

Lecture 13

Structure from motion

So far: camera transformation

$$\mathbf{P} = \mathbf{KR}[\mathbf{I} | -\mathbf{C}]$$

3x3 intrinsics 3x3 3D rotation 3x3 identity 3x1 3D translation

So far: camera calibration

Estimate camera parameters ($K[R|t]$) by measuring


- real world points \mathbf{X}_i in world space
- the same points in pixel space \mathbf{x}_i
- Solve for $K[R | t]$ using SVD after posing the problem as $\mathbf{A}\mathbf{p} = 0$
 - where \mathbf{p} are camera parameters and \mathbf{A} is obtained from mapping \mathbf{X}_i to \mathbf{x}_i

So far: in other words, we can estimate camera motion given multiple images...



https://kornia.readthedocs.io/en/latest/applications/image_matching.html

What we want to do today: extract structure!

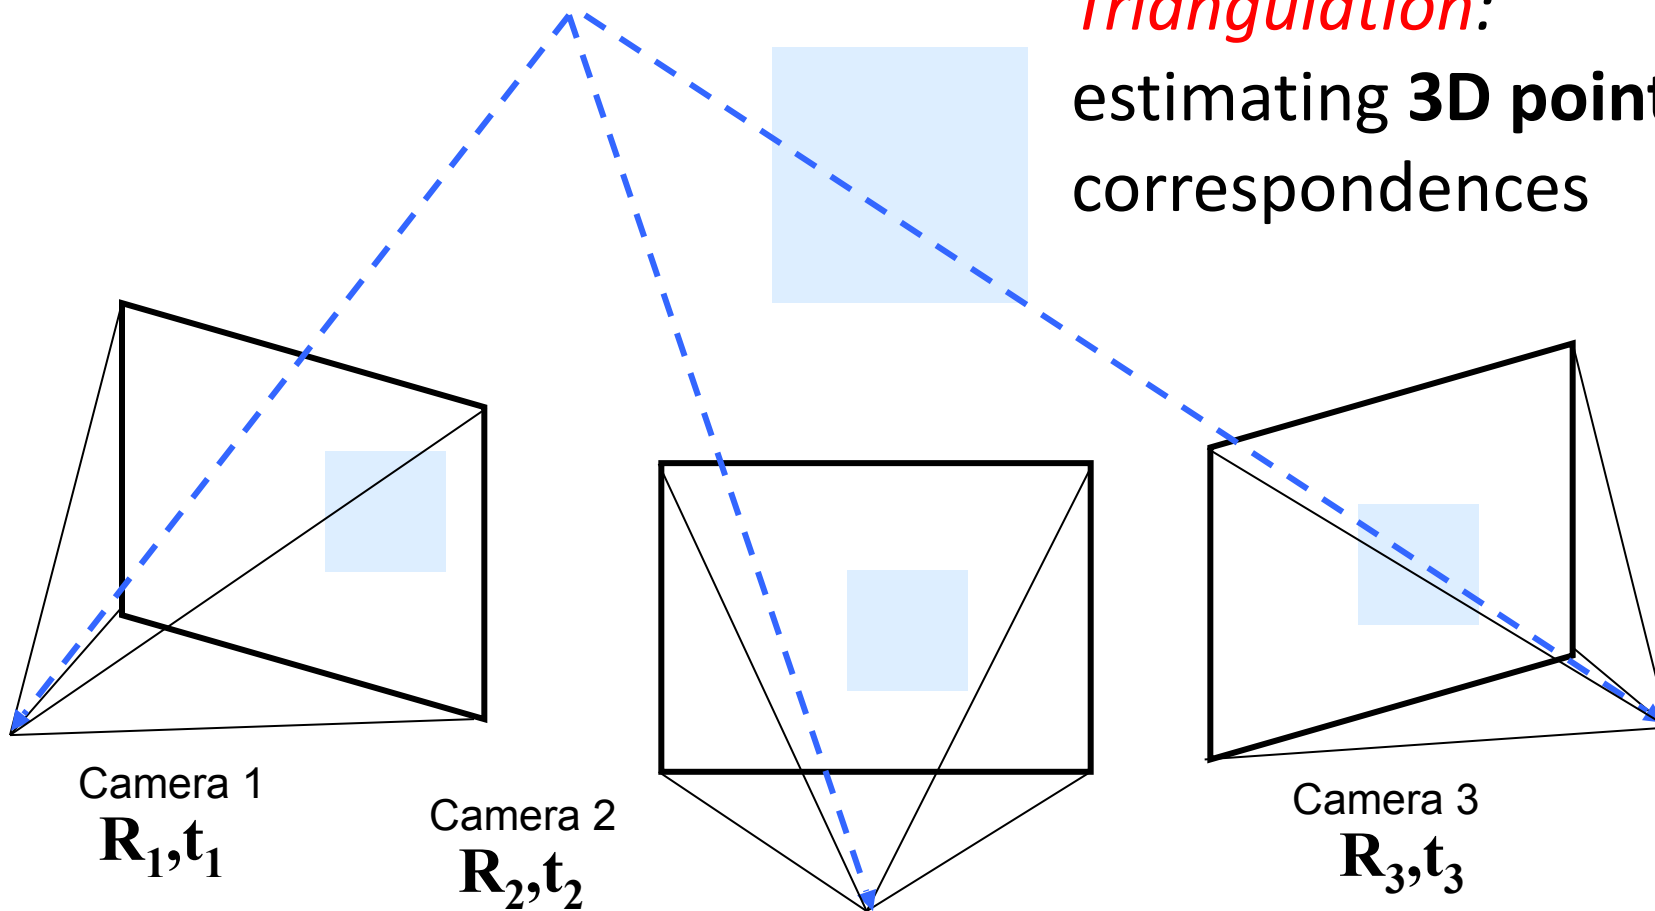


If we knew the camera parameters, we would be able to find the 3D world coordinates of things




Triangulation:

estimating **3D points** from **2D correspondences**



First we need to estimate motion (R, t)
from 2D correspondences (we already did this)




Today's agenda

- Triangulation
- Epipolar geometry
- Essential matrix
- Fundamental matrix
- Structure from motion

