

Epipolar Geometry

EECS 442 – David Fouhey and Justin Johnson
Winter 2021, University of Michigan


<https://web.eecs.umich.edu/~justincj/teaching/eecs442/WI2021/>

Multi-view geometry



Multi-view geometry problems

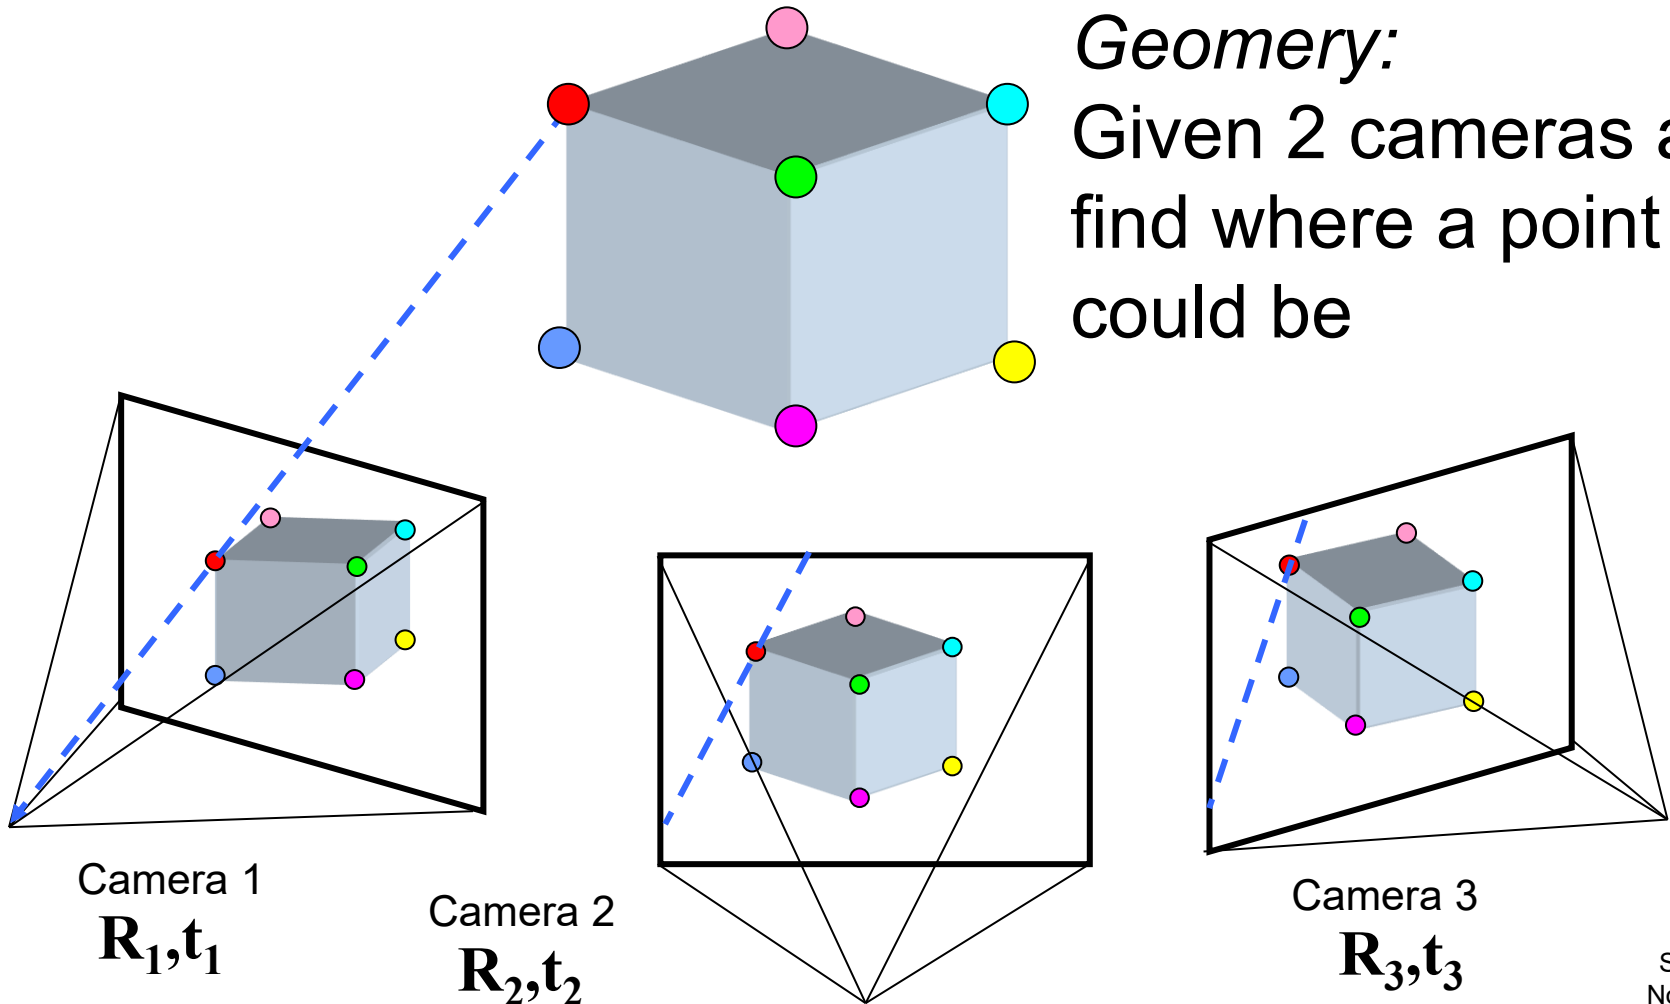
Recovering structure:
Given cameras and correspondences, find 3D.



