

Informal Early Feedback


- **MPs**
 - Provide more detailed comments in code
 - Be clear about the deliverables
 - Provide sample output
- **Lectures**
 - More opportunities for interaction
 - Provide more context / better specify what to do in break-out rooms.
- **Slides**
 - Better pointers to suggested readings
- **Piazza**
 - Monitor and respond more often
- **More office hours**

Will try, but can't promise

- **MPs**
 - Faster grading
- **Slides**
 - Annotated slides. PPT on iPad doesn't work well with videos / equations.
- **Lectures**
 - More examples

Which way is the cube rotating?

- A. Left to right
- B. Right to left




Necker cube


Multi-view geometry



Structure from motion




Structure from motion




- **Structure:** Given *known cameras* and projections of the same 3D point in two or more images, compute the 3D coordinates of that point

Structure from motion



- **Motion:** Given a set of *known* 3D points seen by a camera, compute the camera parameters

Structure from motion




- **Bootstrapping the process:** Given a set of 2D point correspondences in *two images*, compute the camera parameters

Two-view geometry




Epipolar geometry



- **Baseline** – line connecting the two camera centers
- **Epipolar Plane** – plane containing baseline (1D family)
- **Epipoles**
 - = intersections of baseline with image planes
 - = projections of the other camera center
 - = vanishing points of the motion direction


Epipolar geometry



- **Baseline** – line connecting the two camera centers
- **Epipolar Plane** – plane containing baseline (1D family)
- **Epipoles**
 - = intersections of baseline with image planes
 - = projections of the other camera center
 - = vanishing points of the motion direction
- **Epipolar Lines** - intersections of epipolar plane with image planes (always come in corresponding pairs)



Example 1

- Converging cameras



Example 2

- Motion parallel to the image plane



Example 3




Example 3



- Motion is perpendicular to the image plane
- Epipole is the “focus of expansion” and the principal point

Motion perpendicular to image plane




Kubrick // One-Point Perspective

from kognada PLUS 1 year ago NOT YET RATED


<http://vimeo.com/48425421>

Epipolar constraint



- If we observe a point \mathbf{x} in one image, where can the corresponding point \mathbf{x}' be in the other image?

Epipolar constraint




- Potential matches for x have to lie on the corresponding epipolar line l' .
- Potential matches for x' have to lie on the corresponding epipolar line l .


Epipolar constraint example



Epipolar constraint: Calibrated case



Epipolar constraint: Calibrated case




- Intrinsic and extrinsic parameters of the cameras are known, world coordinate system is set to that of the first camera
- Then the projection matrices are given by $K[I \mid \mathbf{0}]$ and $K'[R \mid t]$
- We can multiply the projection matrices (and the image points) by the inverse of the calibration matrices to get *normalized* image coordinates:

$$\mathbf{x}_{\text{norm}} = K^{-1} \mathbf{x}_{\text{pixel}} = [I \mid \mathbf{0}] X, \quad \mathbf{x}'_{\text{norm}} = K'^{-1} \mathbf{x}'_{\text{pixel}} = [R \mid t] X$$


Epipolar constraint: Calibrated case

Derivation




Epipolar constraint: Calibrated case

Derivation



Epipolar constraint: Calibrated case

Simplification




Recall:
$$\mathbf{a} \times \mathbf{b} = \begin{bmatrix} 0 & -a_z & a_y \\ a_z & 0 & -a_x \\ -a_y & a_x & 0 \end{bmatrix} \begin{bmatrix} b_x \\ b_y \\ b_z \end{bmatrix} = [\mathbf{a}_\times] \mathbf{b}$$

The vectors Rx , t , and x' are coplanar

Epipolar constraint: Calibrated case

Simplification



$$x' \cdot [t \times (Rx)] = 0$$


$$x' \cdot [t_{\times}] Rx = 0$$

$$x'^T E x = 0$$

Recall:
$$\mathbf{a} \times \mathbf{b} = \begin{bmatrix} 0 & -a_z & a_y \\ a_z & 0 & -a_x \\ -a_y & a_x & 0 \end{bmatrix} \begin{bmatrix} b_x \\ b_y \\ b_z \end{bmatrix} = [\mathbf{a}_{\times}] \mathbf{b}$$