Today's agenda

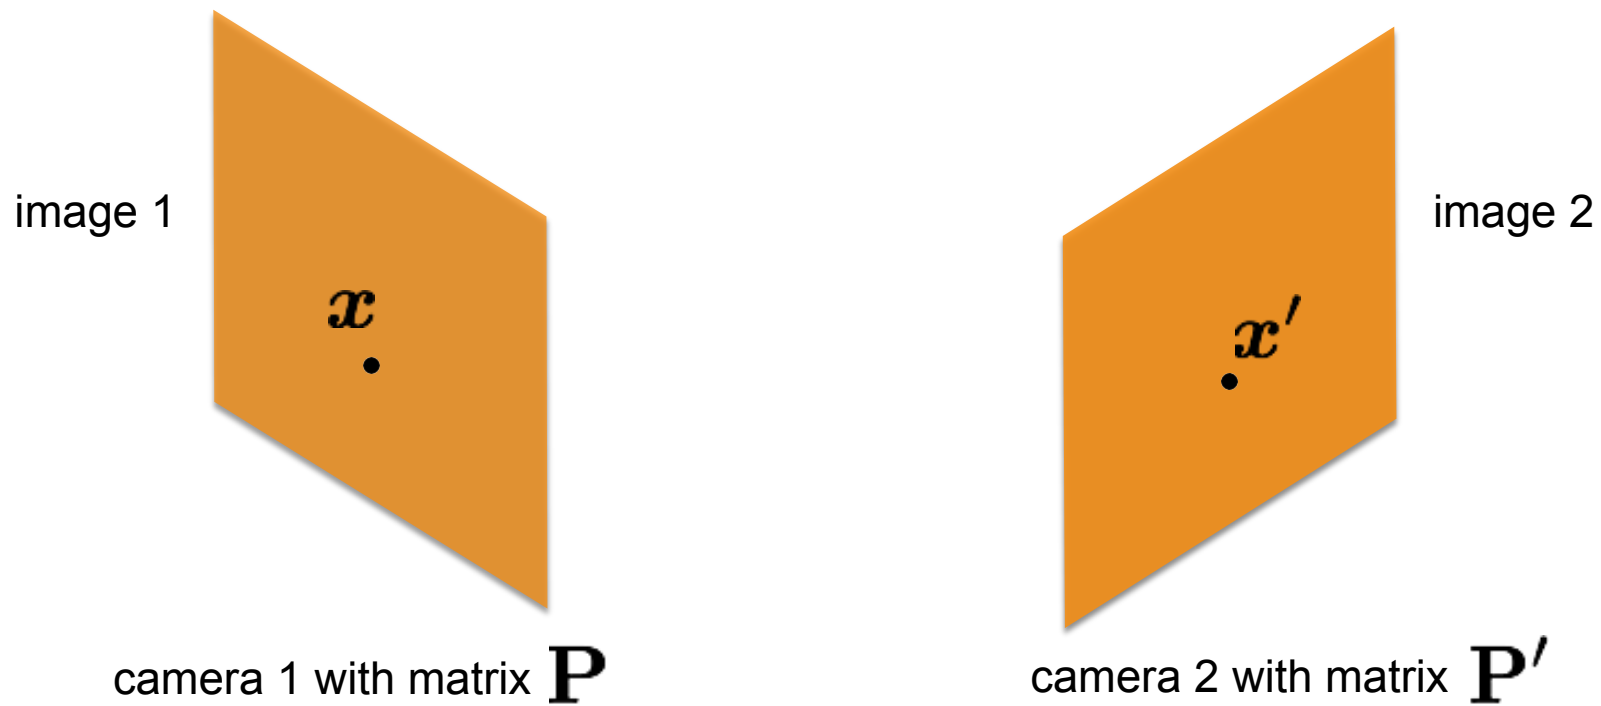
- **Triangulation**
- Epipolar geometry
- Essential matrix
- Fundamental matrix
- Structure from motion

Triangulation



Triangulation

Where is the 3D point that maps to the two x 's?



Triangulation formalization

Given a set of (noisy) matched points

$$\{\mathbf{x}_i, \mathbf{x}'_i\}$$

and camera matrices

$$\mathbf{P}, \mathbf{P}'$$

Estimate the 3D point

$$\mathbf{X}$$


Triangulation equation

$$\mathbf{x} = \mathbf{P}\mathbf{X}$$

known known

Q. Can we compute \mathbf{X} from a single correspondence \mathbf{x} ?


Triangulation



Triangulation

Apply the **pseudo-inverse** of \mathbf{P} on \mathbf{x} . Then connect the two points.

This procedure is called **backprojection**



$$\mathbf{x} = \mathbf{P}\mathbf{X}$$

We lose information going from 3D to 2D.
Specifically, we lose depth information

$$\mathbf{x} = \alpha\mathbf{P}\mathbf{X}$$


Scaling by α . is the same ray direction but differs by a scale factor corresponding to depth

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix} = \alpha \begin{bmatrix} p_1 & p_2 & p_3 & p_4 \\ p_5 & p_6 & p_7 & p_8 \\ p_9 & p_{10} & p_{11} & p_{12} \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix}$$

How do we solve for unknowns in a similarity relation?

Triangulation


Q. How do we find the exact point on the ray?



Reminder: cross products from linear algebra

Vector (cross) product

takes two vectors and returns a vector perpendicular to both



$$\mathbf{a} \times \mathbf{b} = \begin{bmatrix} a_2b_3 - a_3b_2 \\ a_3b_1 - a_1b_3 \\ a_1b_2 - a_2b_1 \end{bmatrix}$$

cross product of two vectors in the same direction is zero vector

$$\mathbf{a} \times \mathbf{a} = \mathbf{0}$$

remember this!!!

$$\mathbf{c} \cdot \mathbf{a} = 0$$

$$\mathbf{c} \cdot \mathbf{b} = 0$$

Reminder: cross products from linear algebra


$$\mathbf{a} \times \mathbf{b} = \begin{bmatrix} a_2b_3 - a_3b_2 \\ a_3b_1 - a_1b_3 \\ a_1b_2 - a_2b_1 \end{bmatrix}$$

Can also be written as a matrix multiplication

$$\mathbf{a} \times \mathbf{b} = [\mathbf{a}]_{\times} \mathbf{b} = \begin{bmatrix} 0 & -a_3 & a_2 \\ a_3 & 0 & -a_1 \\ -a_2 & a_1 & 0 \end{bmatrix} \begin{bmatrix} b_1 \\ b_2 \\ b_3 \end{bmatrix}$$

Skew symmetric

Compare with: dot product



$$c \cdot a = 0$$

$$c \cdot b = 0$$

Dot product of two orthogonal vectors is zero!

Back to triangulation

$$\mathbf{x} = \alpha \mathbf{P} \mathbf{X}$$

Same direction but differs by a
scale factor

*How can we rewrite this using vector
products?*

$$\mathbf{x} = \alpha \mathbf{P} \mathbf{X}$$

Same direction but differs by a scale factor

$$\mathbf{x} \times \mathbf{P} \mathbf{X} = \mathbf{0}$$

Cross product of two vectors of same direction is zero
(this equality removes the scale factor)

$$\mathbf{x} = \alpha \mathbf{P} \mathbf{X}$$

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix} = \alpha \begin{bmatrix} p_1 & p_2 & p_3 & p_4 \\ p_5 & p_6 & p_7 & p_8 \\ p_9 & p_{10} & p_{11} & p_{12} \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix}$$

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix} = \alpha \begin{bmatrix} \text{---} & \mathbf{p}_1^\top & \text{---} \\ \text{---} & \mathbf{p}_2^\top & \text{---} \\ \text{---} & \mathbf{p}_3^\top & \text{---} \end{bmatrix} \begin{bmatrix} | \\ \mathbf{X} \\ | \end{bmatrix}$$