Multi-view geometry problems

*Stereo/Epipolar
Geometry:*



Given 2 cameras and
find where a point
could be




Multi-view geometry problems

Motion:


Figure out R, t for a set of cameras given correspondences



Camera 1
 R_1, t_1 ?



Camera 2
 R_2, t_2 ?



? Camera 3
 R_3, t_3

Two-view geometry



Camera Geometry Reminder

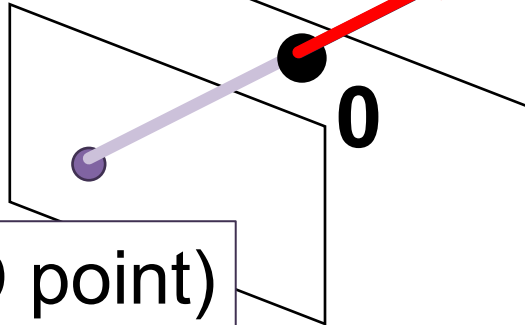
$K^{-1}p$ (Ray)
3 h. coordinates

X (3D point)
4 h. coordinates


p (2D point)
3 h. coordinates
Pretending image
plane is in front

p (2D point)
3 h. coordinates
Actual location

Have camera with pinhole
at origin $\mathbf{0}$




Epipolar Geometry




Suppose we have two cameras at origins o , o'
Baseline is the line connecting the origins

Epipolar Geometry



Now add a **point X** , which projects to p and p'


Epipolar Geometry



The plane formed by X , o , and o' is called the epipolar plane

There is a family of planes per o , o'

