Epipolar geometry

Photo by Frank Dellaert

Many slides adapted from [J. Johnson and D. Fouhey]

Outline

- Motivation
- Epipolar geometry setup
- Epipolar constraint
- Essential matrix
- Fundamental matrix
- Estimating the fundamental matrix
- Monocular visual odometry

What happens in two views

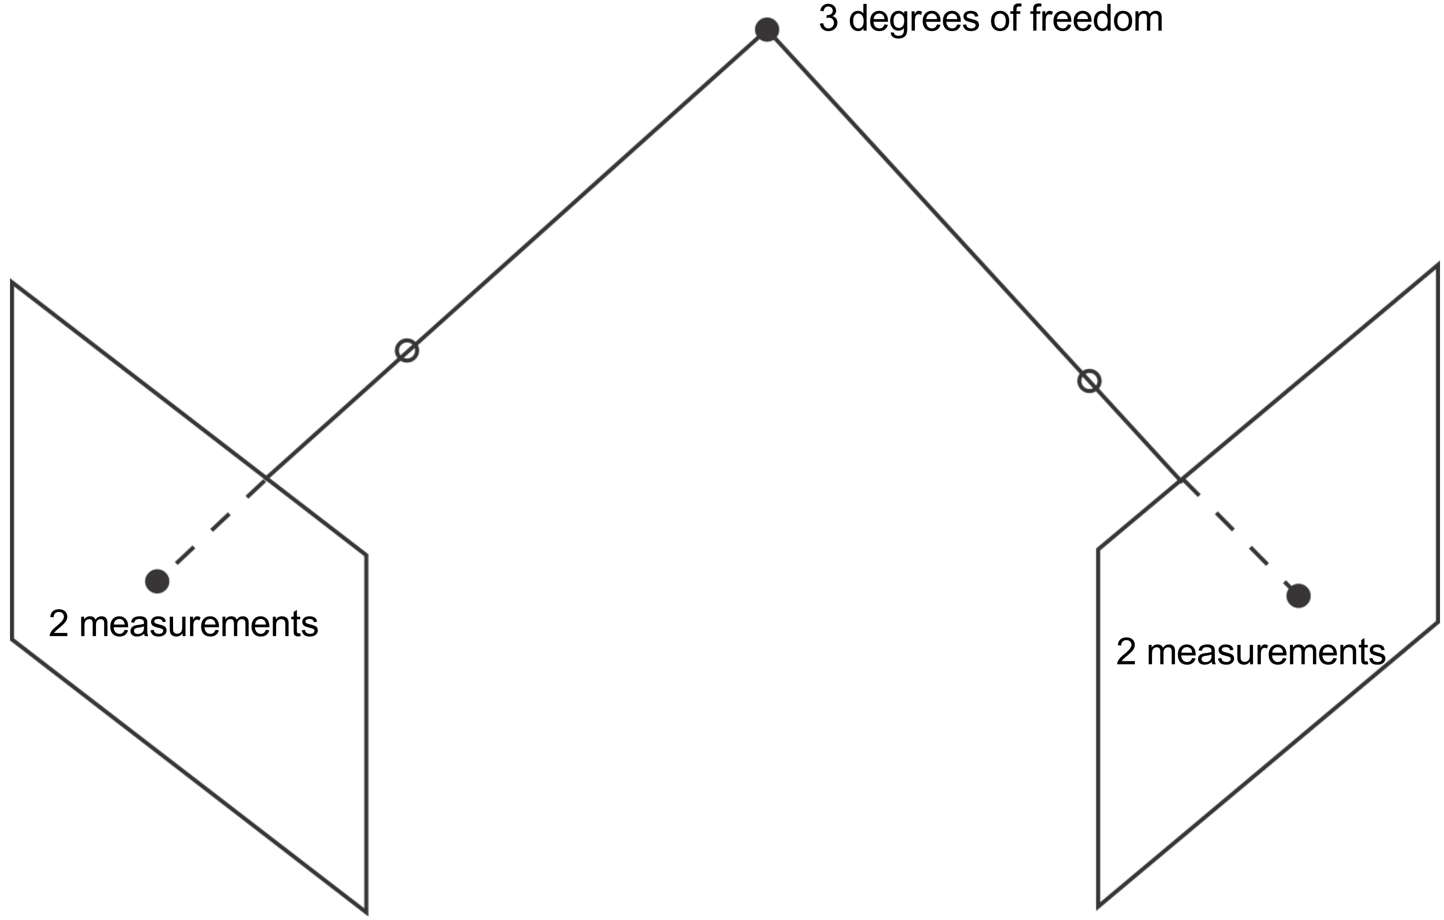

All of Camera Geometry

- From the picture
  - two views of a point give four measurements of three DOF
  - this means
    - correspondence is constrained
    - if we have enough points and enough pix we can recover
      - points
      - cameras

Two views of the same 3D scene, because:

- You have two
  - Eyes
  - Cameras
- OR the camera moves

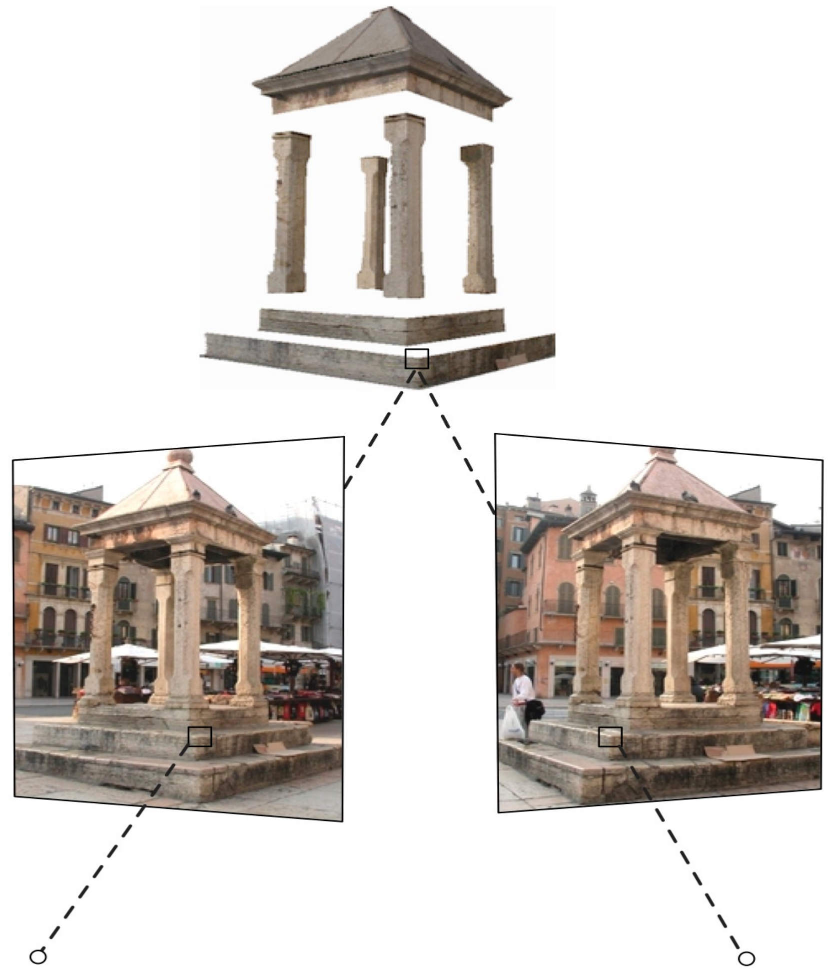

Consider two views of the same 3D scene

- What constraints must hold between two projections of the same 3D point?

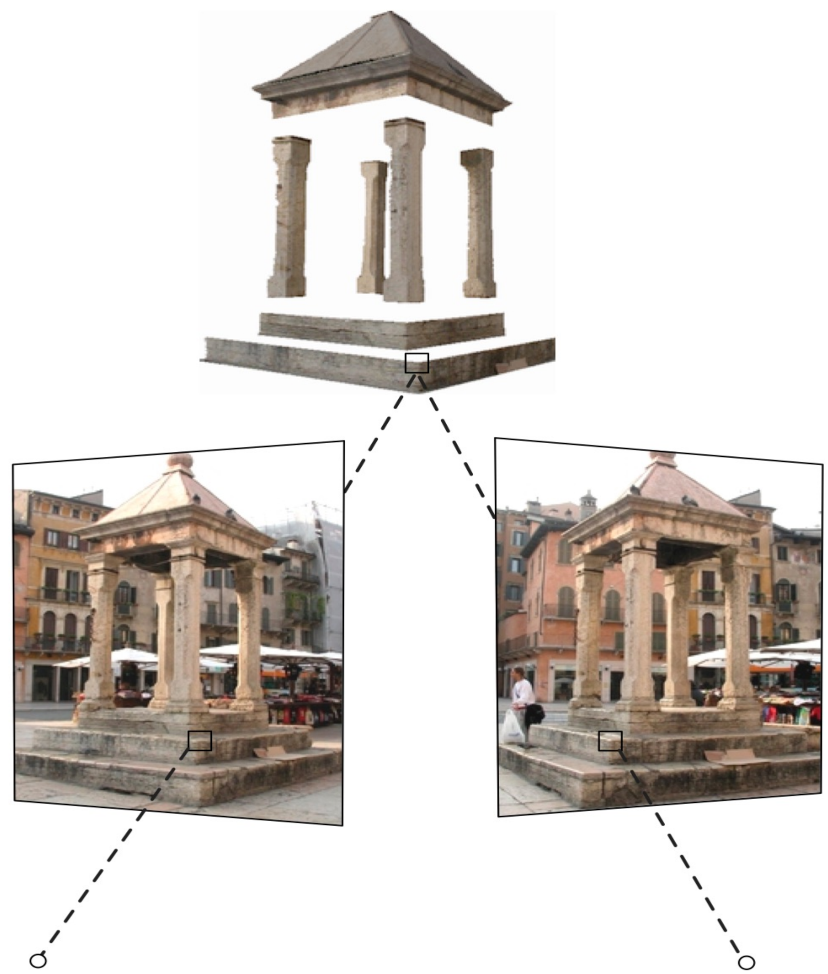

Consider two views of the same 3D scene

- Given a 2D point in one view, where can we find the corresponding point in the other view?

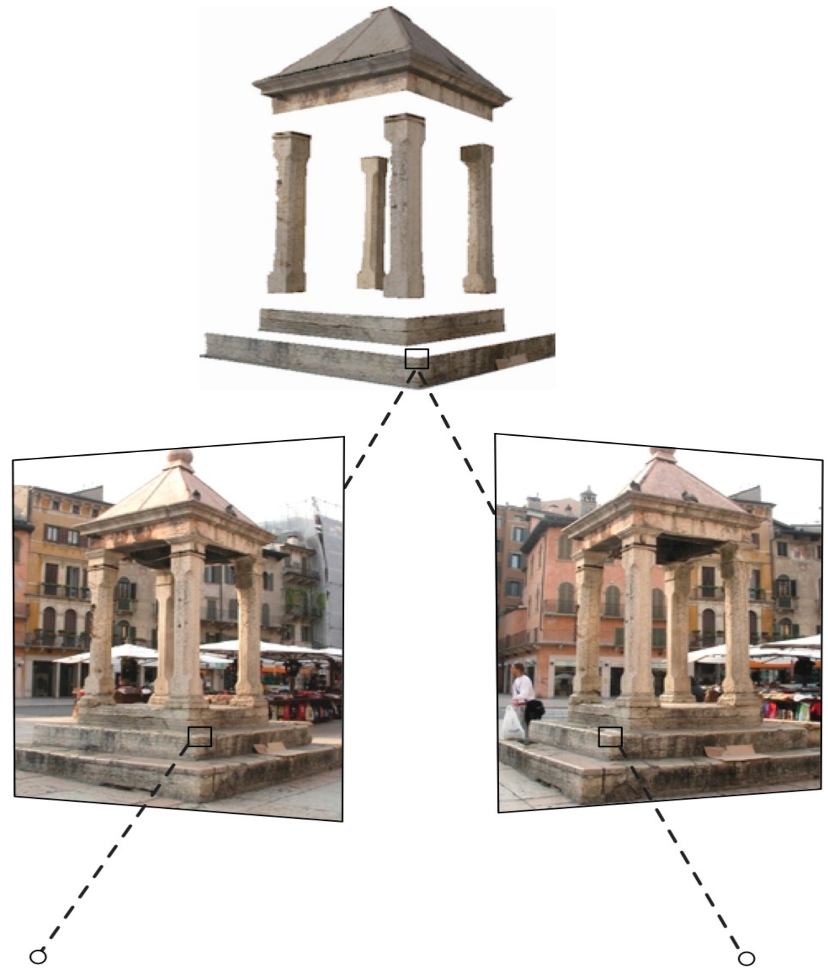

Consider two views of the same 3D scene

- Given only 2D correspondences, how can we calibrate the two cameras, i.e., estimate their relative position and orientation and the intrinsic parameters?