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix} = \alpha \begin{bmatrix} \mathbf{p}_1^\top \mathbf{X} \\ \mathbf{p}_2^\top \mathbf{X} \\ \mathbf{p}_3^\top \mathbf{X} \end{bmatrix}$$

$$\mathbf{x} \times \mathbf{P}\mathbf{X} = \mathbf{0}$$

$$\begin{bmatrix} x \\ y \\ 1 \end{bmatrix} \times \begin{bmatrix} \mathbf{p}_1^\top \mathbf{X} \\ \mathbf{p}_2^\top \mathbf{X} \\ \mathbf{p}_3^\top \mathbf{X} \end{bmatrix} = \begin{bmatrix} y\mathbf{p}_3^\top \mathbf{X} - \mathbf{p}_2^\top \mathbf{X} \\ \mathbf{p}_1^\top \mathbf{X} - x\mathbf{p}_3^\top \mathbf{X} \\ x\mathbf{p}_2^\top \mathbf{X} - y\mathbf{p}_1^\top \mathbf{X} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}$$

$$\begin{bmatrix} y\mathbf{p}_3^\top \mathbf{X} - \mathbf{p}_2^\top \mathbf{X} \\ \mathbf{p}_1^\top \mathbf{X} - x\mathbf{p}_3^\top \mathbf{X} \\ x\mathbf{p}_2^\top \mathbf{X} - y\mathbf{p}_1^\top \mathbf{X} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}$$

$$\mathbf{x} \times \mathbf{P}\mathbf{X} = \mathbf{0}$$

$$\begin{bmatrix} x \\ y \\ 1 \end{bmatrix} \times \begin{bmatrix} \mathbf{p}_1^\top \mathbf{X} \\ \mathbf{p}_2^\top \mathbf{X} \\ \mathbf{p}_3^\top \mathbf{X} \end{bmatrix} = \begin{bmatrix} y\mathbf{p}_3^\top \mathbf{X} - \mathbf{p}_2^\top \mathbf{X} \\ \mathbf{p}_1^\top \mathbf{X} - x\mathbf{p}_3^\top \mathbf{X} \\ x\mathbf{p}_2^\top \mathbf{X} - y\mathbf{p}_1^\top \mathbf{X} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}$$

$$\begin{bmatrix} y\mathbf{p}_3^\top \mathbf{X} - \mathbf{p}_2^\top \mathbf{X} \\ \mathbf{p}_1^\top \mathbf{X} - x\mathbf{p}_3^\top \mathbf{X} \\ x\mathbf{p}_2^\top \mathbf{X} - y\mathbf{p}_1^\top \mathbf{X} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}$$

Third line is a linear combination of the first and second lines. (x times the first line plus y times the second line)

So, we only get **2 equations to calculate 3 unknowns**: X, Y, Z

Remove third
row, and
rearrange as
system on
unknowns

$$\begin{bmatrix} y\mathbf{p}_3^\top \mathbf{X} - \mathbf{p}_2^\top \mathbf{X} \\ \mathbf{p}_1^\top \mathbf{X} - x\mathbf{p}_3^\top \mathbf{X} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$$

$$\begin{bmatrix} y\mathbf{p}_3^\top - \mathbf{p}_2^\top \\ \mathbf{p}_1^\top - x\mathbf{p}_3^\top \end{bmatrix} \mathbf{X} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$$

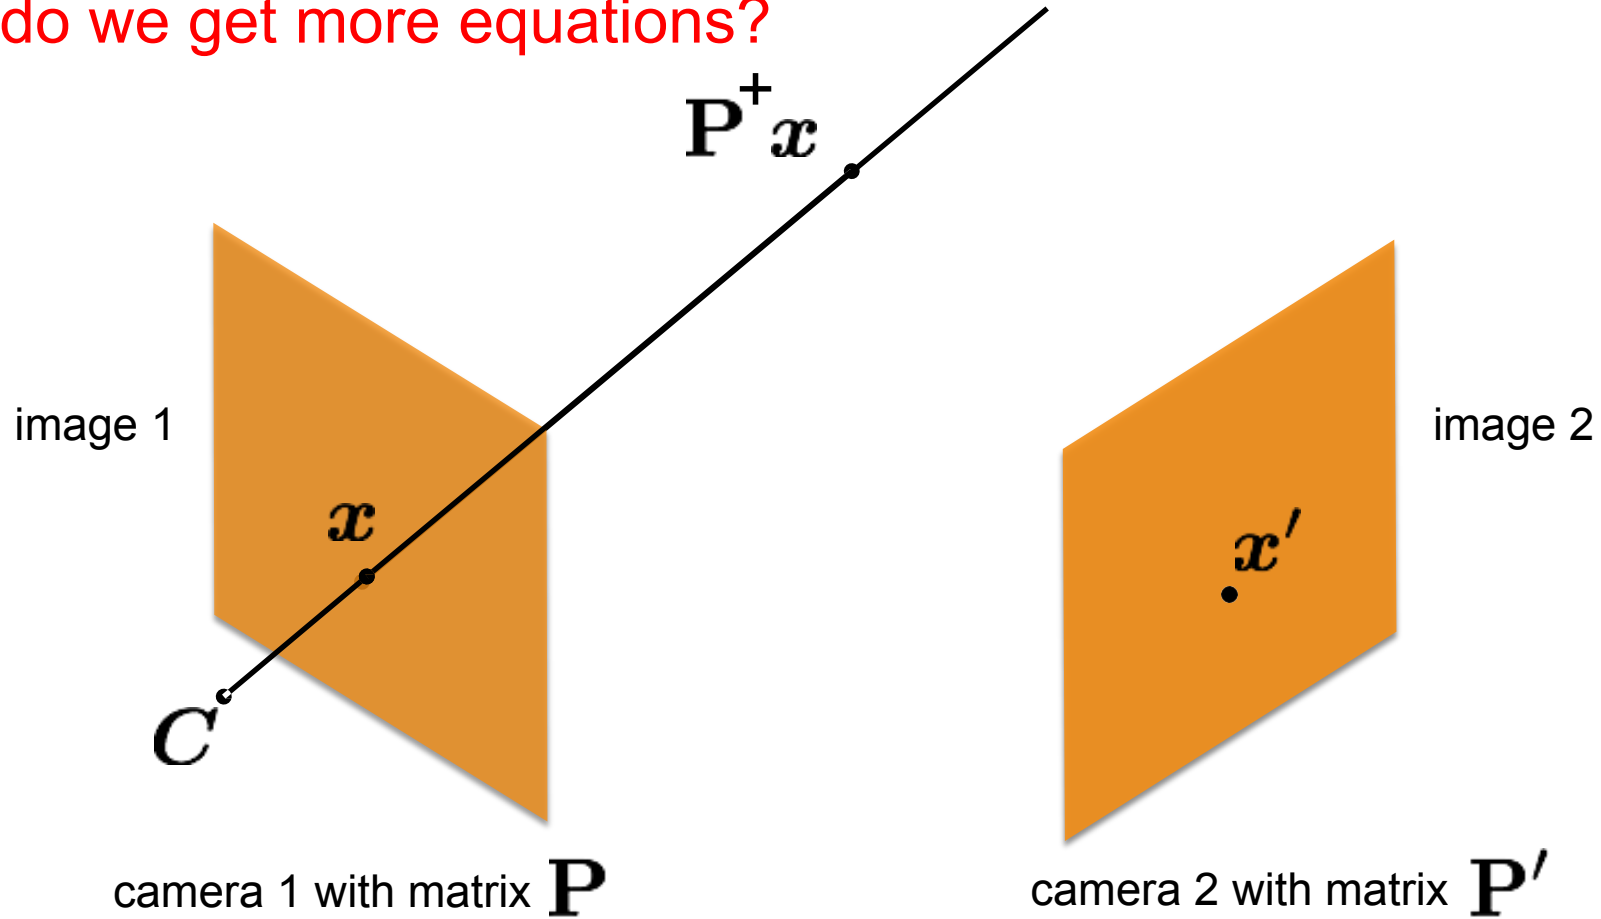
$$\mathbf{A}_i \mathbf{X} = \mathbf{0}$$

This is the proof that we can not solve for X... we only have two equations.


Q. So, how do we get more equations?

Triangulation

Q. So, how do we get more equations?



How do we find this intersection?