Epipolar Geometry



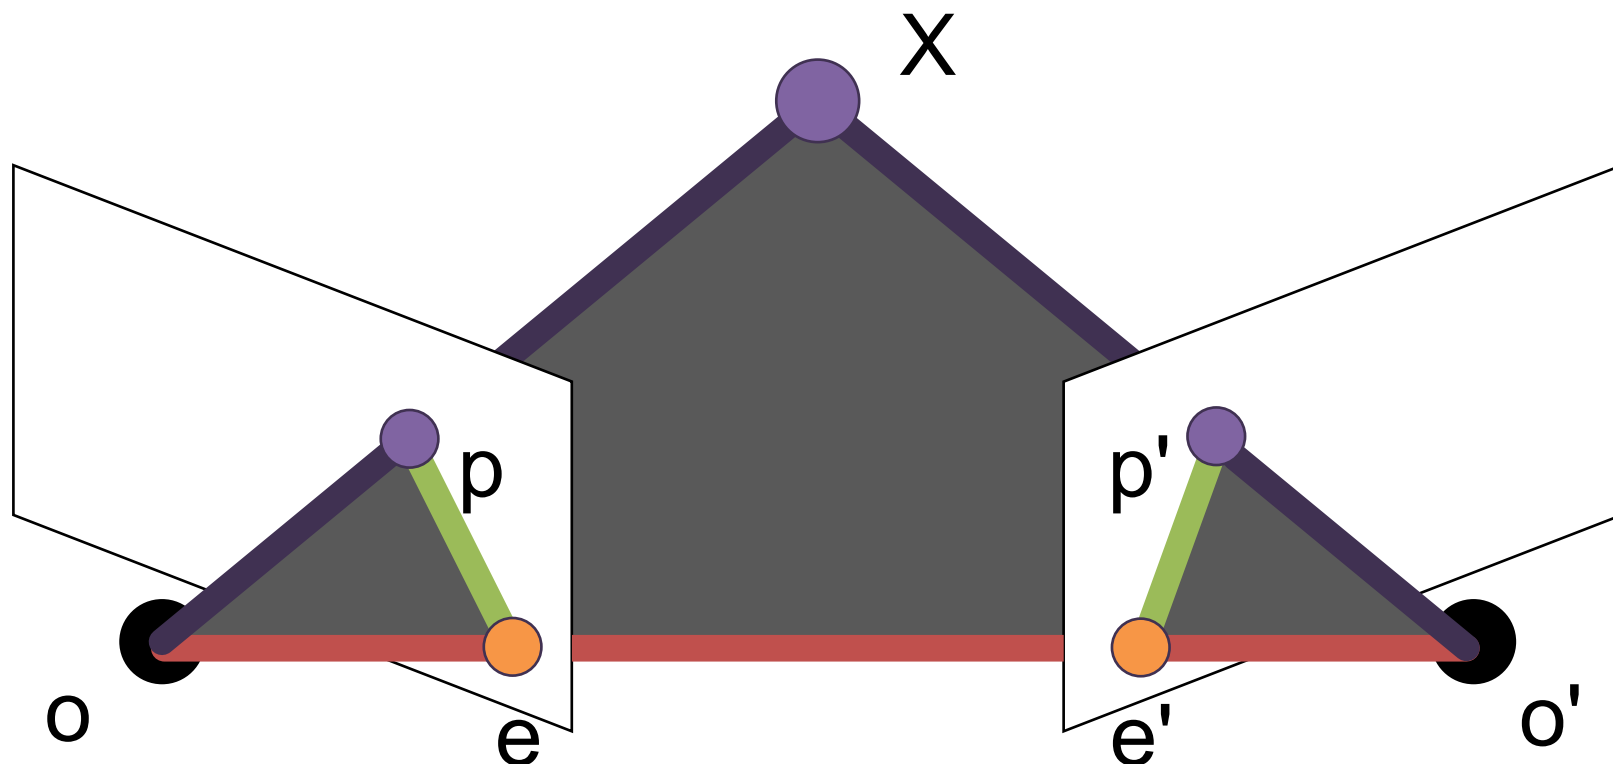
- Epipoles e, e' are where the baseline intersects the image planes
- Projection of other camera in the image plane

The Epipole




Photo by Frank Dellaert

Epipolar Geometry




- Epipolar lines go between the epipoles and the projections of the points.
- Intersection of epipolar plane with image plane

Example: Converging Cameras



Epipoles finite, maybe in image; epipolar lines converge

Example: Converging Cameras




Epipolar lines come in pairs: given a point p , we can construct the epipolar line for p' .