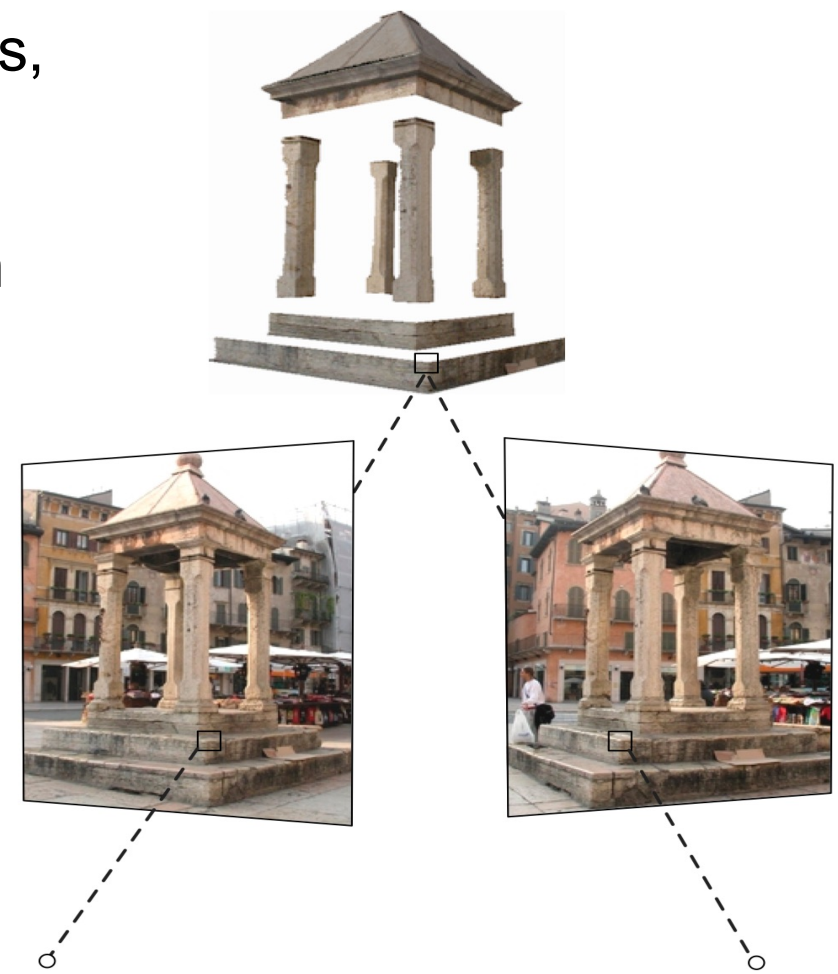

Consider two views of the same 3D scene

- Key idea: we want to answer all these questions without explicit 3D reasoning, by considering the projections of camera centers and visual rays into the other view

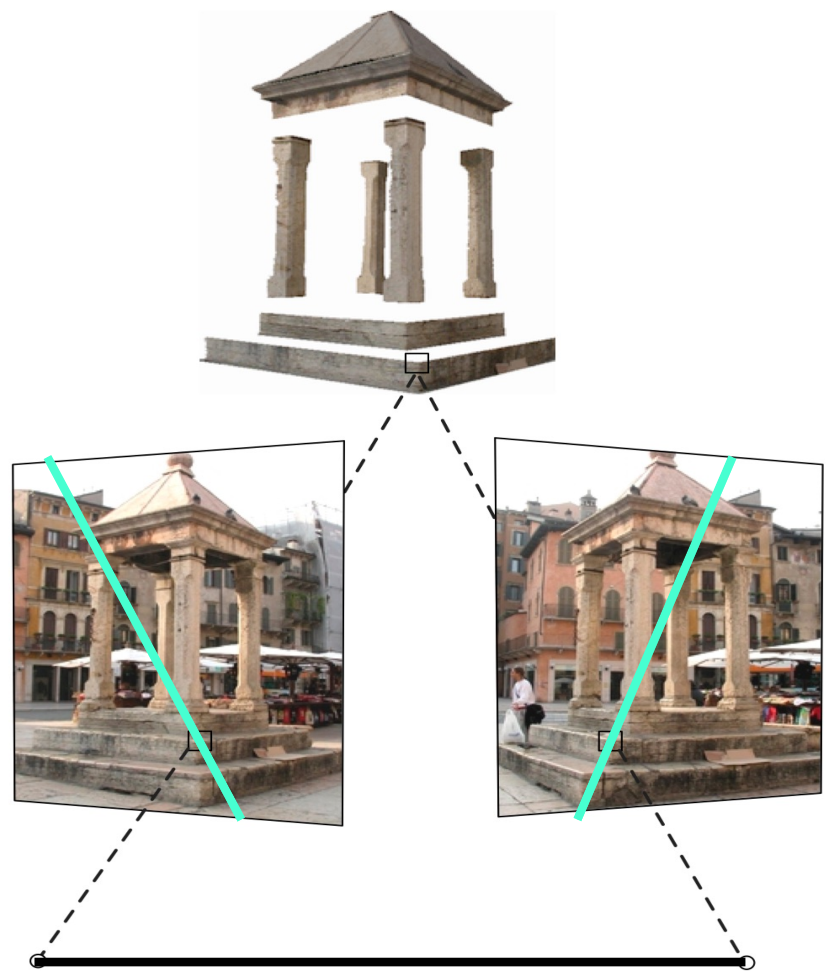

Epipolar geometry setup

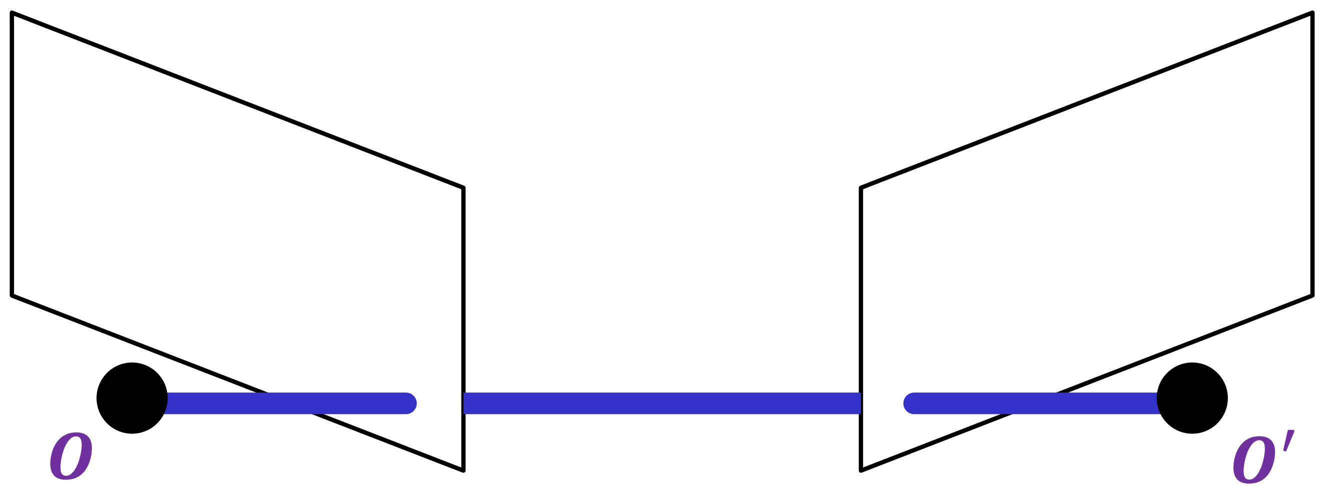

- Suppose we have two cameras with centers O, O'
- The **baseline** is the line connecting the origins

Epipolar geometry setup

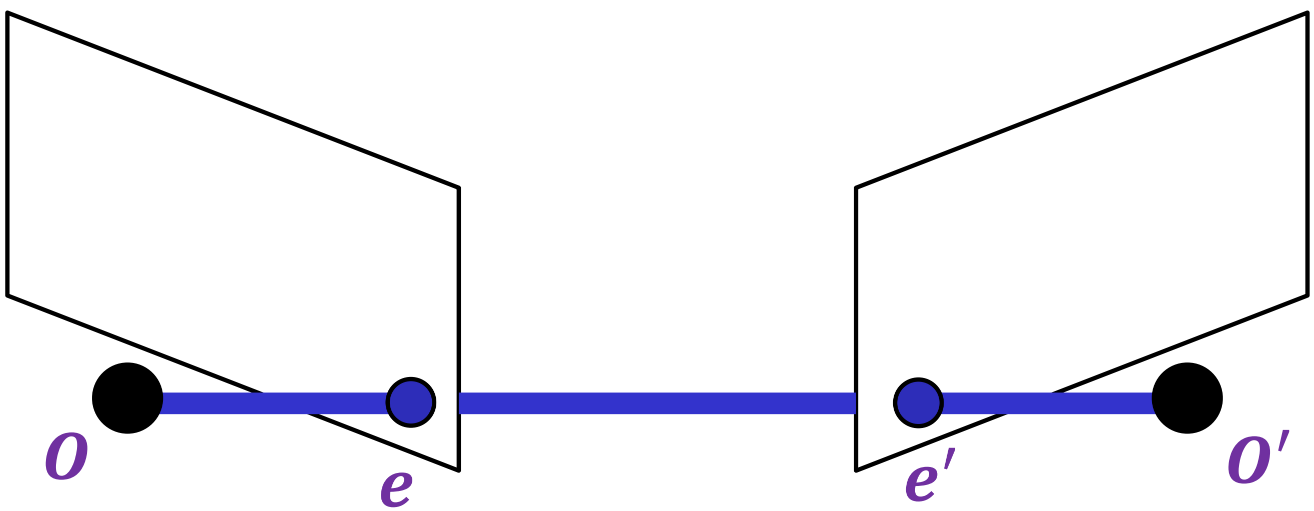

- **Epipoles** e, e' are where the baseline intersects the image planes, or projections of the other camera in each view

Epipolar geometry setup

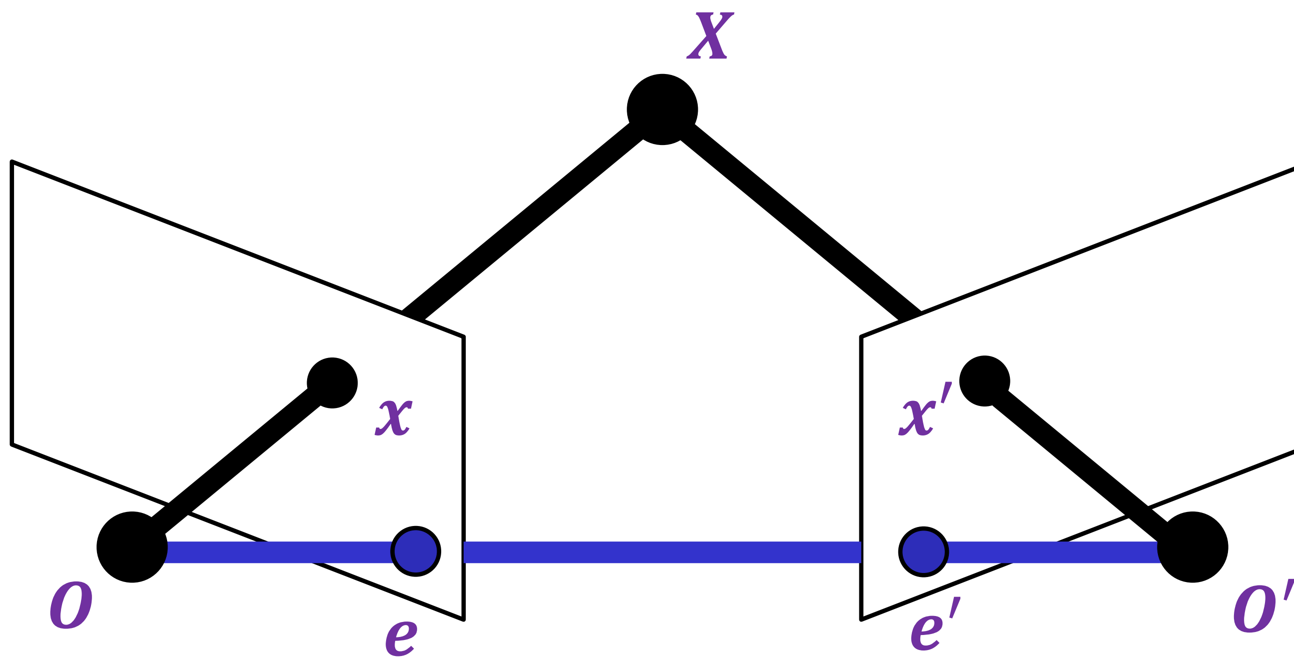

- Consider a **point** X, which projects to x and x'

Epipolar geometry setup

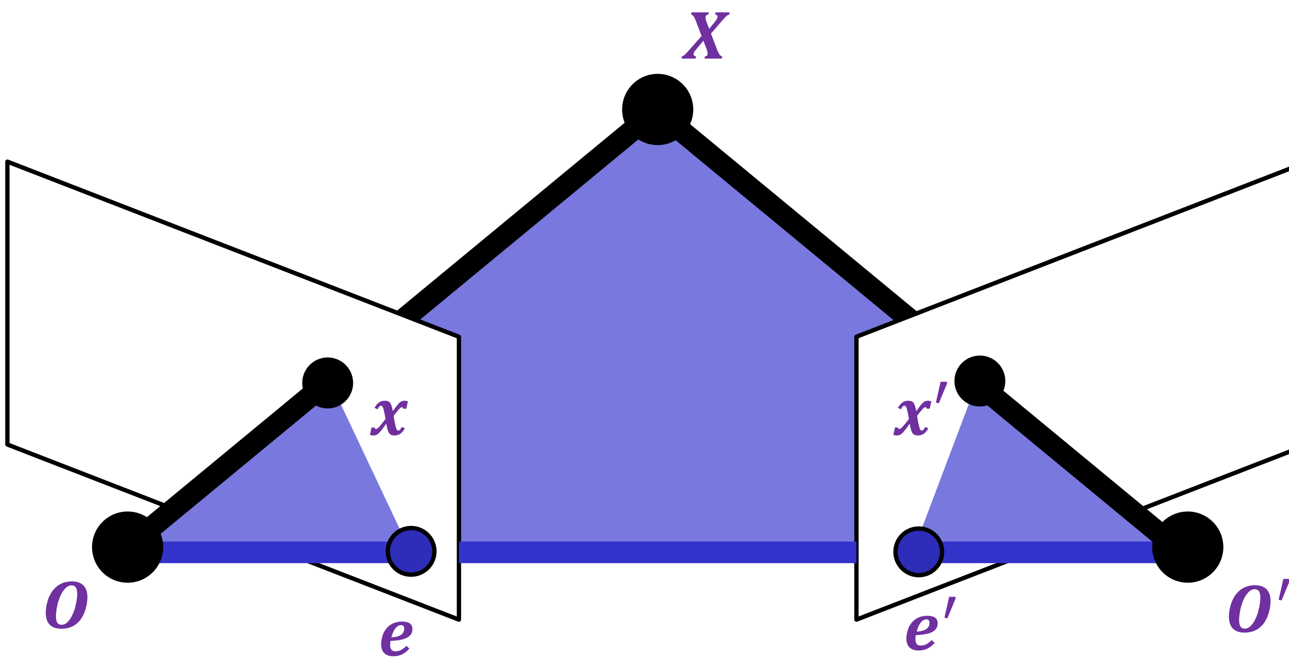

- The plane formed by X, O, and O' is called an **epipolar plane**
- There is a family of planes passing through O and O'

Epipolar geometry setup

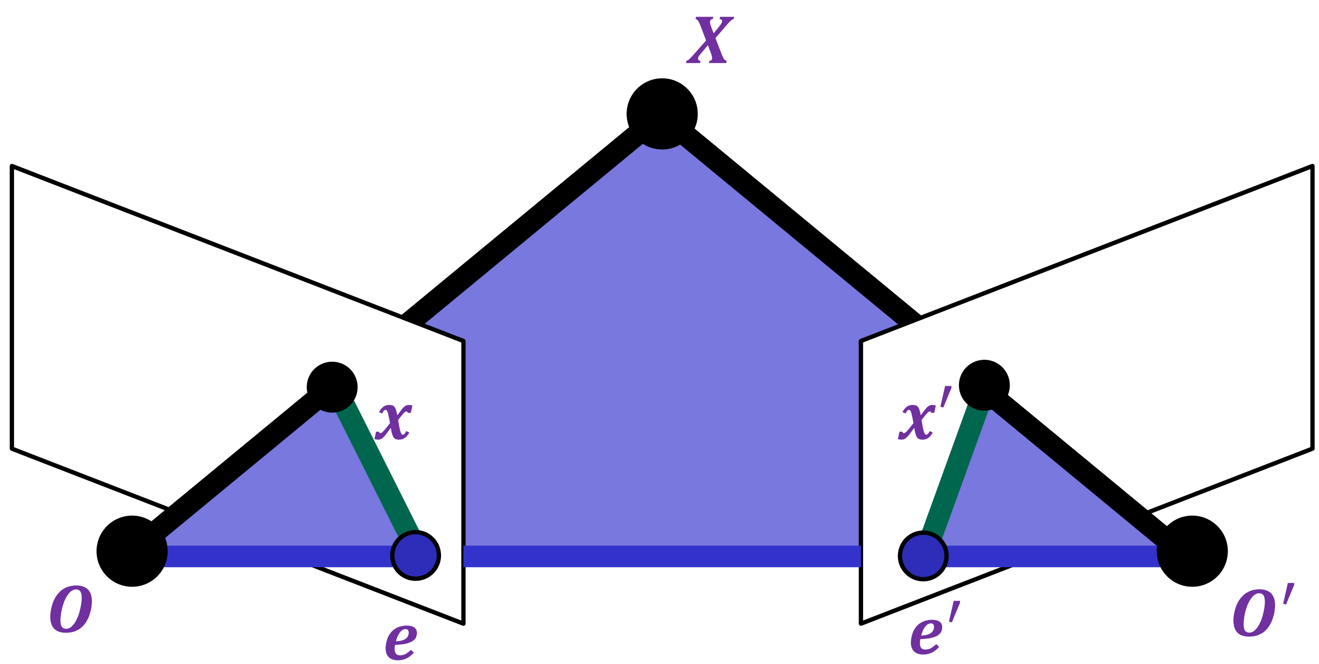

- **Epipolar lines** connect the epipoles to the projections of ***X***
- Equivalently, they are intersections of the epipolar plane with the image planes – thus, they come in matching pairs