Collect more equations from other cameras

Two rows
from camera
one

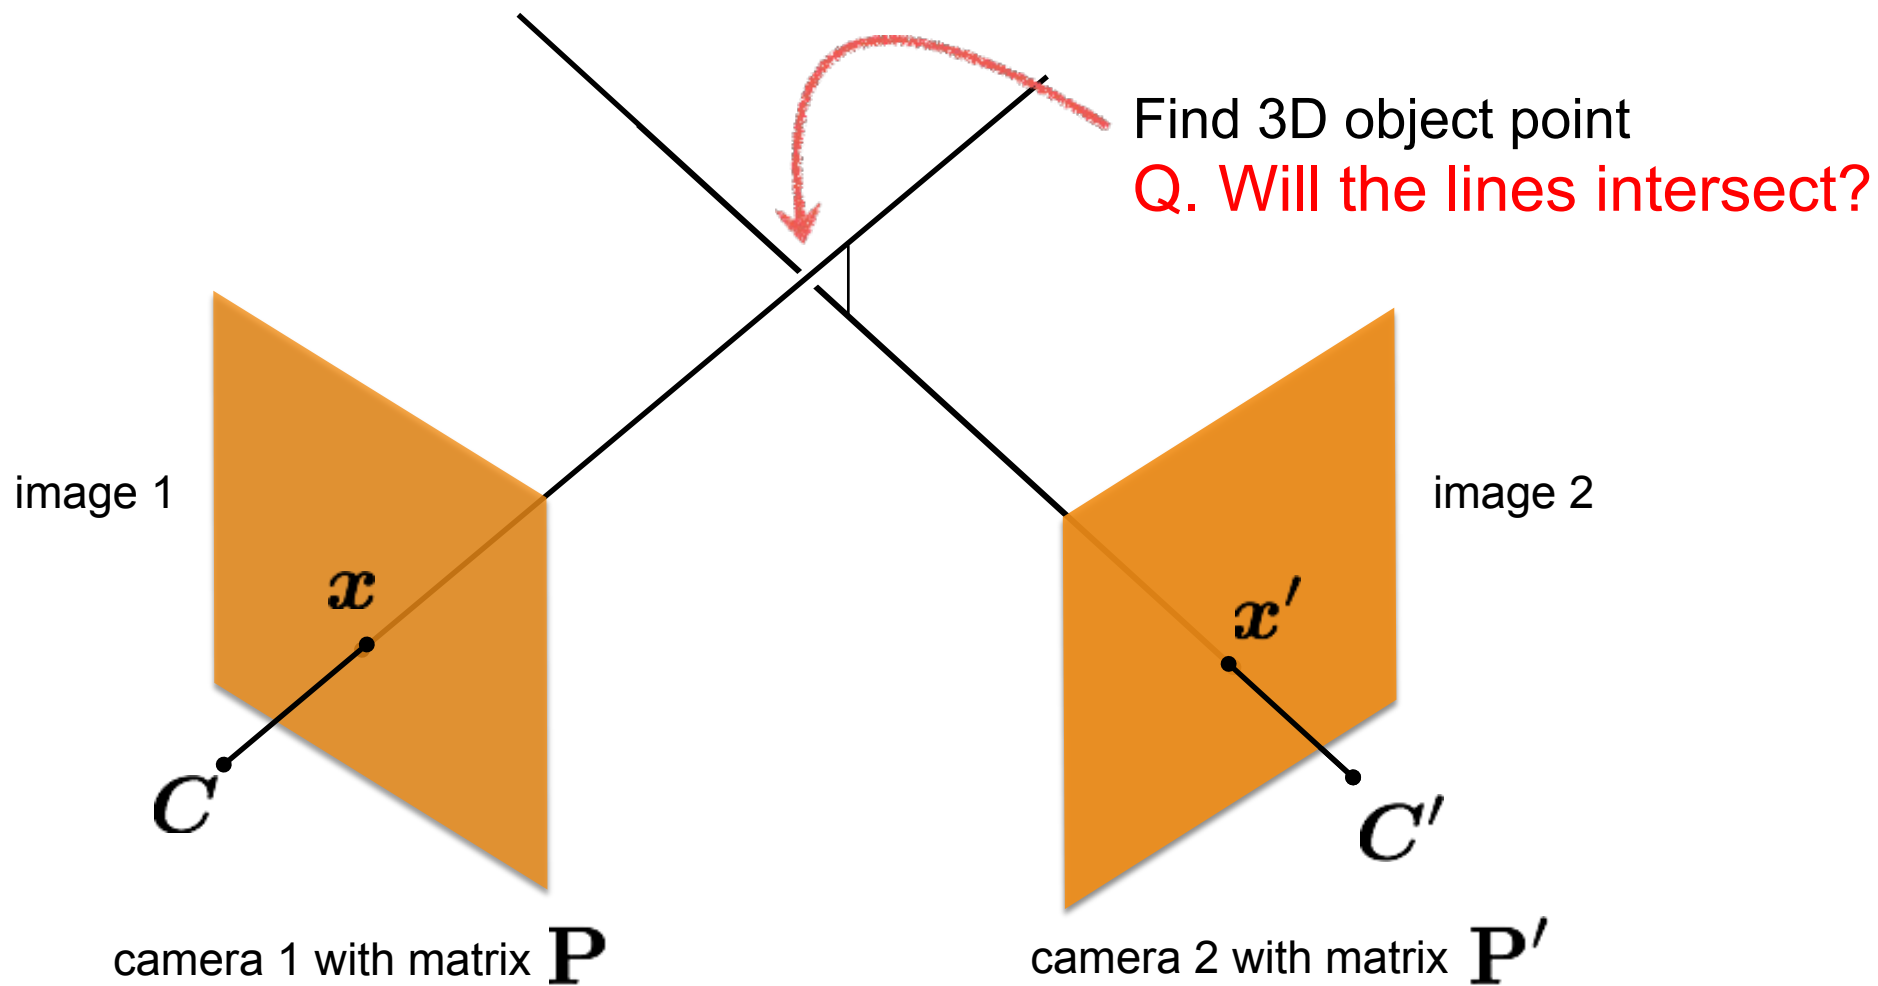
Two rows
from camera
two

$$\begin{bmatrix} y\mathbf{p}_3^\top - \mathbf{p}_2^\top \\ \mathbf{p}_1^\top - x\mathbf{p}_3^\top \\ y'\mathbf{p}'_3{}^\top - \mathbf{p}'_2{}^\top \\ \mathbf{p}'_1{}^\top - x'\mathbf{p}'_3{}^\top \end{bmatrix} \mathbf{X} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}$$

