n & -v_nZ_n & -v_n \end{bmatrix} \begin{bmatrix} m_{00} \\ m_{01} \\ m_{02} \\ m_{03} \\ m_{10} \\ m_{11} \\ m_{12} \\ m_{13} \\ m_{20} \\ m_{21} \\ m_{22} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}$$

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
 - 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