Epipolar geometry setup: Summary

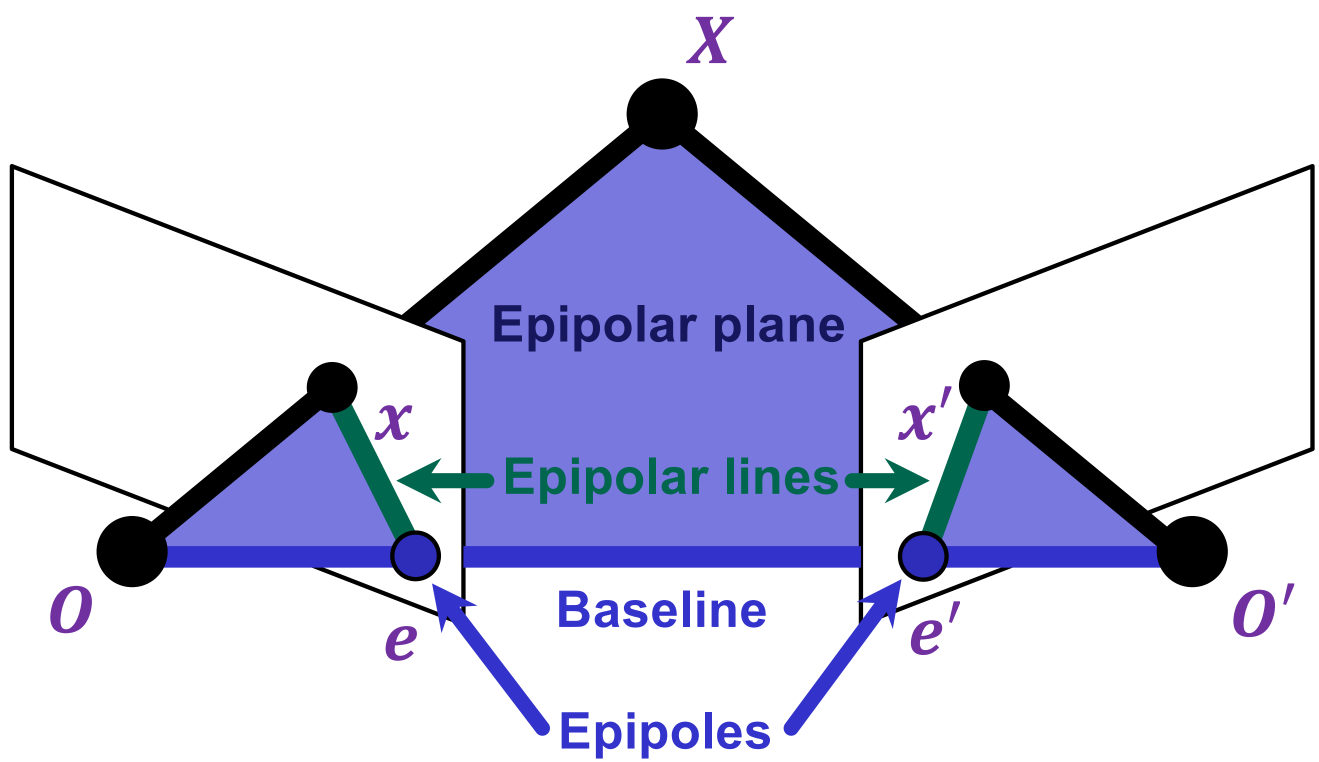

Example configuration: Converging cameras

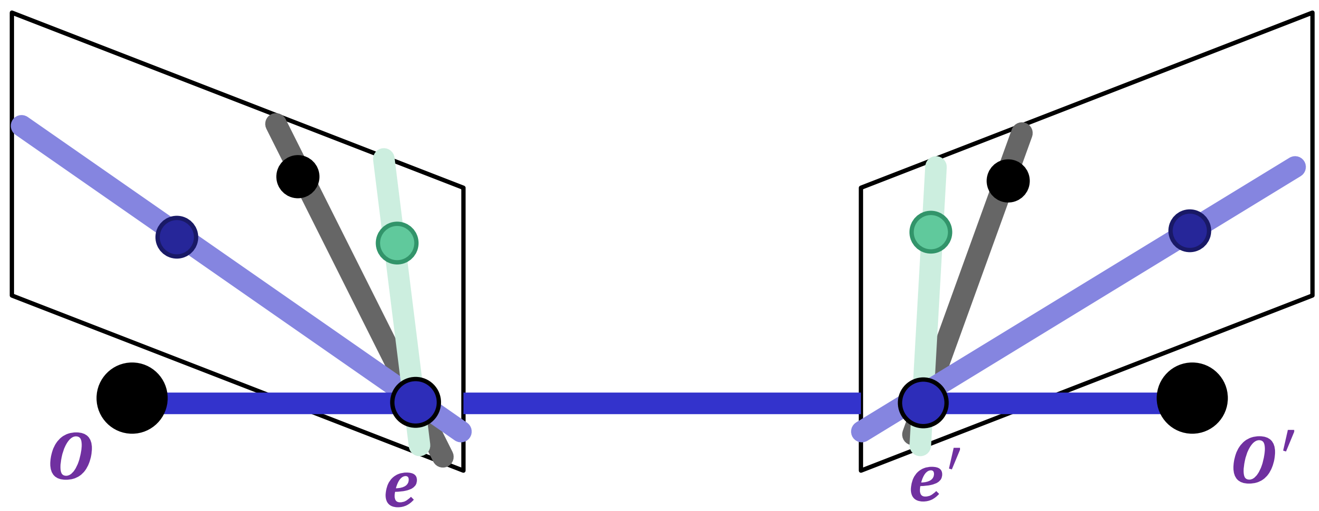

- Epipoles are finite, may be visible in the image

Example configuration: Converging cameras

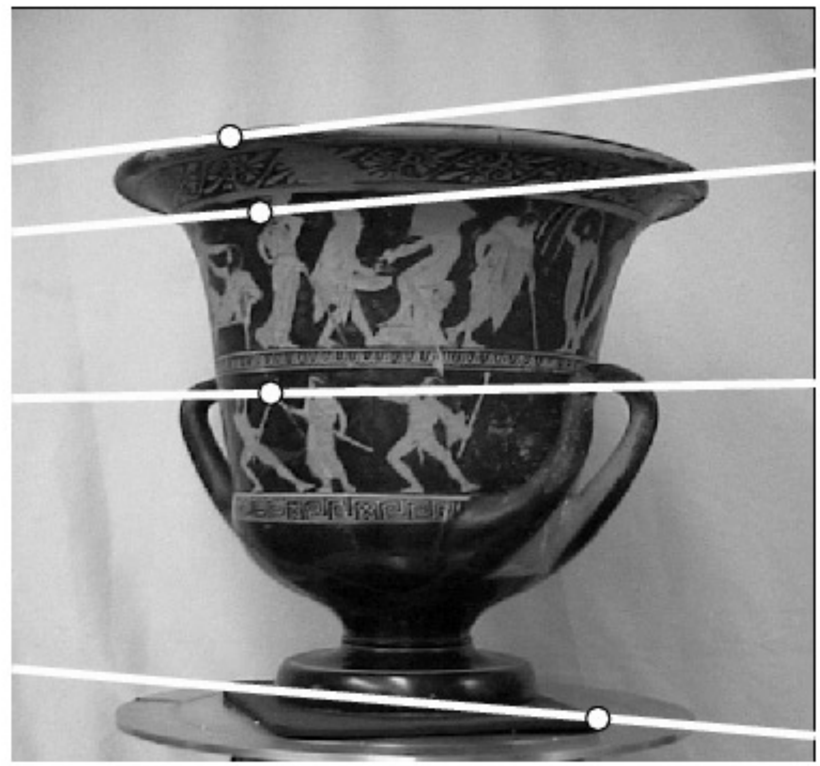

Example configuration: Motion parallel to image plane

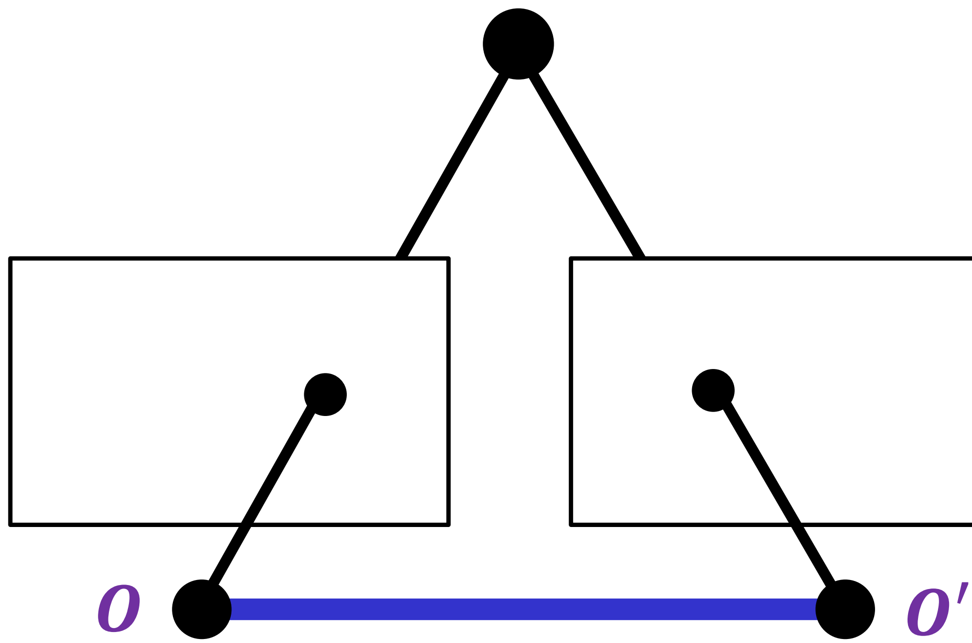

- Where are the epipoles and what do the epipolar lines look like?