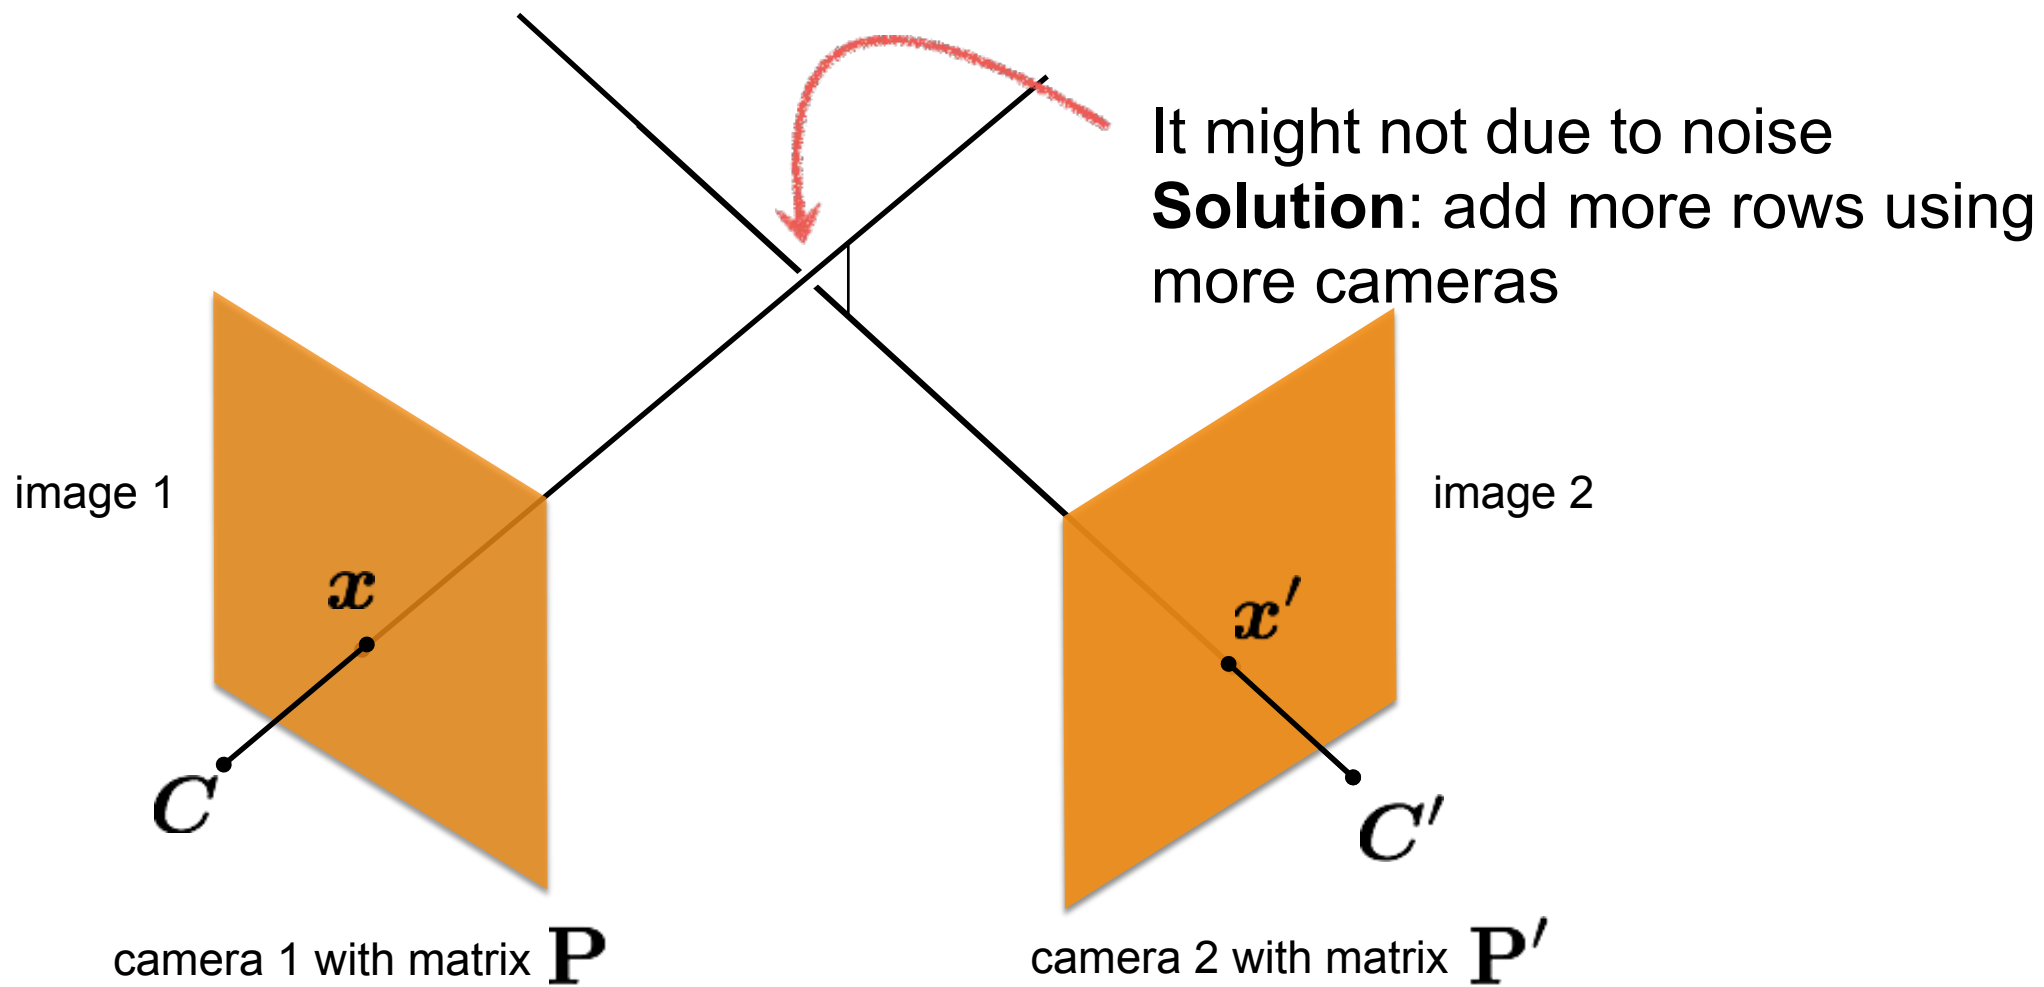
$$\mathbf{A}\mathbf{X} = \mathbf{0}$$

Now, we can solve for X using SVD

Triangulation




Triangulation




Today's agenda

- Triangulation
- **Epipolar geometry**
- Essential matrix
- Fundamental matrix
- Structure from motion