Example 1: Converging Cameras




Example: Parallel to Image Plane




Suppose the cameras are both facing outwards.
Where are the epipoles (proj. of other camera)?

Example: Parallel to Image Plane



Epipoles *infinitely* far away, epipolar lines parallel

Example: Parallel to Image Plane



Example: Forward Motion




Example: Forward Motion




Example: Forward Motion

Epipole is focus of expansion / principal point of the camera.

Epipolar lines go out from principal point

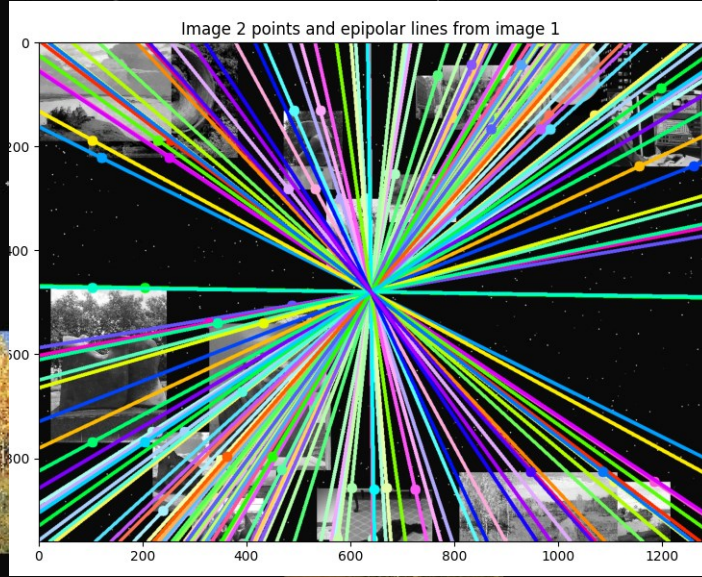




Motion perpendicular to image plane

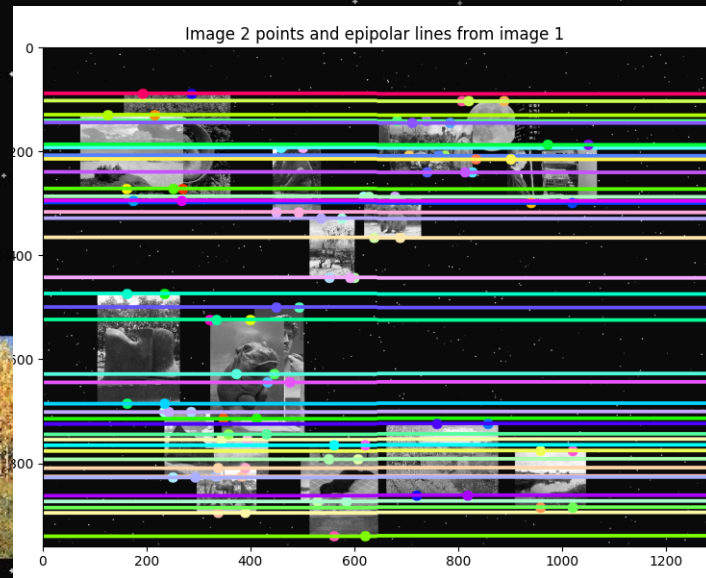


<http://vimeo.com/48425421>

HW6 Preview









So?

Epipolar Geometry




- Suppose we don't know X and just have p
- Can construct the epipolar line in the other image

Epipolar Geometry



- Suppose we don't know X and just have p
- Corresponding p' is on corresponding epipolar line

Epipolar Geometry




- Suppose we don't know X and just have p'
- Corresponding p is on corresponding epipolar line

Epipolar Geometry

- If I want to do stereo, I want to find a corresponding pixel for each pixel in the image:
- Naïve search:
 - For each pixel, search every other pixel
- With epipolar geometry:
 - For each pixel, search along each line (1D search)


Epipolar constraint example



Epipolar Constraint: One Note

- This year, I've updated my slides to match the more common convention. If you read earlier slides from me, they may not match
- I'm pretty sure I got them all right but I apologize for any p vs p' typos!