Example configuration: Motion parallel to image plane

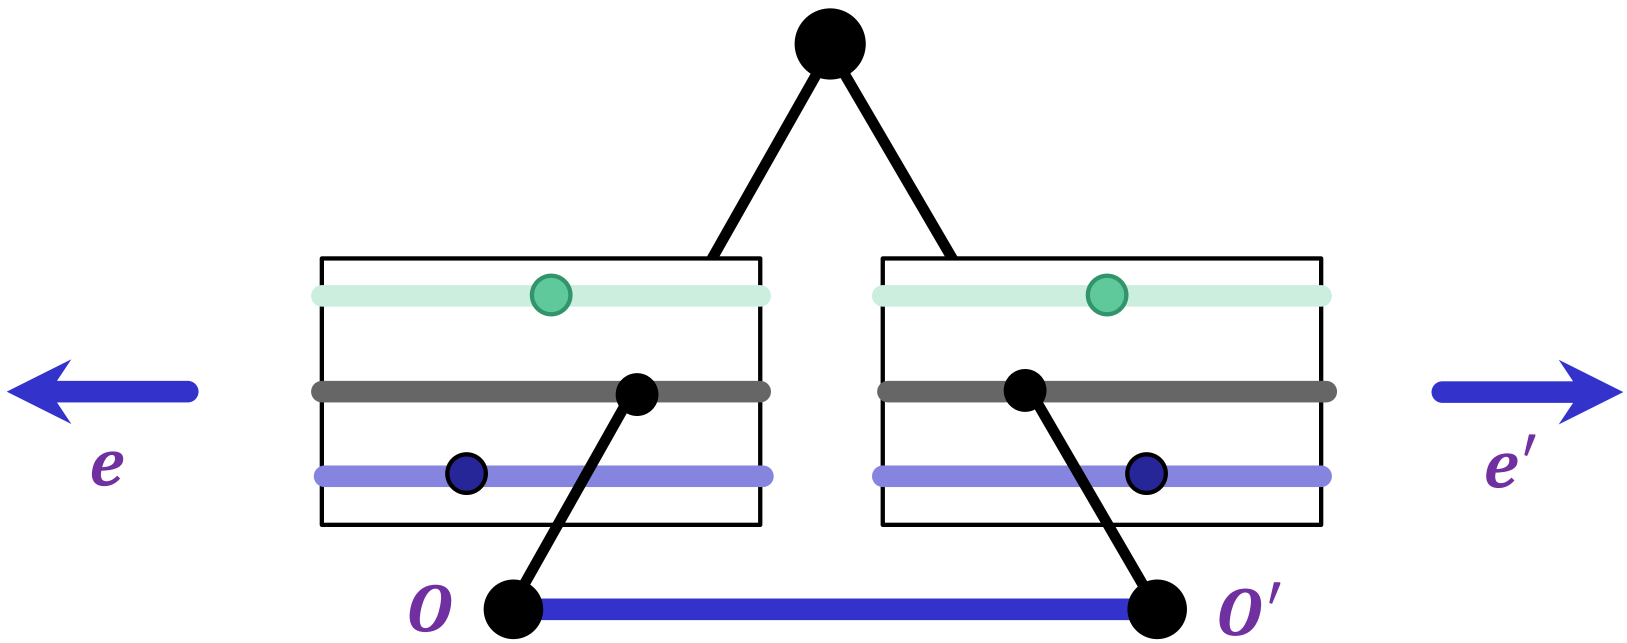

- Epipoles *infinitely* far away, epipolar lines parallel

Example configuration: Motion parallel to image plane

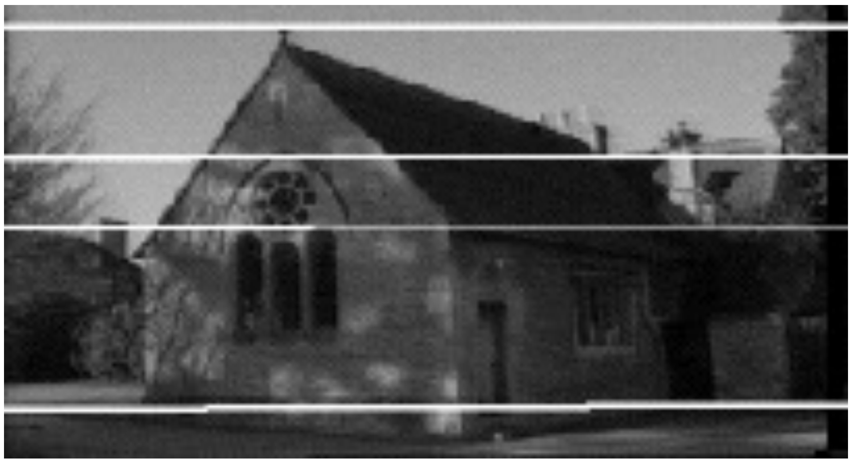

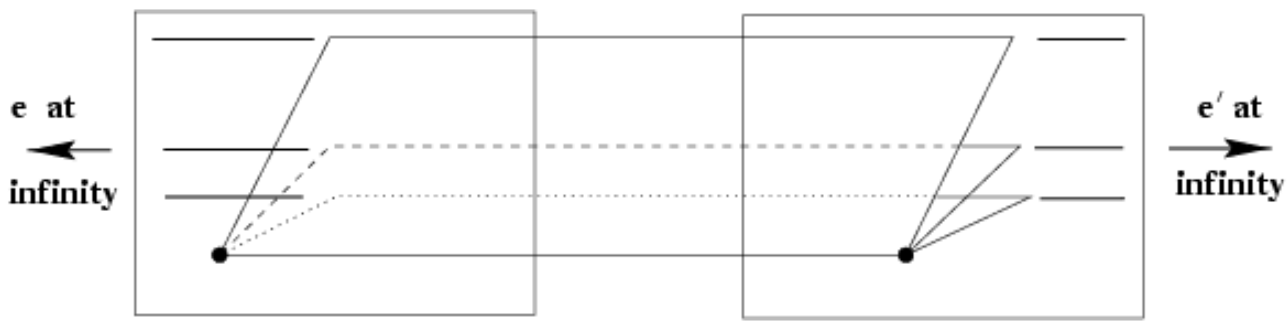

Example configuration: Motion perpendicular to image plane

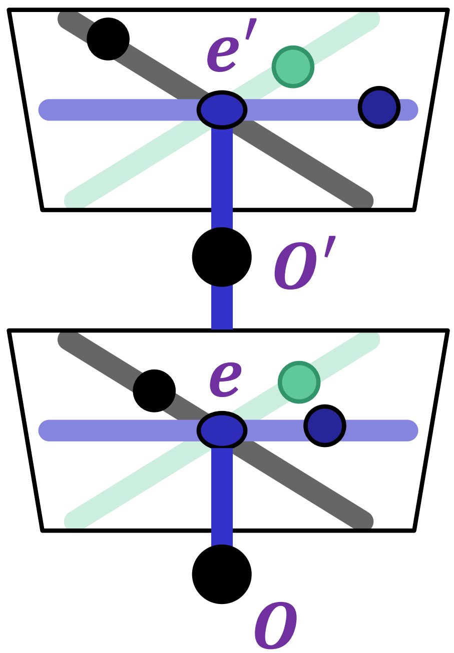

- Epipole is “focus of expansion” and coincides with the principal point of the camera
- Epipolar lines go out from principal point

Example configuration: Motion perpendicular to image plane

Example configuration: Motion perpendicular to image plane

Epipoles (resp. epipolar lines)

- Informative on their own

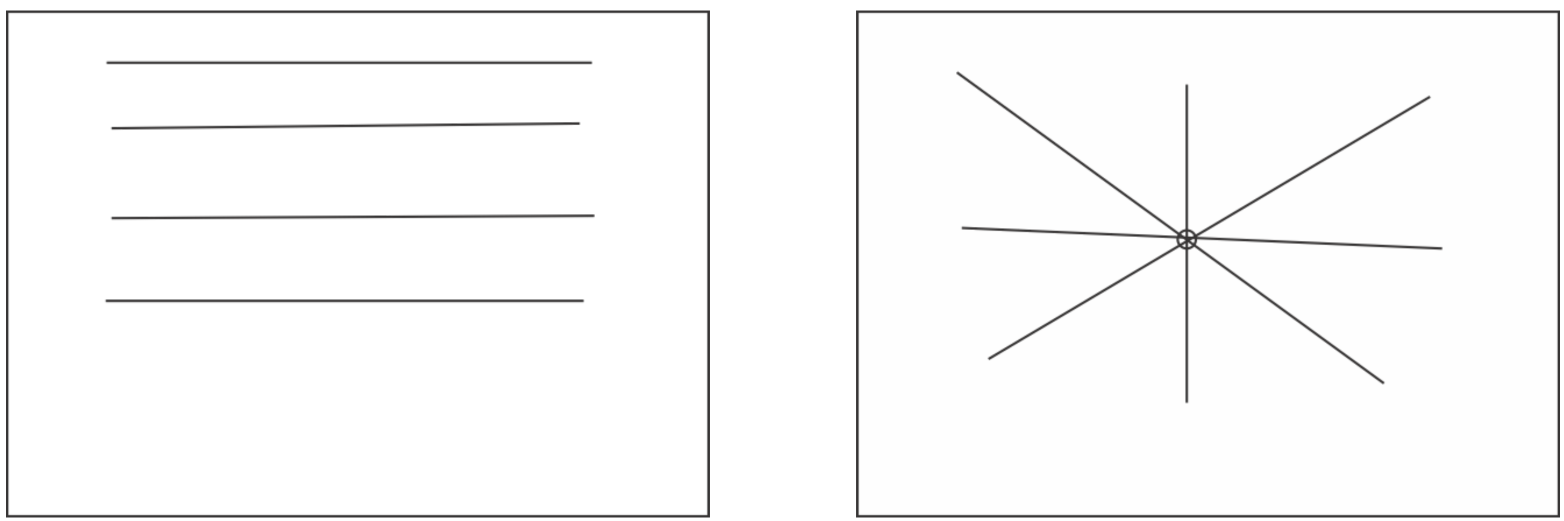

Epipole and epipolar lines in camera 1 - where is camera 2?

Outline

- Motivation
- Epipolar geometry setup
- Epipolar constraint

Epipolar constraint

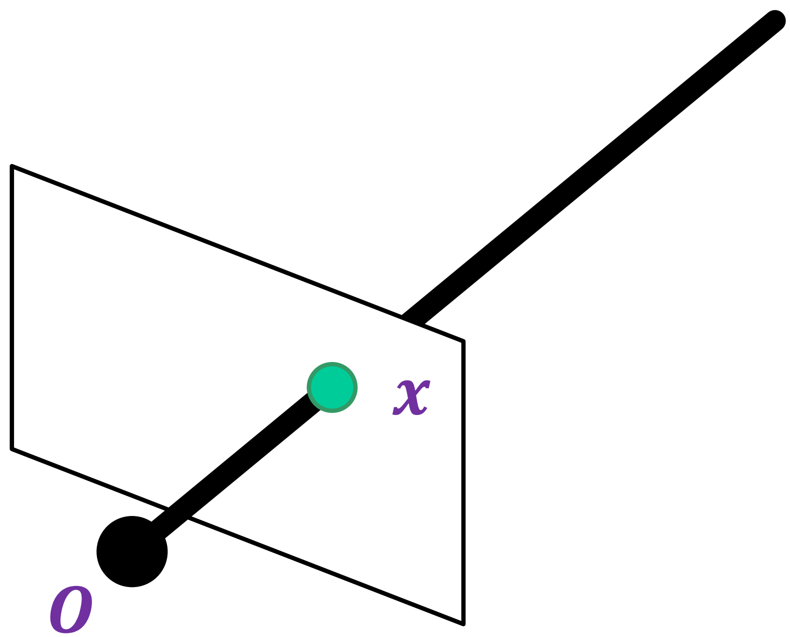

- Suppose we observe a single point x in one image

Epipolar constraint

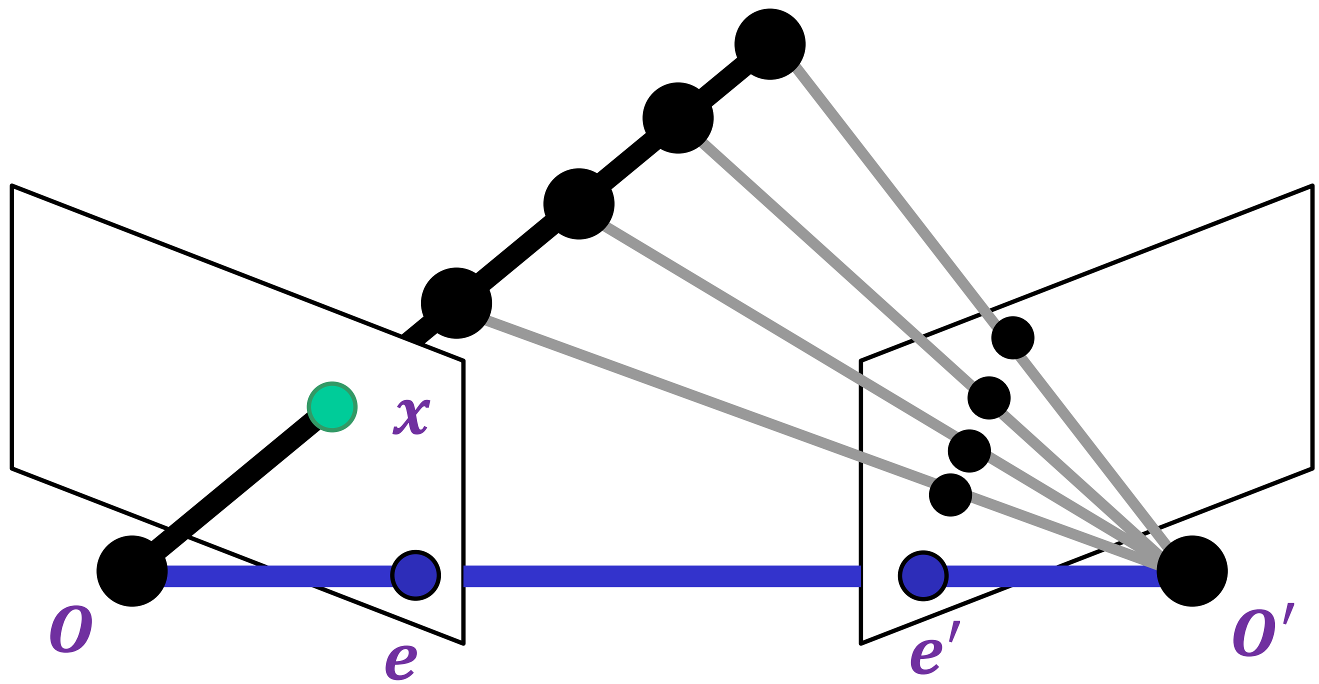

- Where can we find the x' corresponding to x in the other image?

Epipolar constraint

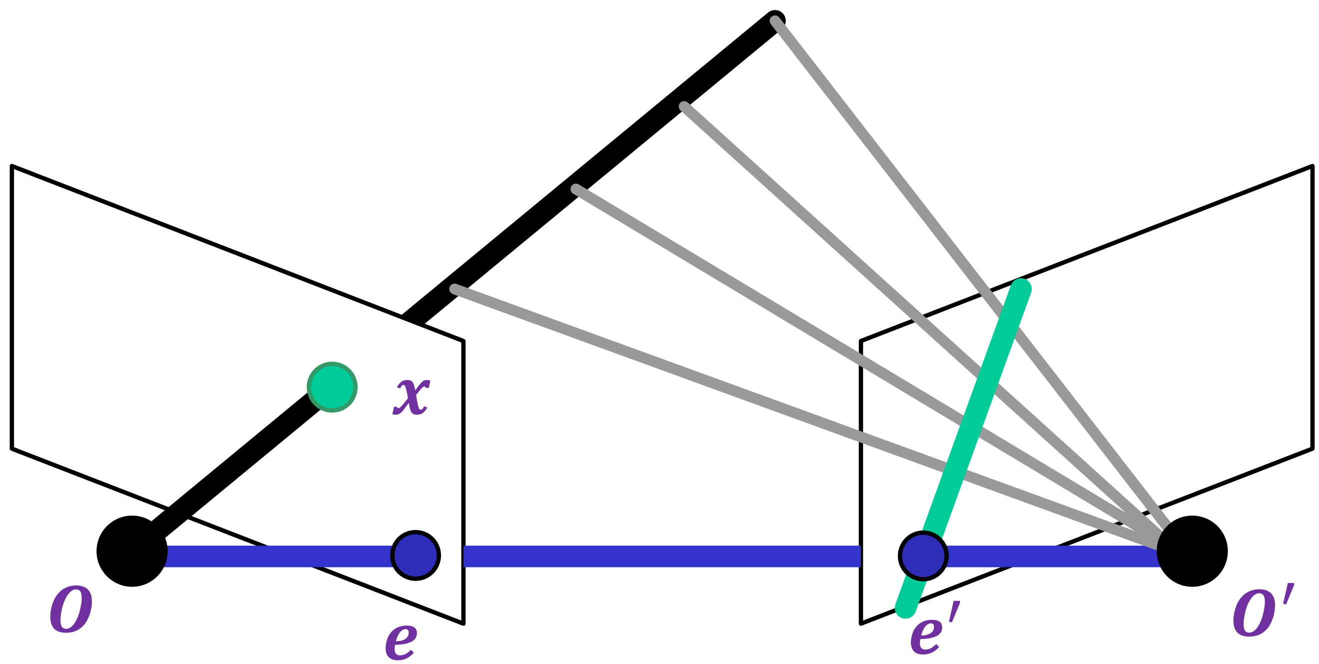

- Where can we find the x' corresponding to x in the other image?
- Along the **epipolar line** corresponding to x (projection of visual ray connecting O with x into the second image plane)