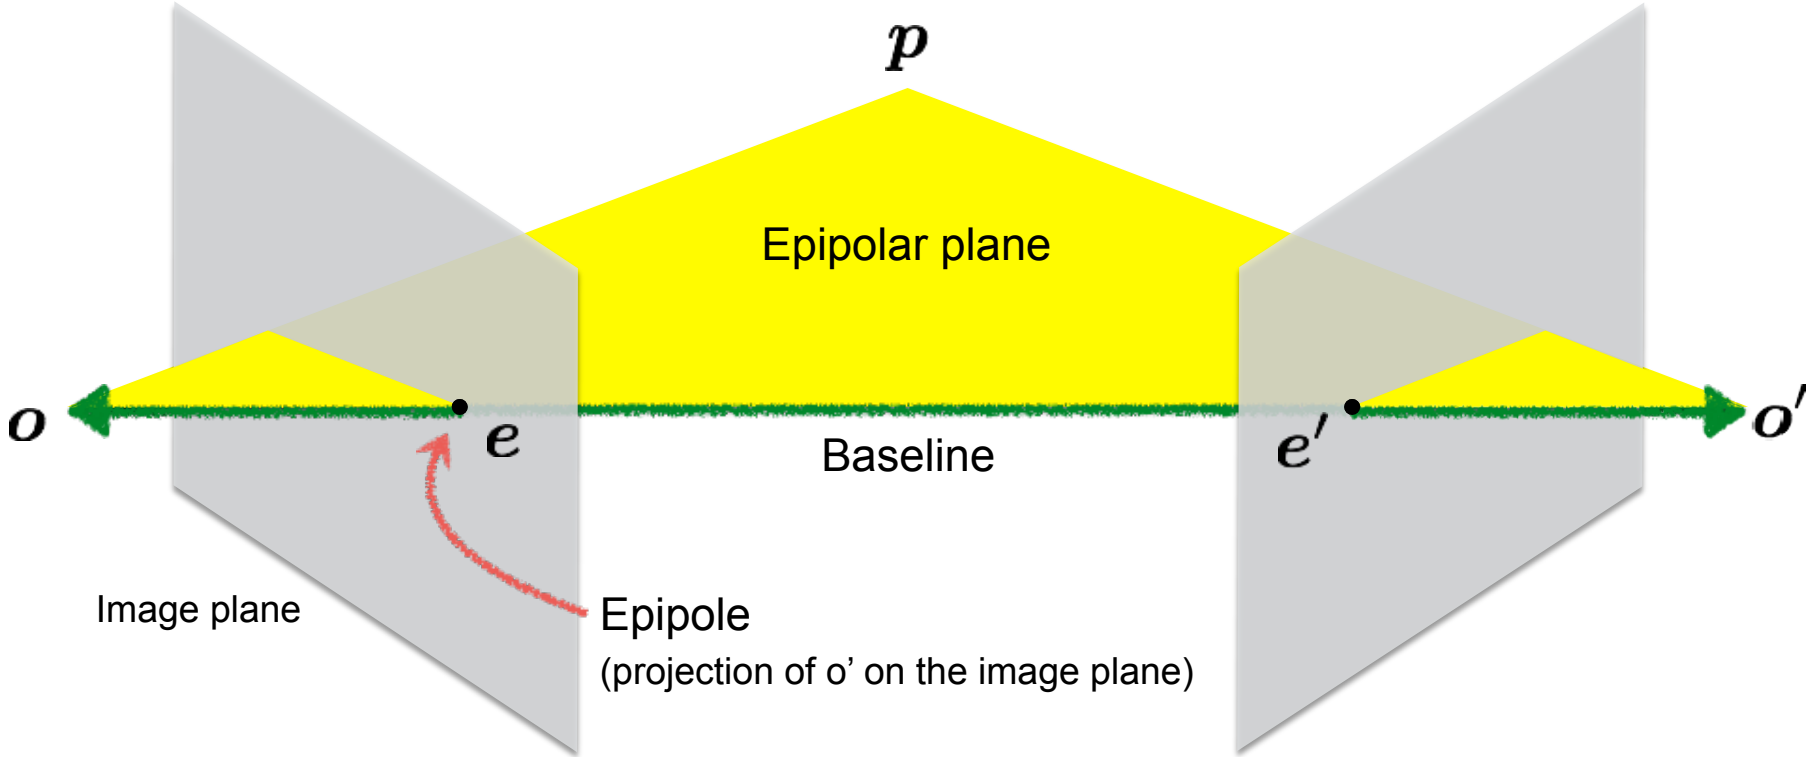
Epipolar geometry



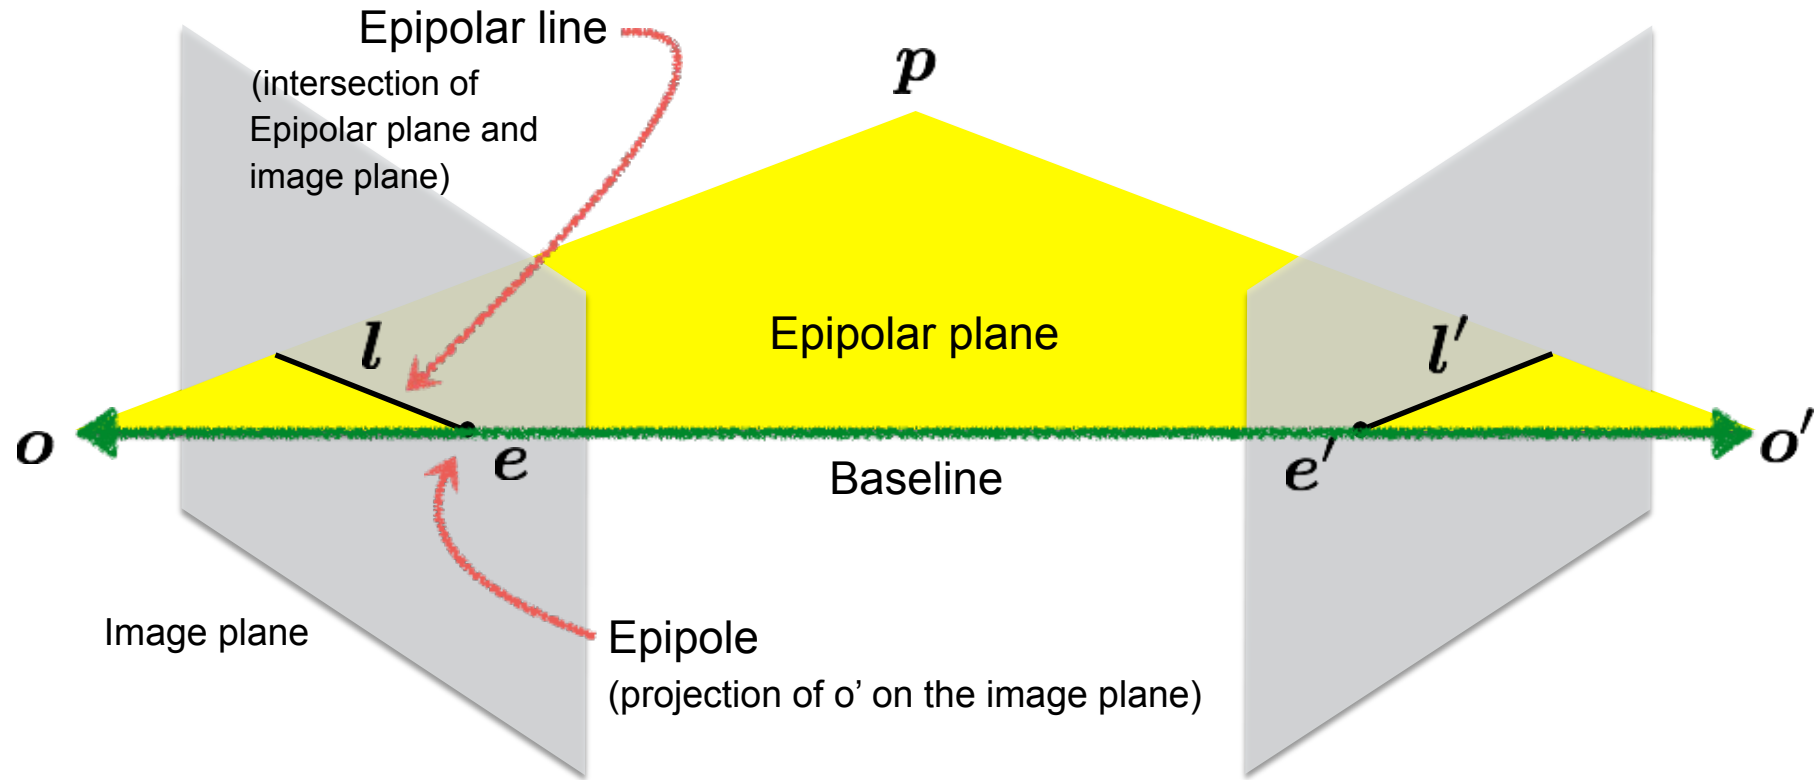
Epipolar geometry



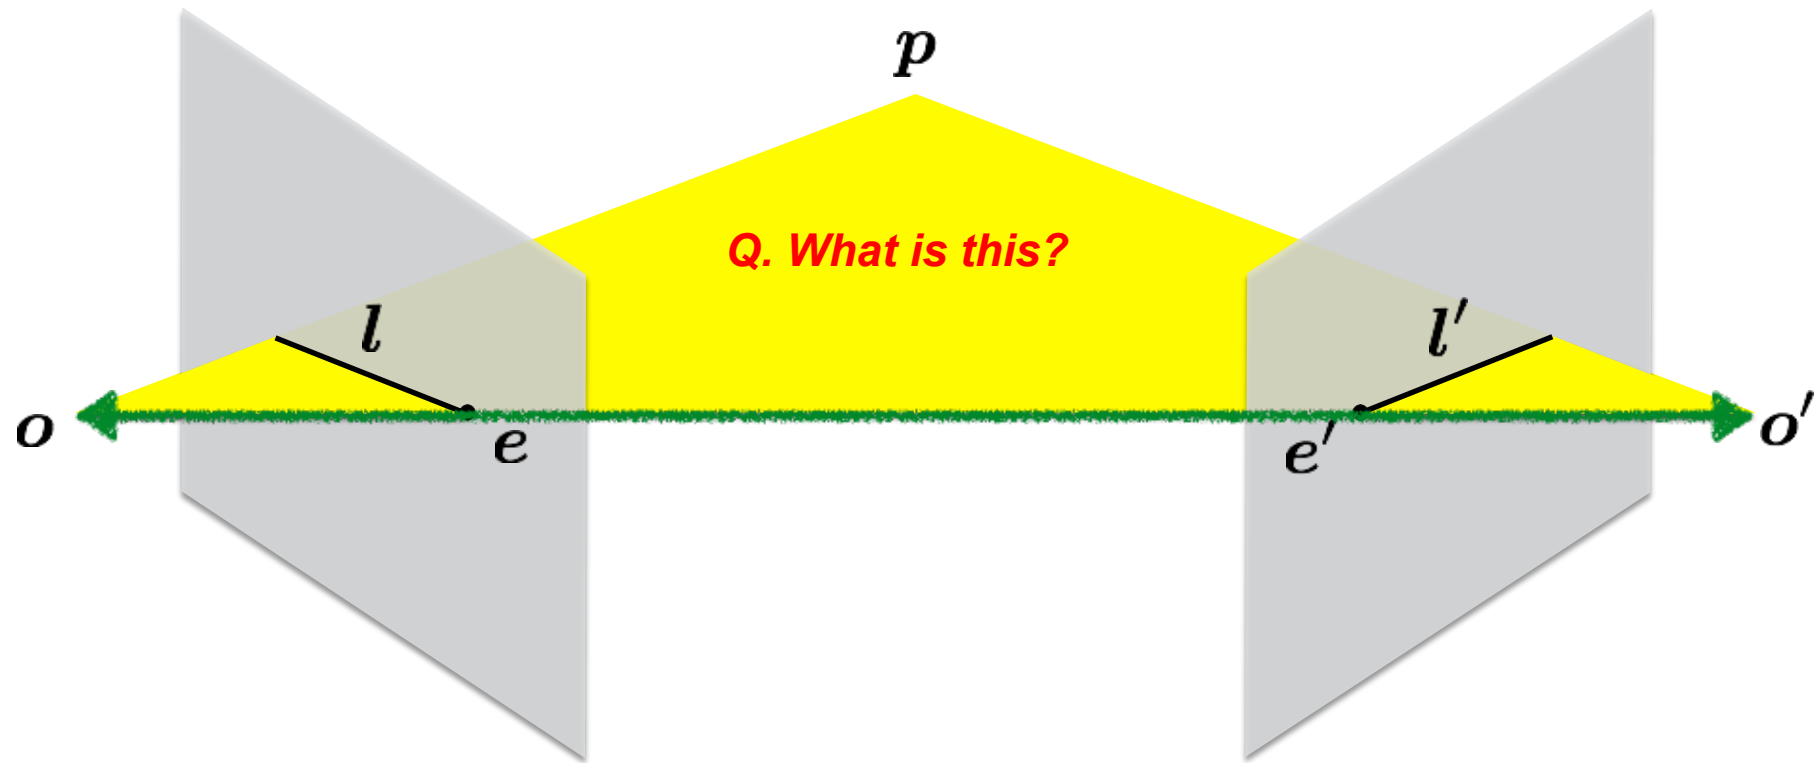
Epipolar geometry




Epipolar geometry




Let's see if you remember:




Let's see if you remember:




Let's see if you remember:




Let's see if you remember:




Let's see if you remember:




Let's see if you remember:



Let's see if you remember:




Let's see if you remember:




Epipolar constraint

Backproject \mathbf{x}
to a ray in 3D




Another way to construct the epipolar plane, this time given \mathbf{x}

Epipolar constraint






Potential matches for x lie on the epipolar line l'

Questions





1. The point \mathbf{x} (left image) maps to _____ in the right image
 2. The baseline connects the _____ and _____
 3. The epipolar line (left image) maps to a _____ in the right image
 4. An epipole \mathbf{e} is a projection of the _____ on the image plane
 5. All epipolar lines in an image intersect at the _____ on the image plane
-

Converging cameras



Where is the epipole in this image?


Converging cameras




here!

It's not always in the image

Where is the epipole when the epipolar lines are parallel?




The epipoles can be at infinity




epipole at infinity

The epipolar constraint is an important concept for stereo vision

Task: Match point in left image to point in right image



Left image




Right image


How would you do it?

The epipolar constraint is an important concept for stereo vision

Task: Match point in left image to point in right image




Left image



Right image