Epipolar Constraint: Calibrated Case



- If we know intrinsic and extrinsic parameters, set coordinate system to first camera
- Projection matrices: $M_1 = K[I, \mathbf{0}]$ and $M_2 = K'[R, t]$
- **What are:**


$$M_1 X \quad M_2 X \quad K^{-1} p \quad K'^{-1} p'$$

Epipolar Constraint: Calibrated Case



- Given calibration, $\hat{p} = K^{-1} p$ and $\hat{p}' = K'^{-1} p'$ are “normalized coordinates”
- Note that \hat{p}' is actually translated and rotated to o'


Epipolar Constraint: Calibrated Case



- The following are all co-planar: $R\hat{p}$, t , \hat{p}' (can ignore translation for co-planarity here)
- One way to check co-planarity (triple product):

$$\hat{p}'^T (t \times R\hat{p}) = 0$$


Epipolar Constraint: Calibrated Case



$$\hat{p}'^T (\mathbf{t} \times R\hat{p}) = 0 \longrightarrow \hat{p}'^T \begin{bmatrix} 0 & [t_x] & t_2 \\ t_3 & 0 & -t_1 \\ -t_2 & t_1 & 0 \end{bmatrix} R\hat{p} = 0$$

Want something like $\mathbf{x}^T \mathbf{A} \mathbf{y} = 0$. What's \mathbf{A} ?


Epipolar Constraint: Calibrated Case



Essential matrix (Longuet-Higgins, 1981): $E = [t_x]R$


If you have a normalized point \hat{p} , its correspondence \hat{p}' must satisfy $\hat{p}'^T E \hat{p} = 0$

Essential Matrix Facts




- Suppose we know \mathbf{E} and $\hat{p}'^T \mathbf{E} \hat{p} = 0$. What is the set $\{\mathbf{x}: \hat{p}'^T \mathbf{E} \mathbf{x} = 0\}$?
- $\hat{p}'^T \mathbf{E}$ gives equation of the epipolar line (in $ax+by+c=0$ form) in image for o .
- What's $\mathbf{E}^T \hat{p}'$?

Essential Matrix Facts




- $E\hat{e} = 0$ and $E^T\hat{e}' = 0$ (epipoles are the nullspace of E – note all epipolar lines pass through epipoles)
- **Degrees of freedom (Recall $E = [t_x]R$)?**
- $5 - 3 (R) + 3 (t) - 1$ due to scale ambiguity
- E is singular (rank 2); it has two non-zero and *identical* singular values

Essential Matrix Facts



- One nice thing: if I estimate E from two images (more on this later), it's unique up to easy symmetries

What if we don't know K ?




Have: $\hat{p} = K^{-1}p$, $\hat{p}' = K'^{-1}p'$, $\hat{p}'^T E \hat{p} = 0$
 $(K'^{-1}p')^T E (K^{-1}p) = 0 \implies p'^T K'^{-T} E K^{-1} p = 0$

Set: $\underbrace{F = K'^{-T} E K^{-1}}$ Then: $p'^T F p = 0$

Fundamental Matrix (Faugeras and Luong, 1992)

Fundamental Matrix Fundamentals



- $Fp, F^T p'$ are epipolar lines for p', p
- $Fe = 0, F^T e' = 0$
- F is singular (rank 2)
- F has seven degrees of freedom
- F definitely not unique

Estimating the fundamental matrix



Estimating the fundamental matrix

- F has 7 degrees of freedom so it's in principle possible to fit F with seven correspondences, but it's a slightly more complex and typically not taught in regular vision classes

Estimating the fundamental matrix

Given correspondences $\mathbf{p} = [u, v, 1]$ and $\mathbf{p}' = [u', v', 1]$ (e.g., via SIFT) we know: $\mathbf{p}'^T \mathbf{F} \mathbf{p} = 0$

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

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

How do we solve for f?