Epipolar constraint

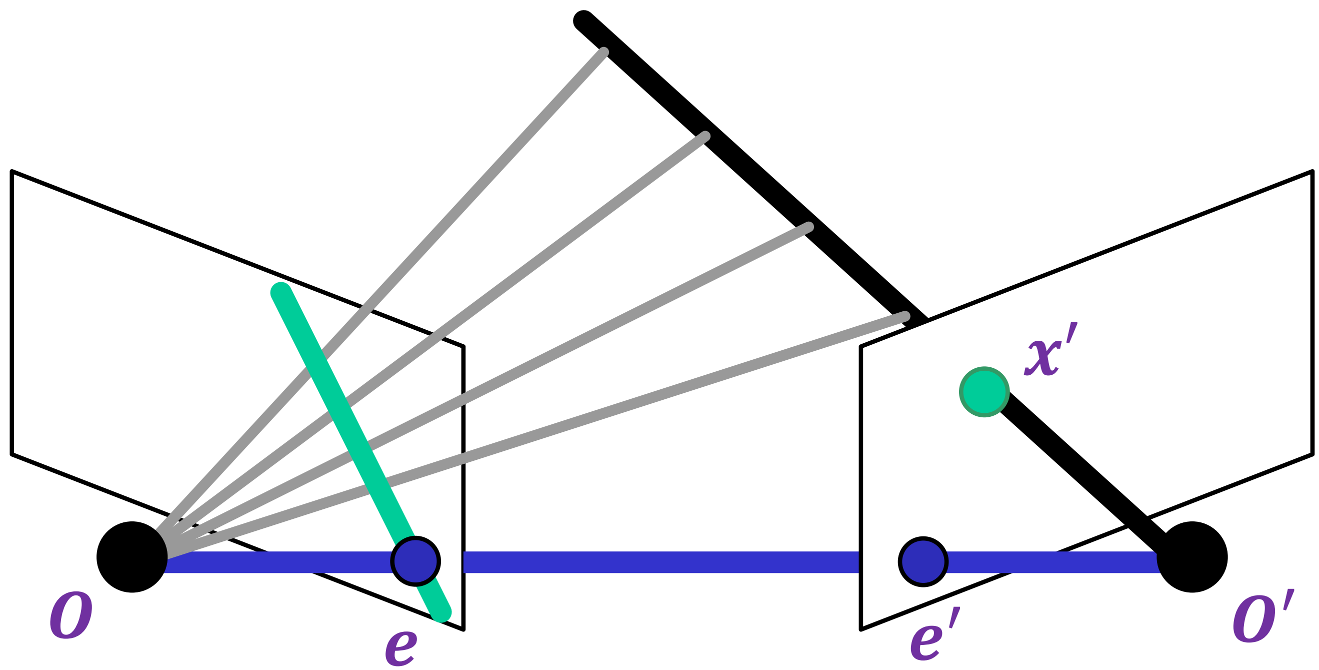

- Similarly, all points in the left image corresponding to x' have to lie along the epipolar line corresponding to x'

Epipolar constraint

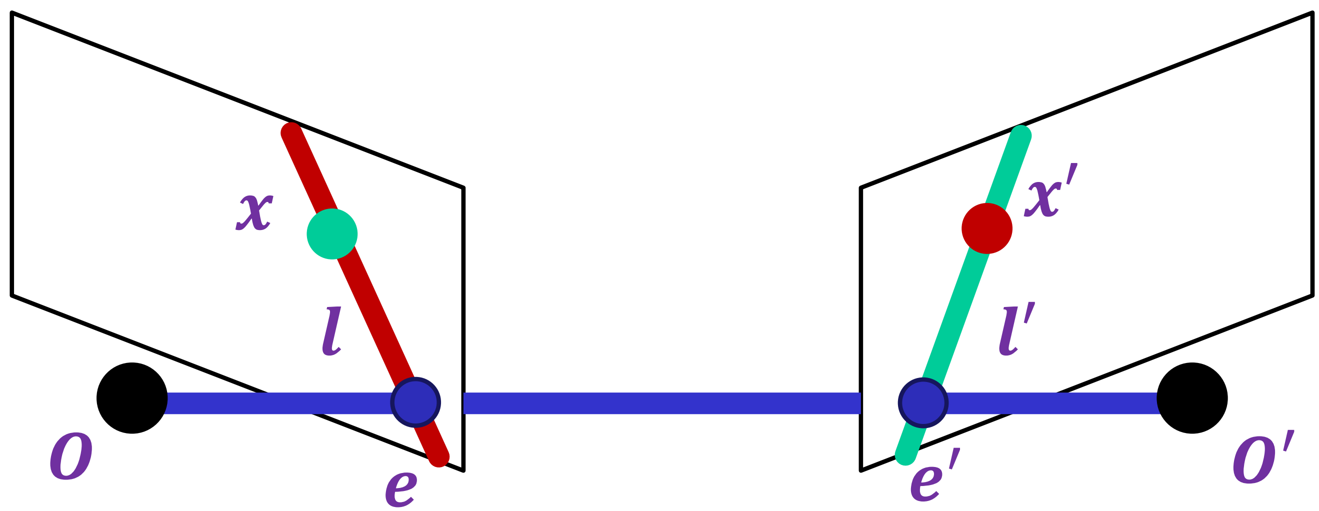

- Potential matches for x have to lie on the matching epipolar line l'
- Potential matches for x' have to lie on the matching epipolar line l

Epipolar constraint: Example

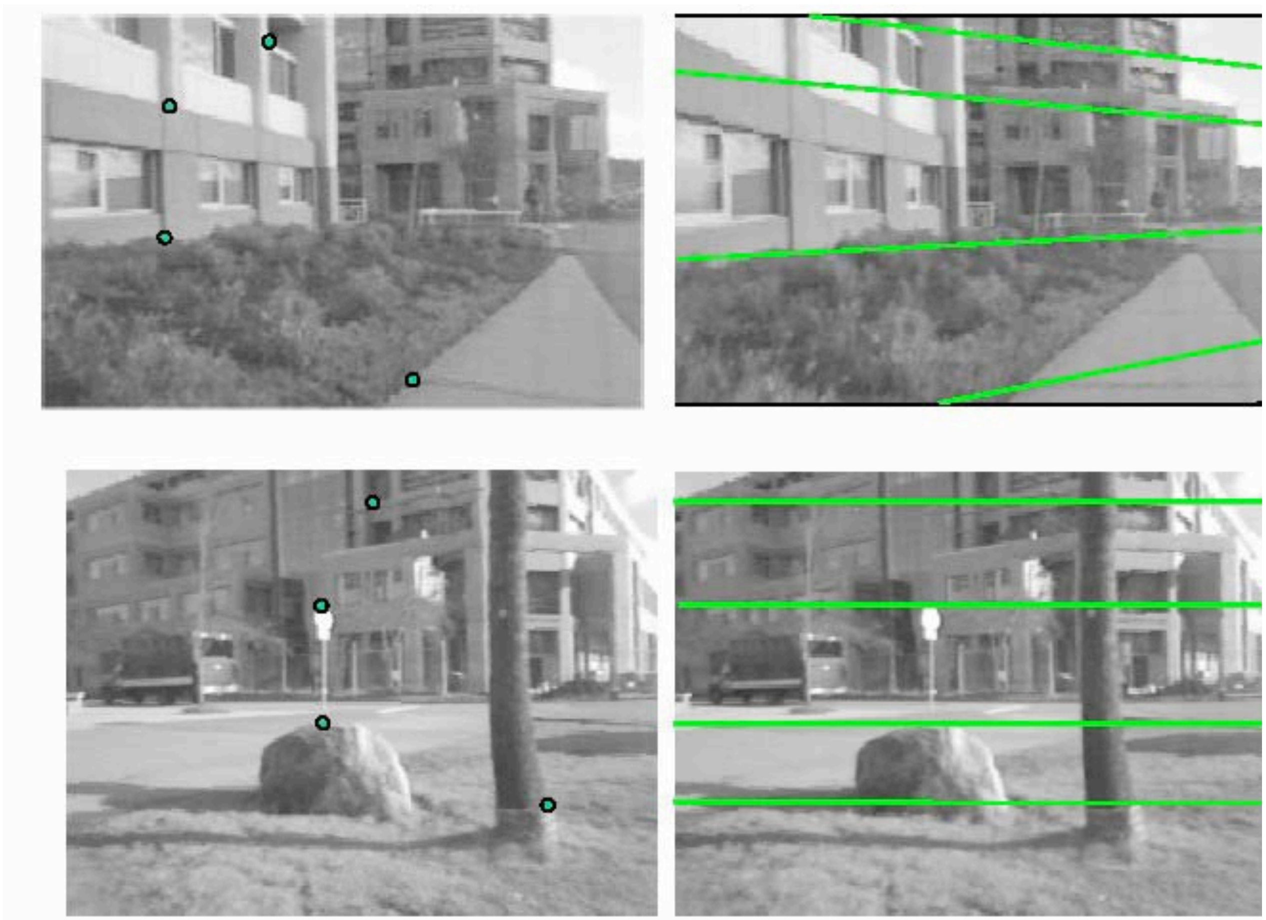

Epipolar constraint

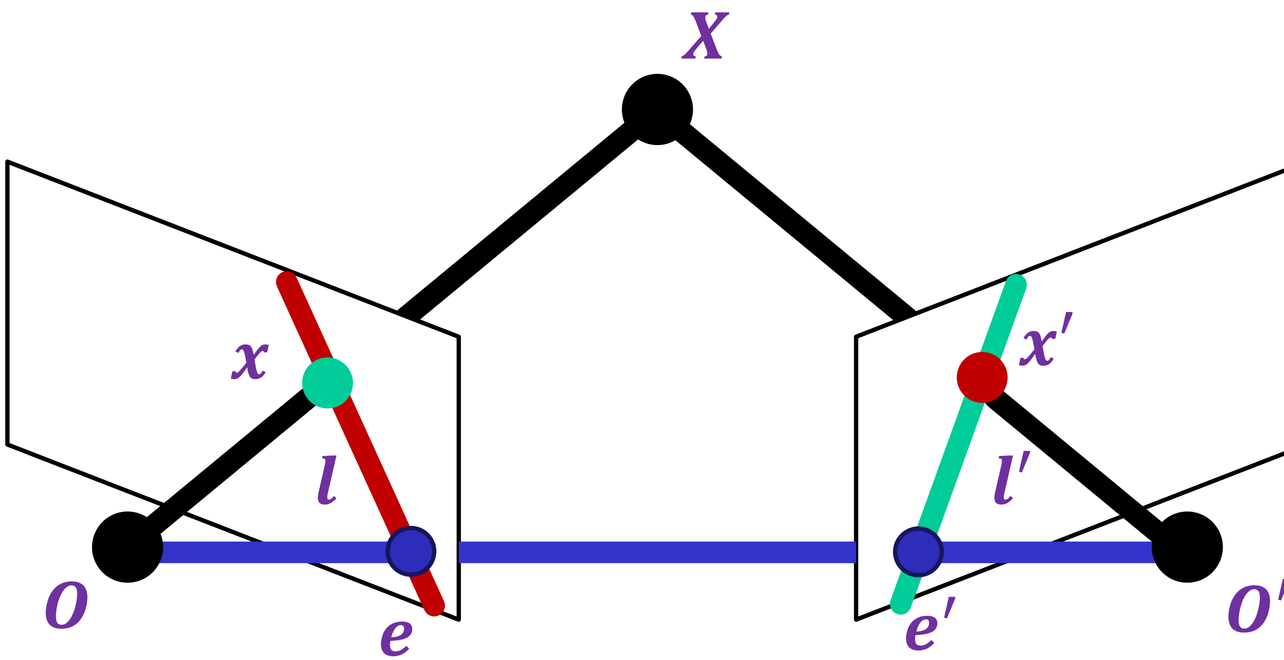

- Whenever two points x and x' lie on matching epipolar lines l and l', the visual rays corresponding to them meet in space, i.e., x and x' could be projections of the same 3D point X

Epipolar constraint

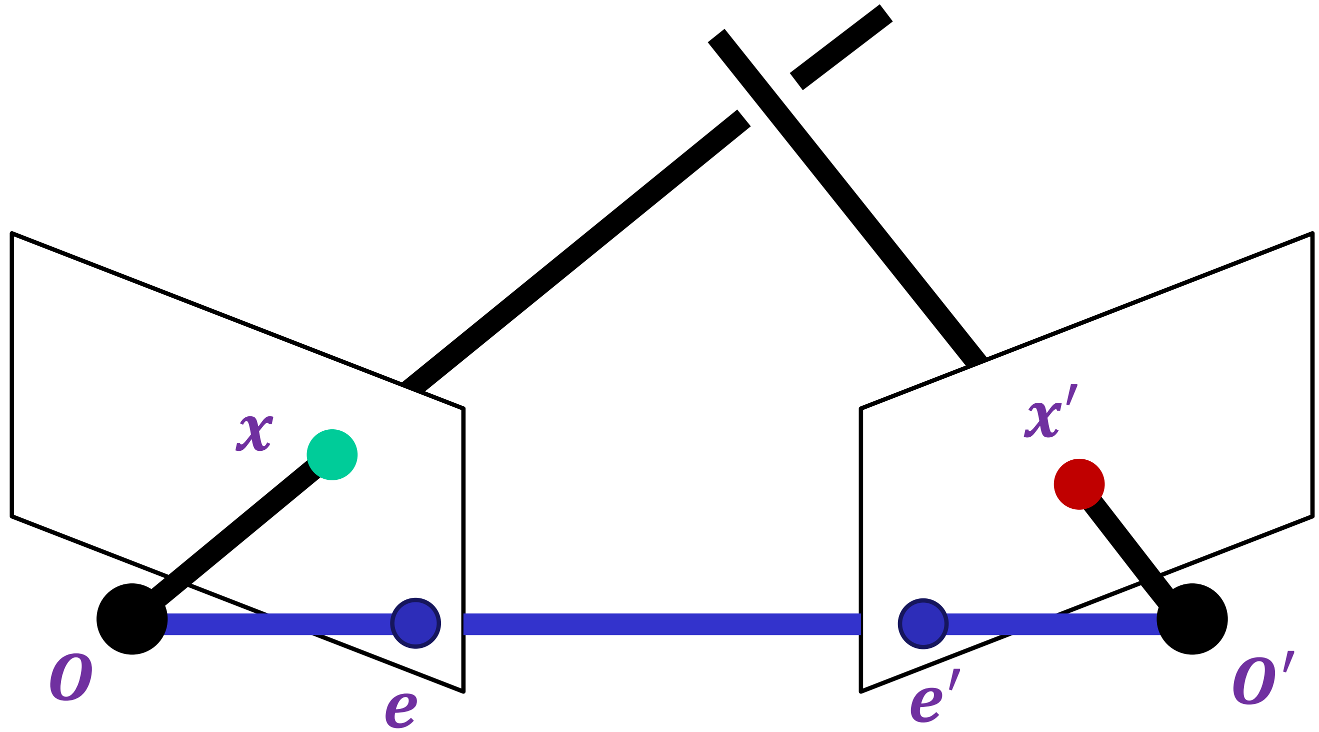

- Remember: in general, two rays *do not* meet in space!