How would you do it using epipolar geometry?


Let's use the epipolar constraint



Potential matches for x lie on the epipolar line l'

How do you compute the epipolar line?

Task: Match point in left image to point in right image



Left image Right image


Want to avoid search over entire image

Epipolar constraint reduces search to a single line

Today's agenda

- Triangulation
- Epipolar geometry
- **Essential matrix**
- Fundamental matrix
- Structure from motion

Given a point in one image, multiplying by the **essential matrix** will tell us the **epipolar line** in the second view.




Motivation

The Essential Matrix is a 3×3 matrix that encodes **epipolar geometry**

Given a point in one image, multiplying by the **essential matrix** will tell us the **epipolar line** in the second image.

Epipolar Line

$$ax + by + c = 0 \quad \text{in vector form} \quad \mathbf{l} = \begin{bmatrix} a \\ b \\ c \end{bmatrix}$$



If the point \mathbf{x} is on the epipolar line \mathbf{l} then


$$\mathbf{x}^\top \mathbf{l} = ?$$

Epipolar Line

$$ax + by + c = 0$$

in vector form

$$\mathbf{l} = \begin{bmatrix} a \\ b \\ c \end{bmatrix}$$




If the point \mathbf{x} is on the epipolar line \mathbf{l} then

$$\mathbf{x}^\top \mathbf{l} = 0$$


So if $\mathbf{x}'^\top \mathbf{l}' = 0$ and $\mathbf{E}\mathbf{x} = \mathbf{l}'$ then

$$\mathbf{x}'^\top \mathbf{E}\mathbf{x} = ?$$



So if $\mathbf{x}'^\top \mathbf{l}' = 0$ and $\mathbf{E}\mathbf{x} = \mathbf{l}'$ then

$$\mathbf{x}'^\top \mathbf{E}\mathbf{x} = 0$$



Essential Matrix vs Homography

What's the difference between the essential matrix and a homography?