How many correspondences do we need?

Leads to the **eight point algorithm**

Eight Point Algorithm

Each point gives an equation:

$$\begin{bmatrix} u'u, u'v, u', v'u, v'v, v', u, v, 1 \\ f_{11}, f_{12}, f_{13}, f_{21}, f_{22}, f_{23}, f_{31}, f_{32}, f_{33} \end{bmatrix} \cdot \mathbf{f} = 0$$

Stack equations to yield \mathbf{U} :

$$\mathbf{U} = \begin{bmatrix} \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ u'_i u_i & u'_i v_i & u'_i & v'_i u_i & v'_i v_i & v'_i & u_i & v_i & 1 \end{bmatrix}$$

Usual eigenvalue stuff to find \mathbf{f} (\mathbf{F} unrolled):

$$\arg \min_{\|\mathbf{f}\|=1} \|\mathbf{U}\mathbf{f}\|_2^2 \longrightarrow \text{Eigenvector of } \mathbf{U}^T \mathbf{U} \text{ with smallest eigenvalue}$$

Eight Point Algorithm – Difficulty 1

If we estimate F , we get some 3×3 matrix F .
We know F needs to be singular/rank 2. How do we force F to be singular?


$$U\Sigma V^T = F_{init}$$



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



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



$$F = U\Sigma'V^T$$

Open it up with SVD, mess with singular values, put it back together.

See Eckart–Young–Mirsky theorem if you're interested

Eight Point Algorithm – Difficulty 1

Estimated F
(Wrong)



Estimated+SVD'd F
(Correct)



Eight Point Algorithm – Difficulty 2

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

Recall: u, u' are in pixels. Suppose image is 1Kx1K

How big might uu' be? How big might u' be?

Each row looks like:

$$U = \begin{bmatrix} 10^6 & 10^6 & 10^3 & 10^6 & \vdots & 10^6 & 10^3 & 10^3 & 10^3 & 1 \\ & & & & \vdots & & & & & \end{bmatrix}$$

Then: $U^T U_{1,1}$ is $\sim 10^{12}$, $U^T U_{2,9}$ is $\sim 10^3$

Eight Point Algorithm – Difficulty 2

Numbers of varying magnitude → instability

Remember: a floating point number (float/double) isn't a "real" number: for sign, coefficient, exponent integers
 $(-1)^{\text{sign}} * \text{coefficient} * 2^{\text{exponent}}$

Exercise to see how this screws up: add up Gaussian noise (mean=100, std=10), divide by number you added up


Remember Numerical Instability?

Code :

```
x += N(100, 10)
i += 1
mean = x/I
```

Only change is the
of bits in
accumulator x

Note: 50M is 50
1Kx1K images



Solution: Normalized 8-point

- 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 F from the normalized points
- Enforce the rank-2 constraint (for example, take SVD of F and throw out the smallest singular value)
- Transform fundamental matrix back to original units: if T and T' are the normalizing transformations in the two images, then the fundamental matrix in original coordinates is $T'^T F T$

R. Hartley


Last Trick

Minimizing via $U^T U$ minimizes sum of squared *algebraic* distances between points \mathbf{p}_i and epipolar lines $\mathbf{F}\mathbf{p}'_i$ (or points \mathbf{p}'_i and epipolar lines $\mathbf{F}^T\mathbf{p}_i$):


$$\sum_i (\mathbf{p}'_i{}^T \mathbf{F} \mathbf{p}_i)^2$$

May want to minimize *geometric* distance:

$$\sum_i d(\mathbf{p}'_i, \mathbf{F}\mathbf{p}_i)^2 + d(\mathbf{p}_i, \mathbf{F}^T\mathbf{p}'_i)^2$$



Comparison



	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 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