Epipolar constraint

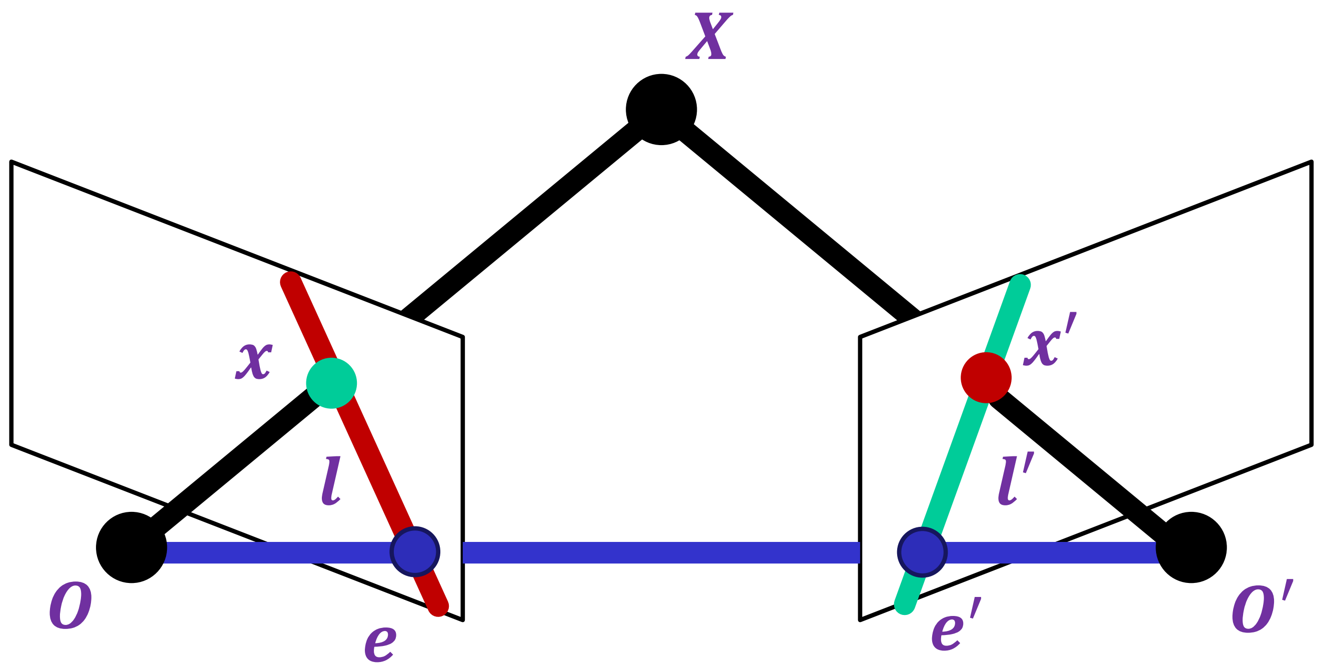

- X, x, x’, O, O’ are coplanar
- Know O, O’
  - Choose x – This yields the plane and so l’ and x’ lies on l’
    - So there is some vector function f(x; O, O’) such that
    - f(x; O, O’)^T x’=0

Epipolar constraint

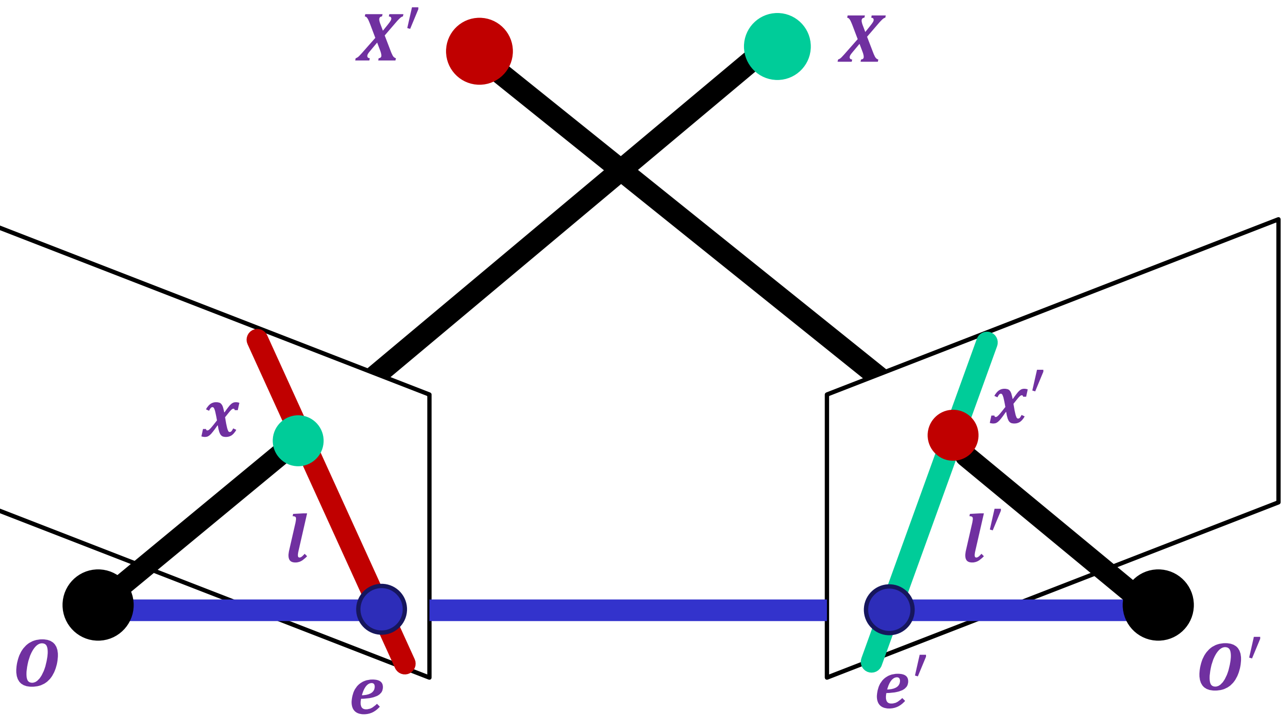

- Caveat: if x and x' satisfy the epipolar constraint, this doesn't mean they *have to be* projections of the same 3D point

Outline

- Motivation
- Epipolar geometry setup
- Epipolar constraint
- Essential matrix

Math of the epipolar constraint: Calibrated case

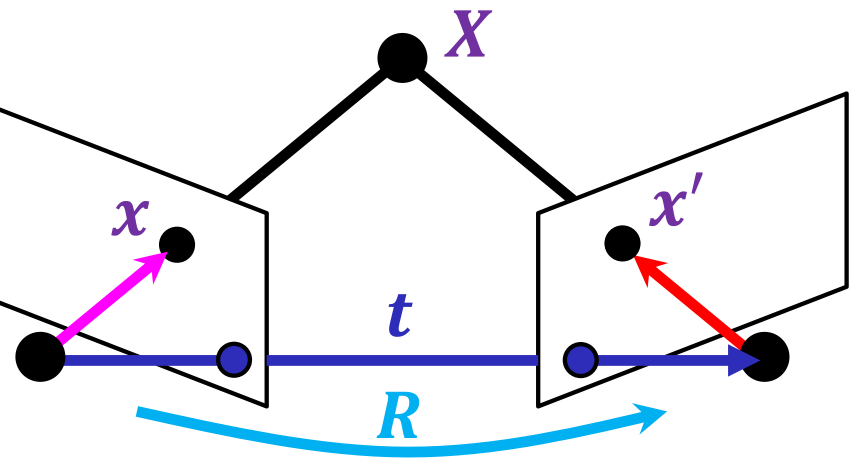

- Assume the intrinsic and extrinsic parameters of the cameras are known, world coordinate system is set to that of the first camera
- Then the projection matrices are given by $\boldsymbol{K}[\boldsymbol{I} \mid \boldsymbol{0}]$ and $\boldsymbol{K}'[\boldsymbol{R} \mid \boldsymbol{t}]$
- We can pre-multiply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

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

Math of the epipolar constraint: Calibrated case

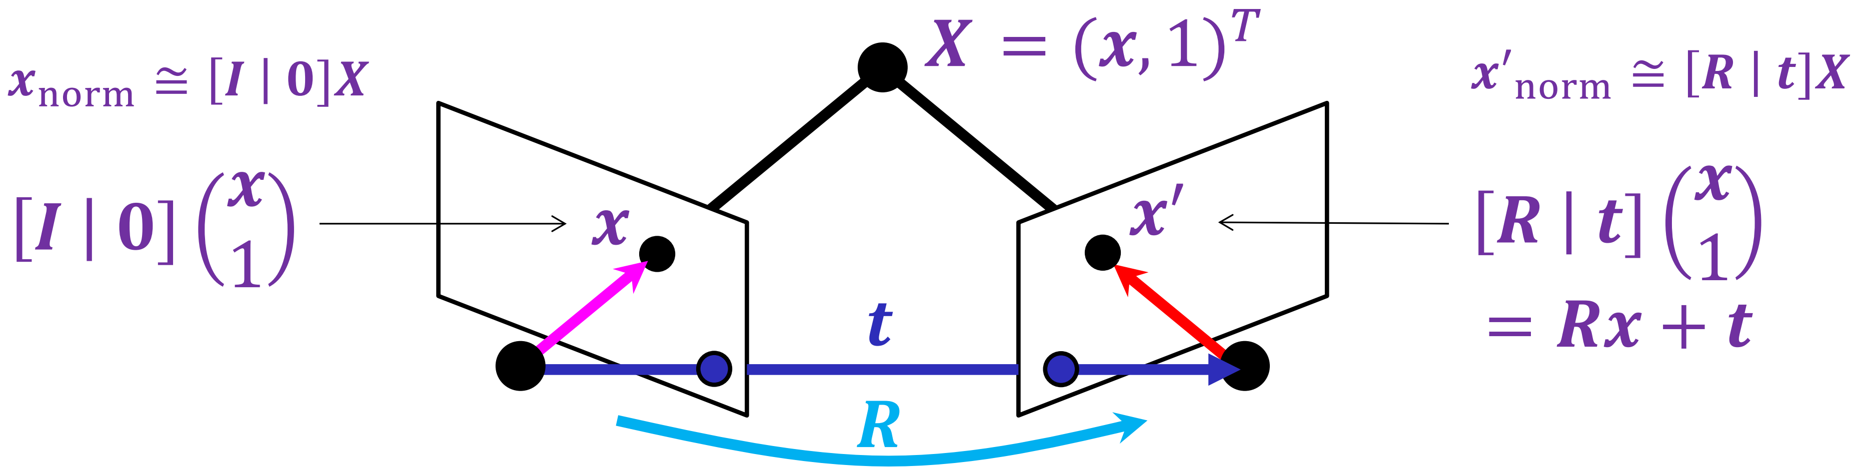

- We have $\boldsymbol{x}' \cong \boldsymbol{R}\boldsymbol{x} + \boldsymbol{t}$
- This means the three vectors $\boldsymbol{x}'$, $\boldsymbol{R}\boldsymbol{x}$, and $\boldsymbol{t}$ are linearly dependent
- This constraint can be written using the *triple product*

$$\boldsymbol{x}' \cdot [\boldsymbol{t} \times (\boldsymbol{R}\boldsymbol{x})] = 0$$

Math of the epipolar constraint: Calibrated case

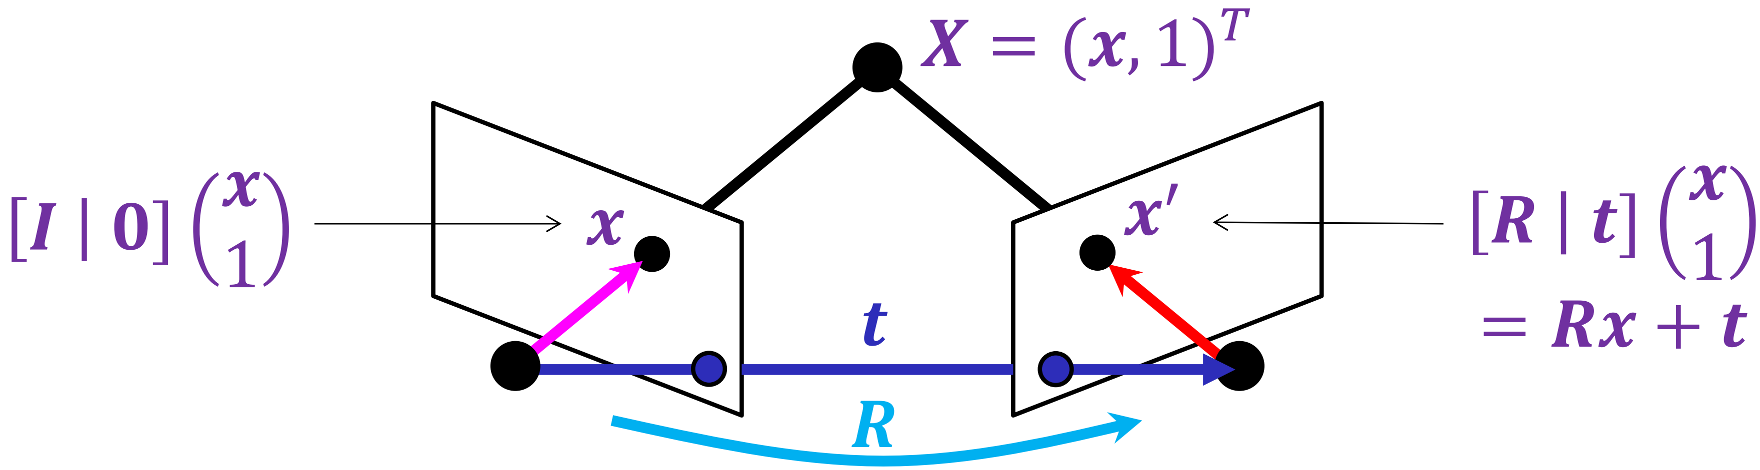

$$x' \cdot [t \times (Rx)] = 0 \quad \Rightarrow \quad x'^T [t_\times] R x = 0$$

Recall: $a \times b = \begin{bmatrix} 0 & -a_3 & a_2 \\ a_3 & 0 & -a_1 \\ -a_2 & a_1 & 0 \end{bmatrix} \begin{pmatrix} b_1 \\ b_2 \\ b_3 \end{pmatrix} = [a_\times] b$

Math of the epipolar constraint: Calibrated case

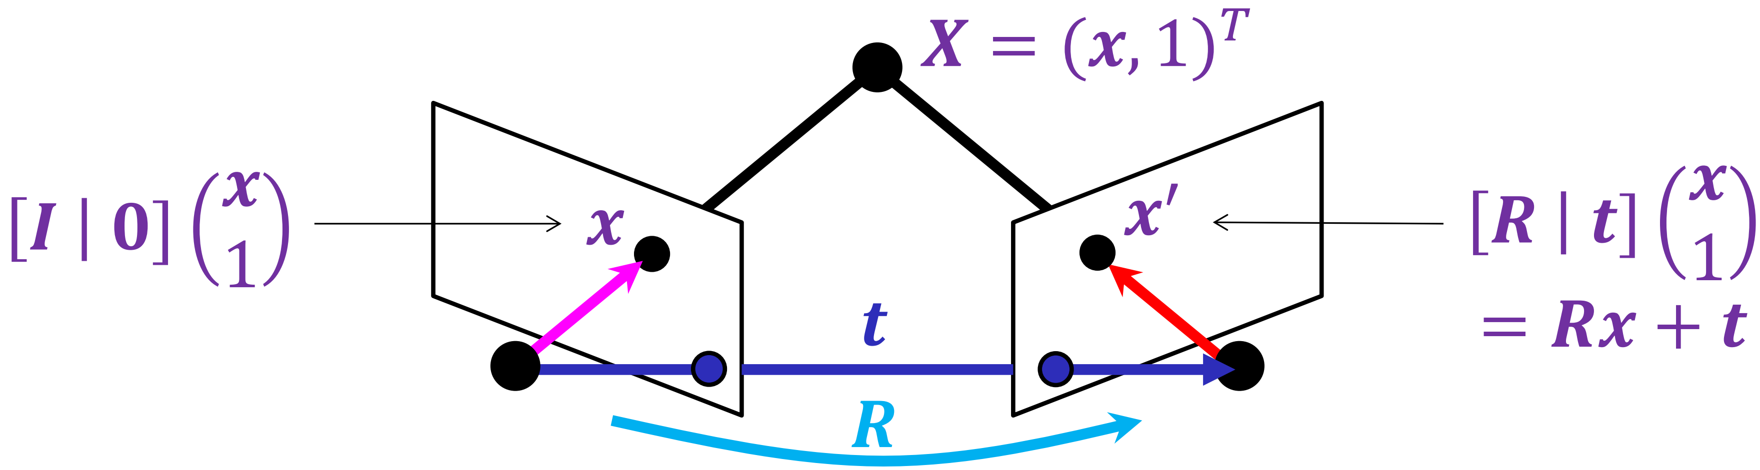

$\mathbf{x}' \cdot [\mathbf{t} \times (\mathbf{R}\mathbf{x})] = 0 \quad \Rightarrow \quad \mathbf{x}'^T [\mathbf{t}_\times] \mathbf{R}\mathbf{x} = 0 \quad \Rightarrow \quad \mathbf{x}'^T \mathbf{E}\mathbf{x} = 0$