The vectors Rx , t , and x' are coplanar

Epipolar constraint: Calibrated case




$$x' \cdot [t \times (Rx)] = 0 \quad \Rightarrow \quad x'^T [t_x] Rx = 0 \quad \Rightarrow \quad x'^T E x = 0$$



Essential Matrix
(Longuet-Higgins, 1981)

The vectors Rx , t , and x' are coplanar

Epipolar constraint: Calibrated case




$$\mathbf{x}'^T \mathbf{E} \mathbf{x} = 0$$

- $\mathbf{E} \mathbf{x}$ is the epipolar line associated with \mathbf{x} ($l' = \mathbf{E} \mathbf{x}$)
 - Recall: a line is given by $ax + by + c = 0$ or


$$\mathbf{l}^T \mathbf{x} = 0 \quad \text{where} \quad \mathbf{l} = \begin{bmatrix} a \\ b \\ c \end{bmatrix}, \quad \mathbf{x} = \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

Epipolar constraint: Calibrated case




- $E x$ is the epipolar line associated with x ($l' = E x$)
- $E^T x'$ is the epipolar line associated with x' ($l = E^T x'$)
- $E e = 0$ and $E^T e' = 0$
- E is singular (rank two)
- E has five degrees of freedom

Epipolar constraint: Uncalibrated case



- The calibration matrices \mathbf{K} and \mathbf{K}' of the two cameras are unknown
- We can write the epipolar constraint in terms of *unknown* normalized coordinates:


Epipolar constraint: Uncalibrated case



- The calibration matrices \mathbf{K} and \mathbf{K}' of the two cameras are unknown
- We can write the epipolar constraint in terms of *unknown* normalized coordinates:

$$\hat{\mathbf{x}}'^T \mathbf{E} \hat{\mathbf{x}} = 0 \quad \hat{\mathbf{x}} = \mathbf{K}^{-1} \mathbf{x}, \quad \hat{\mathbf{x}}' = \mathbf{K}'^{-1} \mathbf{x}'$$

Epipolar constraint: Uncalibrated case




$$\hat{x}'^T E \hat{x} = 0 \quad \Rightarrow \quad x'^T F x = 0 \quad \text{with} \quad F = K'^{-T} E K^{-1}$$

$$\hat{x} = K^{-1} x$$

$$\hat{x}' = K'^{-1} x'$$

Fundamental Matrix
(Faugeras and Luong, 1992)

Epipolar constraint: Uncalibrated case



$$\hat{\mathbf{x}}'^T \mathbf{E} \hat{\mathbf{x}} = 0 \quad \longrightarrow \quad \mathbf{x}'^T \mathbf{F} \mathbf{x} = 0 \quad \text{with} \quad \mathbf{F} = \mathbf{K}'^{-T} \mathbf{E} \mathbf{K}^{-1}$$

- $\mathbf{F} \mathbf{x}$ is the epipolar line associated with \mathbf{x} ($l' = \mathbf{F} \mathbf{x}$)
- $\mathbf{F}^T \mathbf{x}'$ is the epipolar line associated with \mathbf{x}' ($l = \mathbf{F}^T \mathbf{x}'$)
- $\mathbf{F} \mathbf{e} = 0$ and $\mathbf{F}^T \mathbf{e}' = 0$
- \mathbf{F} is singular (rank two)
- \mathbf{F} has *seven* degrees of freedom

Estimating the fundamental matrix



The eight-point algorithm

$$\mathbf{x} = (u, v, 1)^T, \quad \mathbf{x}' = (u', v', 1)$$



The eight-point algorithm

$$\mathbf{x} = (u, v, 1)^T, \quad \mathbf{x}' = (u', v', 1)$$

$$\begin{bmatrix} u' & v' & 1 \end{bmatrix} \begin{bmatrix} f_{11} & f_{12} & f_{13} \\ f_{21} & f_{22} & f_{23} \\ f_{31} & f_{32} & f_{33} \end{bmatrix} \begin{bmatrix} u \\ v \\ 1 \end{bmatrix} = 0 \quad \Rightarrow \quad \begin{bmatrix} u'u & u'v & u' & v'u & v'v & v' & u & v & 1 \end{bmatrix} \begin{bmatrix} f_{11} \\ f_{12} \\ f_{13} \\ f_{21} \\ f_{22} \\ f_{23} \\ f_{31} \\ f_{32} \\ f_{33} \end{bmatrix} = 0$$