Essential Matrix

H. C. Longuet-Higgins. A computer algorithm for reconstructing a scene from two projections. Nature 293 (5828): 133–135, September 1981

The essential matrix

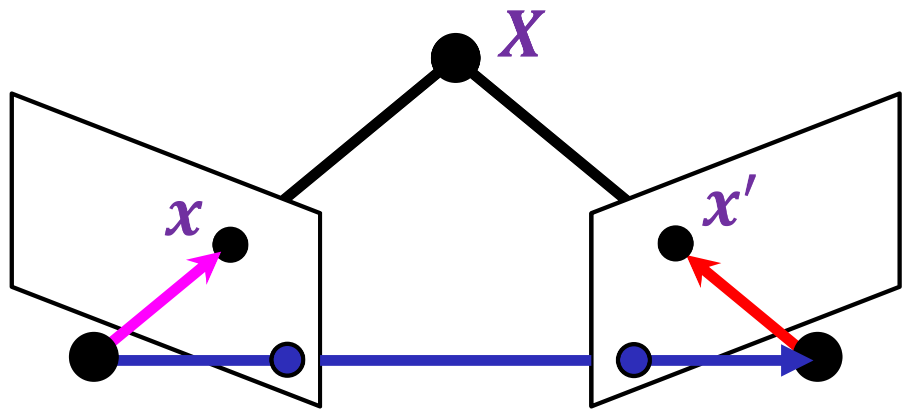

$$\boldsymbol{x}'^T \boldsymbol{E} \boldsymbol{x} = 0$$

$$(x', y', 1) \begin{bmatrix} e_{11} & e_{12} & e_{13} \\ e_{21} & e_{22} & e_{23} \\ e_{31} & e_{32} & e_{33} \end{bmatrix} \begin{pmatrix} x \\ y \\ 1 \end{pmatrix} = 0$$

The essential matrix: Properties

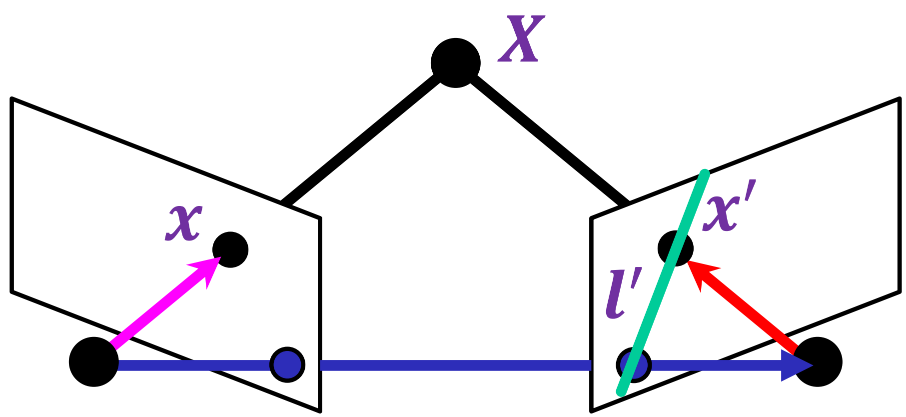

$$\boldsymbol{x}'^T \boldsymbol{E} \boldsymbol{x} = 0$$

- $\boldsymbol{E}\boldsymbol{x}$ is the epipolar line associated with $\boldsymbol{x}$ ($\boldsymbol{l}' = \boldsymbol{E}\boldsymbol{x}$)

  Recall: a line is given by $ax + by + c = 0$ or $\boldsymbol{l}^T \boldsymbol{x} = 0$
  where $\boldsymbol{l} = (a, b, c)^T$ and $\boldsymbol{x} = (x, y, 1)^T$

The essential matrix: Properties

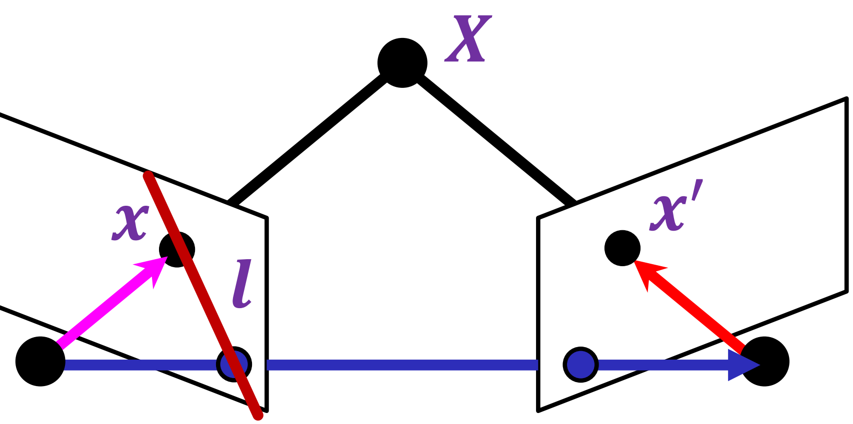

$$x'^T E x = 0$$

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

Outline

- Motivation
- Epipolar geometry setup
- Epipolar constraint
- Essential matrix
- Fundamental matrix

Epipolar constraint: Uncalibrated case

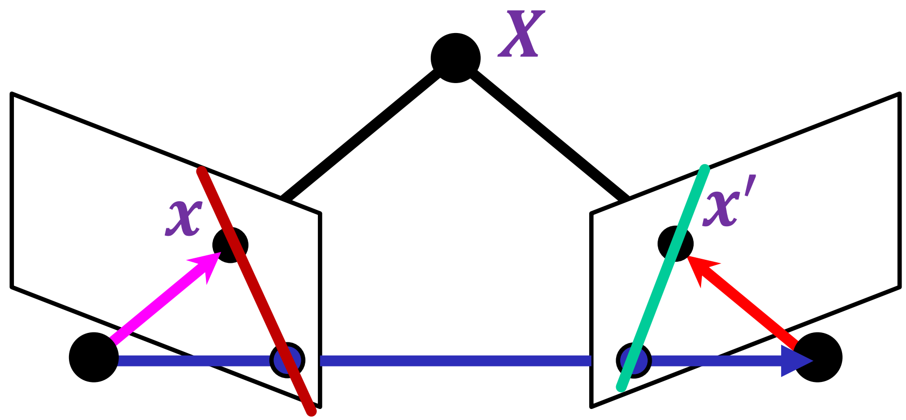

- The calibration matrices $\boldsymbol{K}$ and $\boldsymbol{K}'$ of the two cameras are unknown
- We can write the epipolar constraint in terms of *unknown* normalized coordinates:

$$\boldsymbol{x}'^T_{\text{norm}} \boldsymbol{E} \boldsymbol{x}_{\text{norm}} = 0,$$

where $\boldsymbol{x}_{\text{norm}} = \boldsymbol{K}^{-1}\boldsymbol{x}$, $\boldsymbol{x}'_{\text{norm}} = \boldsymbol{K}'^{-1}\boldsymbol{x}'$

Epipolar constraint: Uncalibrated case

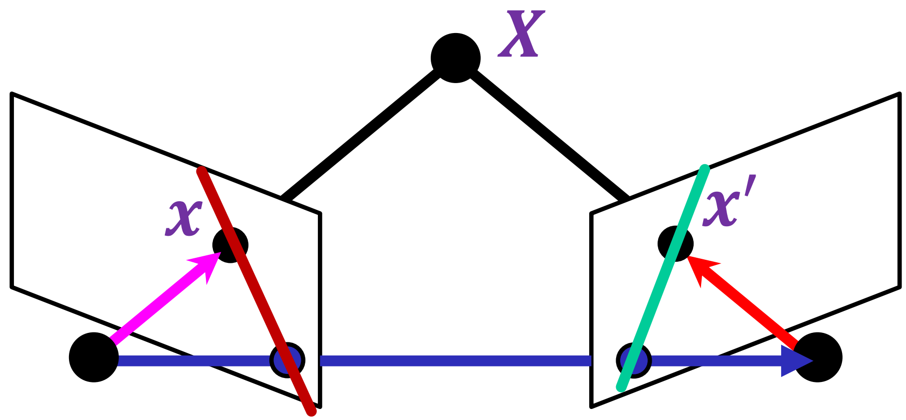

$$\boldsymbol{x}'^{T}_{\text{norm}}\boldsymbol{E}\boldsymbol{x}_{\text{norm}} = 0 \quad \Rightarrow \quad \boldsymbol{x}'^{T}\boldsymbol{F}\boldsymbol{x} = 0, \text{ where } \boldsymbol{F} = \boldsymbol{K}'^{-T}\boldsymbol{E}\boldsymbol{K}^{-1}$$

$$\boldsymbol{x}_{\text{norm}} = \boldsymbol{K}^{-1}\boldsymbol{x}$$

$$\boldsymbol{x}'_{\text{norm}} = \boldsymbol{K}'^{-1}\boldsymbol{x}'$$

Fundamental Matrix

Faugeras et al., (1992), Hartley (1992)

The fundamental matrix

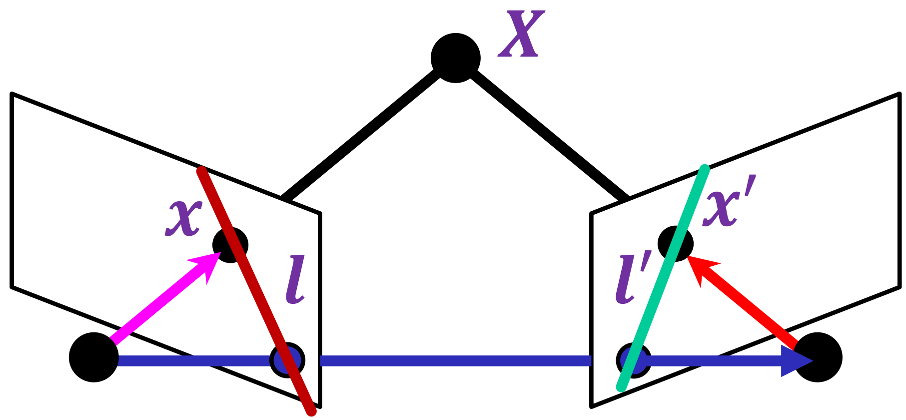

$$\mathbf{x}'^T \mathbf{F} \mathbf{x} = 0$$

$$(x', y', 1) \begin{bmatrix} f_{11} & f_{12} & f_{13} \\ f_{21} & f_{22} & f_{23} \\ f_{31} & f_{32} & f_{33} \end{bmatrix} \begin{pmatrix} x \\ y \\ 1 \end{pmatrix} = 0$$

The fundamental matrix: Properties

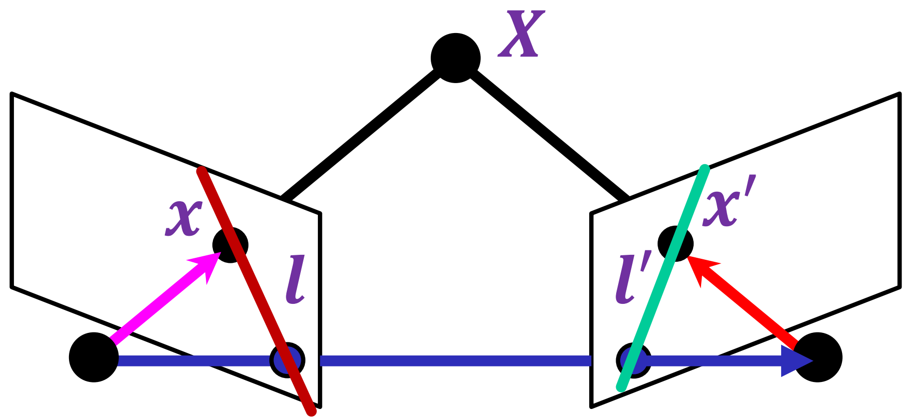

$$\boldsymbol{x}'^T \boldsymbol{F} \boldsymbol{x} = 0$$

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

The fundamental matrix: Properties

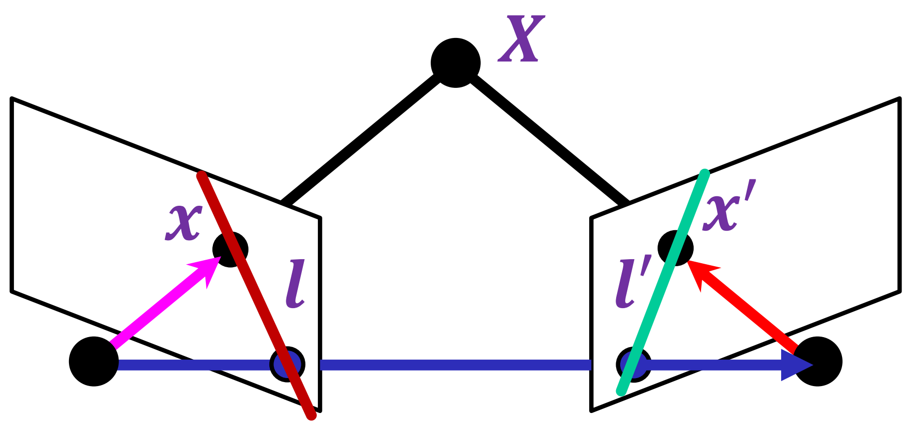

$$x'^T F x = 0$$

- F is singular (rank two) and has seven degrees of freedom
  - Why singular?
    - F maps any point on one side to a line in through epipole on other
      - NOT to a general line
      - So 2DOF pt -> 1DOF family of lines

Outline

- Motivation
- Epipolar geometry setup
- Epipolar constraint
- Essential matrix
- Fundamental matrix
- Estimating the fundamental matrix

Estimating the fundamental matrix

- Given: correspondences $\boldsymbol{x} = (x, y, 1)^T$ and $\boldsymbol{x}' = (x', y', 1)^T$