Solve homogeneous linear system using eight or more matches

Enforce rank-2 constraint (take SVD of \mathbf{F} and throw out the smallest singular value)



Problem with eight-point algorithm

$$\begin{bmatrix} u'u & u'v & u' & v'u & v'v & v' & u & v \end{bmatrix} \begin{bmatrix} f_{11} \\ f_{12} \\ f_{13} \\ f_{21} \\ f_{22} \\ f_{23} \\ f_{31} \\ f_{32} \end{bmatrix} = -1$$

Problem with eight-point algorithm

250906.36	183269.57	921.81	200931.10	146766.13	738.21	272.19	198.81
2692.28	131633.03	176.27	6196.73	302975.59	405.71	15.27	746.79
416374.23	871684.30	935.47	408110.89	854384.92	916.90	445.10	931.81
191183.60	171759.40	410.27	416435.62	374125.90	893.65	465.99	418.65
48988.86	30401.76	57.89	298604.57	185309.58	352.87	846.22	525.15
164786.04	546559.67	813.17	1998.37	6628.15	9.86	202.65	672.14
116407.01	2727.75	138.89	169941.27	3982.21	202.77	838.12	19.64
135384.58	75411.13	198.72	411350.03	229127.78	603.79	681.28	379.48

$$\begin{bmatrix} f_{11} \\ f_{12} \\ f_{13} \\ f_{21} \\ f_{22} \\ f_{23} \\ f_{31} \\ f_{32} \end{bmatrix} = -1$$

Poor numerical conditioning

Can be fixed by rescaling the data

The normalized eight-point algorithm

(Hartley, 1995)

- Center the image data at the origin, and scale it so the mean squared distance between the origin and the data points is 2 pixels
- Use the eight-point algorithm to compute \mathbf{F} from the normalized points
- Enforce the rank-2 constraint (for example, take SVD of \mathbf{F} and throw out the smallest singular value)
- Transform fundamental matrix back to original units: if \mathbf{T} and \mathbf{T}' are the normalizing transformations in the two images, then the fundamental matrix in original coordinates is $\mathbf{T}'^T \mathbf{F} \mathbf{T}$

Seven-point algorithm

- Set up least squares system with seven pairs of correspondences and solve for null space (two vectors) using SVD
- Solve for linear combination of null space vectors that satisfies $\det(F)=0$



Nonlinear estimation

- Linear estimation minimizes the sum of squared *algebraic* distances between points \mathbf{x}'_i and epipolar lines $F \mathbf{x}_i$ (or points \mathbf{x}_i and epipolar lines $F^T \mathbf{x}'_i$):


$$\sum_{i=1}^N (\mathbf{x}'_i{}^T F \mathbf{x}_i)^2$$

- Nonlinear approach: minimize sum of squared *geometric* distances

$$\sum_{i=1}^N \left[d^2(\mathbf{x}'_i, F \mathbf{x}_i) + d^2(\mathbf{x}_i, F^T \mathbf{x}'_i) \right]$$




Comparison of estimation algorithms



	8-point	Normalized 8-point	Nonlinear least squares
Av. Dist. 1	2.33 pixels	0.92 pixel	0.86 pixel
Av. Dist. 2	2.18 pixels	0.85 pixel	0.80 pixel

The Fundamental Matrix Song



<http://danielwedge.com/fmatrix/>

From epipolar geometry to camera calibration

- Estimating the fundamental matrix is known as “weak calibration”
- If we know the calibration matrices of the two cameras, we can estimate the essential matrix: $\mathbf{E} = \mathbf{K}'^T \mathbf{F} \mathbf{K}$
- The essential matrix gives us the relative rotation and translation between the cameras, or their extrinsic parameters
- Alternatively, if the calibration matrices are known, the [five-point algorithm](#) can be used to estimate relative camera pose