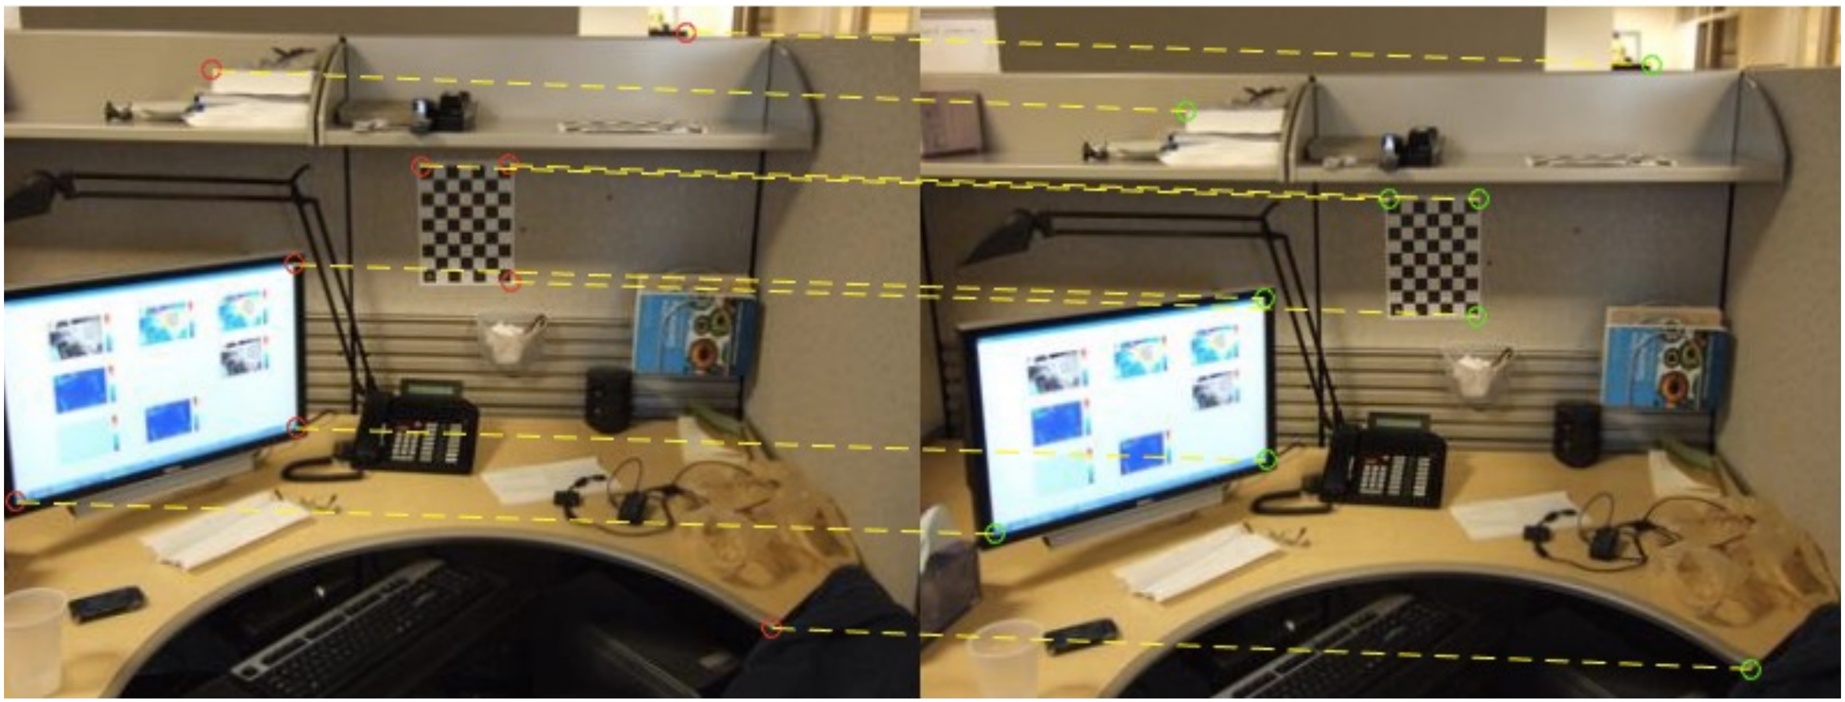

Estimating the fundamental matrix

- Given: correspondences $\boldsymbol{x} = (x, y, 1)^T$ and $\boldsymbol{x}' = (x', y', 1)^T$
- Constraint: $\boldsymbol{x}'^T \boldsymbol{F} \boldsymbol{x} = 0$

$$(x', y', 1) \begin{bmatrix} f_{11} & f_{12} & f_{13} \\ f_{21} & f_{22} & f_{23} \\ f_{31} & f_{32} & f_{33} \end{bmatrix} \begin{pmatrix} x \\ y \\ 1 \end{pmatrix} = 0 \quad \Rightarrow \quad (x'x, x'y, x', y'x, y'y, y', x, y, 1) \begin{pmatrix} f_{11} \\ f_{12} \\ f_{13} \\ f_{21} \\ f_{22} \\ f_{23} \\ f_{31} \\ f_{32} \\ f_{33} \end{pmatrix} = 0$$

The eight point algorithm

$$\underbrace{\begin{bmatrix} & & & \vdots & & & & & \\ x'x & x'y & x' & y'x & y'y & y' & x & y & 1 \\ & & & \vdots & & & & & \end{bmatrix}}_{\boldsymbol{U}} \begin{pmatrix} f_{11} \\ f_{12} \\ f_{13} \\ f_{21} \\ f_{22} \\ f_{23} \\ f_{31} \\ f_{32} \\ f_{33} \end{pmatrix} = \mathbf{0}$$

- Homogeneous least squares to find $\boldsymbol{f}$:

$$\underset{\|\boldsymbol{f}\|=1}{\arg\min} \|\boldsymbol{U}\boldsymbol{f}\|_2^2$$ ➡ Eigenvector of $\boldsymbol{U}^T\boldsymbol{U}$ with smallest eigenvalue

Enforcing rank-2 constraint

- We know $\boldsymbol{F}$ needs to be singular/rank 2. How do we force it to be singular?
- Solution: take SVD of the initial estimate and throw out the smallest singular value

$$\boldsymbol{F}_{\text{init}} = \boldsymbol{U\Sigma V}^T$$

$$\boldsymbol{\Sigma} = \begin{bmatrix} \sigma_1 & 0 & 0 \\ 0 & \sigma_2 & 0 \\ 0 & 0 & \sigma_3 \end{bmatrix} \rightarrow \boldsymbol{\Sigma}' = \begin{bmatrix} \sigma_1 & 0 & 0 \\ 0 & \sigma_2 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

$$\boldsymbol{F} = \boldsymbol{U\Sigma' V}^T$$

Enforcing rank-2 constraint

Initial ***F*** estimate

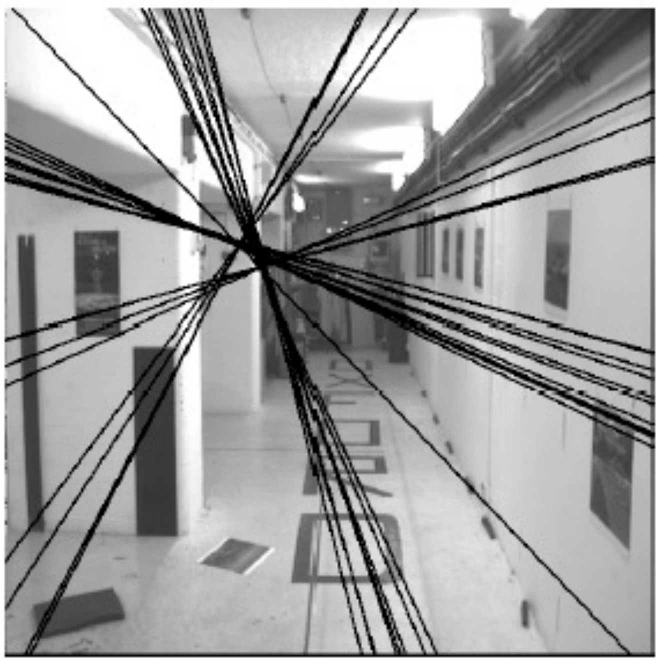

Rank-2 estimate

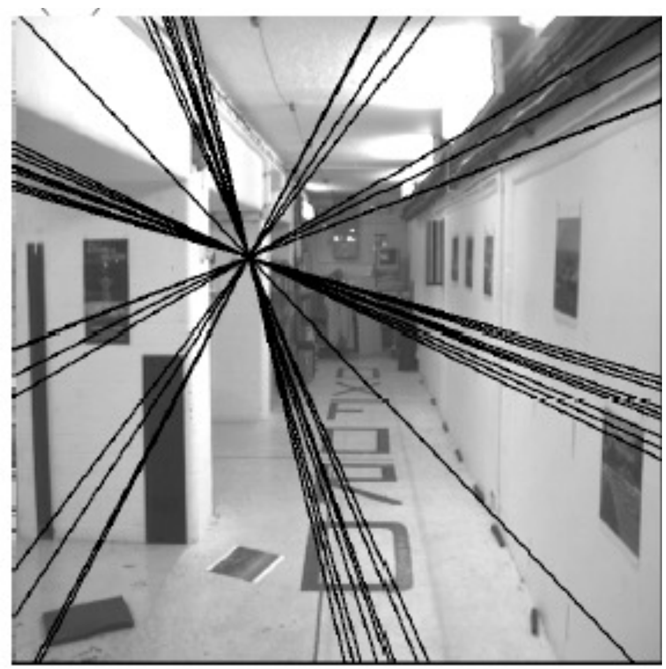

Normalized eight point algorithm

$$
\begin{array}{ccccccccc}
10^6 & 10^6 & 10^3 & 10^6 & 10^6 & 10^3 & 10^3 & 10^3 & 1
\end{array}
$$

$$
\underbrace{\begin{bmatrix} & & & \vdots & & & & & \\ x'x & x'y & x' & y'x & y'y & y' & x & y & 1 \\ & & & \vdots & & & & & \end{bmatrix}}_{\boldsymbol{U}}
\begin{pmatrix} f_{11} \\ f_{12} \\ f_{13} \\ f_{21} \\ f_{22} \\ f_{23} \\ f_{31} \\ f_{32} \\ f_{33} \end{pmatrix} = \mathbf{0}
$$

- Recall that x, y, x', y' are pixel coordinates. What might be the order of magnitude of each column of $\boldsymbol{U}$?
- This causes numerical instability!

The normalized eight-point algorithm

- In each image, center the set of points at the origin, and scale it so the mean squared distance between the origin and the points is 2 pixels
- Use the eight-point algorithm to compute $\boldsymbol{F}$ from the normalized points
- Enforce the rank-2 constraint
- Transform fundamental matrix back to original units: if $\boldsymbol{T}$ and $\boldsymbol{T'}$ are the normalizing transformations in the two images, then the fundamental matrix in original coordinates is $\boldsymbol{T'}^T\boldsymbol{F}\boldsymbol{T}$

R. Hartley. In defense of the eight-point algorithm. TPAMI 1997

Nonlinear estimation

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

$$\sum_i \left(\boldsymbol{x}_i'^T \boldsymbol{F} \boldsymbol{x}_i\right)^2$$

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

$$\sum_i \left[\mathrm{dist}(\boldsymbol{x}'_i, \boldsymbol{F}\boldsymbol{x}_i)^2 + \mathrm{dist}(\boldsymbol{x}_i, \boldsymbol{F}^T\boldsymbol{x}'_i)^2\right]$$

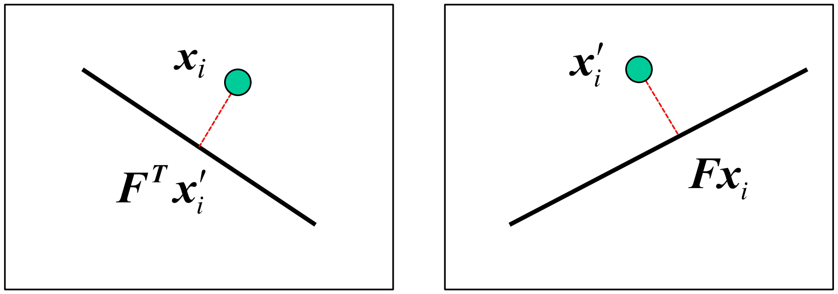

Comparison of estimation algorithms

	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

Seven-point algorithm

- Set up least squares system with seven pairs of matches and solve for null space (two vectors) using SVD
- Solve for polynomial equation to get coefficients of linear combination of null space vectors that satisfies $\det(\boldsymbol{F}) = 0$

Source: e.g., M. Pollefeys tutorial (2000)

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: $\boldsymbol{E} = \boldsymbol{K}'^{T}\boldsymbol{F}\boldsymbol{K}$
- The essential matrix gives us the relative rotation and translation between the cameras, or their extrinsic parameters
- Alternatively, if the calibration matrices are known (or in practice, if good initial guesses of the intrinsics are available), the five-point algorithm can be used to estimate relative camera pose

D. Nister. An efficient solution to the five-point relative pose problem. IEEE Trans. PAMI, 2004

The Fundamental Matrix Song

http://danielwedge.com/fmatrix/

Monocular visual odometry

- A calibrated camera
  - views a static scene,
  - moves,
  - views again

- Q: how did it move?
  - We care, because this allows us to recover movement from single cameras
  - We should be able to tell
    - Recall epipoles etc are quite informative about movement

Visual odometry

- Use eight point algorithm, recover fundamental matrix
- Recall:

$$\mathcal{F} = k\mathcal{C}_L^{-T}\mathcal{R}\mathcal{S}\mathcal{C}_R^{-1}$$

- But we know calibration, which yields essential matrix

$$\mathcal{E} = k\mathcal{T}\mathcal{R}$$

- Q: what can we get out of essential matrix?

Visual odometry, II

Antisymmetric

$$\mathcal{E} = k\mathcal{T}\mathcal{R}$$

Unknown constant

Rotation

Visual odometry, III

- Recall singular value decomposition

Orthonormal

$$\mathcal{M} = \mathcal{U}\Sigma\mathcal{V}^T$$

Diagonal matrix
Of non-negative singular values

Visual odometry, IV

$$\mathcal{M} = \mathcal{U}\Sigma\mathcal{V}^T$$

$$\mathcal{E} = k\mathcal{T}\mathcal{R}$$

- Notice

$\mathcal{E}$ and $\mathcal{E}\mathcal{R}^T$ have the same singular values

- So that

Singular values($\mathcal{E}$) = singular values($\mathcal{T}$)

Visual odometry, V

Check

$$\text{singular values}(\mathcal{T}) = \begin{pmatrix} s & 0 & 0 \\ 0 & s & 0 \\ 0 & 0 & 0 \end{pmatrix}$$

Write

$$\mathcal{U}_E \Sigma_E \mathcal{V}_E^T = \mathcal{E}$$

Visual odometry, VI

- Write

$$\mathcal{W} = \begin{pmatrix} 0 & -1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

- Check

$$\mathcal{U}_E \Sigma_E \mathcal{W} \mathcal{U}_E^T$$ is antisymmetric

$$\mathcal{U}_E \mathcal{W}^{-1} \mathcal{V}_E^T$$ is a rotation

Visual odometry, VII

So we can recover

Rotation exactly

Translation up to scale

From an essential matrix