They are both 3 x 3 matrices but ...


$$l' = \mathbf{E}x$$

Essential matrix maps a
point to a line

$$x' = \mathbf{H}x$$


Homography maps a
point to a point

Where does the essential matrix come from?




$$x' = \mathbf{R}(x - t)$$

Camera-camera transformation is just like world-camera transformation




$$x' = \mathbf{R}(x - t)$$




These three vectors are coplanar

$$x, t, x'$$



If these three vectors are coplanar $\mathbf{x}, \mathbf{t}, \mathbf{x}'$ then

$$\mathbf{x}^\top (\mathbf{t} \times \mathbf{x}) = ?$$




If these three vectors are coplanar $\mathbf{x}, \mathbf{t}, \mathbf{x}'$ then

dot product of
orthogonal vectors

$$\mathbf{x}^\top (\mathbf{t} \times \mathbf{x}) = 0$$


cross-product: vector
orthogonal to plane

What is $(\mathbf{x} - \mathbf{t})$? It is another line on the epipolar plane



If these three vectors are coplanar $\mathbf{x}, \mathbf{t}, \mathbf{x}'$ then

$$(\mathbf{x} - \mathbf{t})^\top (\mathbf{t} \times \mathbf{x}) = ?$$



If these three vectors are coplanar $\mathbf{x}, \mathbf{t}, \mathbf{x}'$ then

$$(\mathbf{x} - \mathbf{t})^\top (\mathbf{t} \times \mathbf{x}) = 0$$

putting it together

$$\mathbf{x}' = \mathbf{R}(\mathbf{x} - \mathbf{t})$$

$$(\mathbf{x} - \mathbf{t})^\top (\mathbf{t} \times \mathbf{x}) = 0$$

Substituting $(\mathbf{x}-\mathbf{t})$:

$$\mathbf{x}' = \mathbf{R}(\mathbf{x} - \mathbf{t}) \quad (\mathbf{x} - \mathbf{t})^\top (\mathbf{t} \times \mathbf{x}) = 0$$

$$(\mathbf{x}'^\top \mathbf{R})(\mathbf{t} \times \mathbf{x}) = 0$$

Cross product reminder:

$$\mathbf{a} \times \mathbf{b} = \begin{bmatrix} a_2b_3 - a_3b_2 \\ a_3b_1 - a_1b_3 \\ a_1b_2 - a_2b_1 \end{bmatrix}$$

Can also be written as a matrix multiplication

$$\mathbf{a} \times \mathbf{b} = [\mathbf{a}]_{\times} \mathbf{b} = \begin{bmatrix} 0 & -a_3 & a_2 \\ a_3 & 0 & -a_1 \\ -a_2 & a_1 & 0 \end{bmatrix} \begin{bmatrix} b_1 \\ b_2 \\ b_3 \end{bmatrix}$$

Skew symmetric

Use the skew-symmetric matrix to represent the cross product

$$\mathbf{x}' = \mathbf{R}(\mathbf{x} - \mathbf{t}) \quad (\mathbf{x} - \mathbf{t})^\top (\mathbf{t} \times \mathbf{x}) = 0$$

$$(\mathbf{x}'^\top \mathbf{R})(\mathbf{t} \times \mathbf{x}) = 0$$

$$(\mathbf{x}'^\top \mathbf{R})([\mathbf{t}_\times] \mathbf{x}) = 0$$

Use the skew-symmetric matrix to represent the cross product

$$\mathbf{x}' = \mathbf{R}(\mathbf{x} - \mathbf{t}) \quad (\mathbf{x} - \mathbf{t})^\top (\mathbf{t} \times \mathbf{x}) = 0$$

$$(\mathbf{x}'^\top \mathbf{R})(\mathbf{t} \times \mathbf{x}) = 0$$

$$(\mathbf{x}'^\top \mathbf{R})([\mathbf{t}_\times] \mathbf{x}) = 0$$

$$\mathbf{x}'^\top (\mathbf{R}[\mathbf{t}_\times]) \mathbf{x} = 0$$

This is the essential matrix

$$\mathbf{x}' = \mathbf{R}(\mathbf{x} - \mathbf{t}) \quad (\mathbf{x} - \mathbf{t})^\top (\mathbf{t} \times \mathbf{x}) = 0$$

$$(\mathbf{x}'^\top \mathbf{R})(\mathbf{t} \times \mathbf{x}) = 0$$

$$(\mathbf{x}'^\top \mathbf{R})([\mathbf{t}_\times] \mathbf{x}) = 0$$

$$\mathbf{x}'^\top (\mathbf{R}[\mathbf{t}_\times]) \mathbf{x} = 0$$

$$\mathbf{x}'^\top \mathbf{E} \mathbf{x} = 0$$

Essential Matrix
[Longuet-Higgins 1981]

Summary

Longuet-Higgins equation

$$\mathbf{x}'^{\top} \mathbf{E} \mathbf{x} = 0$$

Epipolar lines

$$\mathbf{x}^{\top} \mathbf{l} = 0$$

$$\mathbf{x}'^{\top} \mathbf{l}' = 0$$


$$\mathbf{l}' = \mathbf{E} \mathbf{x}$$

$$\mathbf{l} = \mathbf{E}^{\top} \mathbf{x}'$$

$$\mathbf{e}'^{\top} \mathbf{E} = 0$$

$$\mathbf{E} \mathbf{e} = 0$$

Everything we have done so far assumes:
we have camera coordinates of pixels



Today's agenda

- Triangulation
- Epipolar geometry
- Essential matrix
- **Fundamental matrix**
- Structure from motion

$$\hat{\mathbf{x}}'^{\top} \mathbf{E} \hat{\mathbf{x}} = 0$$

The essential matrix operates on **2D points coordinates** in the camera coordinate system

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

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

camera point image point

$$\hat{\mathbf{x}}'^{\top} \mathbf{E} \hat{\mathbf{x}} = 0$$

The essential matrix operates on **2D points coordinates** in the camera coordinate system

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

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

camera point image point

Writing out the epipolar constraint in terms of image coordinates

$$\mathbf{x}'^{\top} (\mathbf{K}'^{-\top} \mathbf{E} \mathbf{K}^{-1}) \mathbf{x} = 0$$

$$\mathbf{x}'^{\top} \mathbf{F} \mathbf{x} = 0$$

Same equation works in image coordinates!

$$\mathbf{x}'^{\top} \mathbf{F} \mathbf{x} = 0$$

it maps pixels to epipolar lines



Longuet-Higgins equation $x'^{\top} \mathbf{E} x = 0$

Epipolar lines $x^{\top} l = 0$ $x'^{\top} l' = 0$
 $l' = \mathbf{E} x$ $l = \mathbf{E}^T x'$

Epipoles $e'^{\top} \mathbf{E} = 0$ $\mathbf{E} e = 0$

(points in **image** coordinates)

Breaking down the fundamental matrix

$$\mathbf{F} = \mathbf{K}'^{-\top} \mathbf{E} \mathbf{K}^{-1}$$

$$\mathbf{F} = \mathbf{K}'^{-\top} [\mathbf{t}_x] \mathbf{R} \mathbf{K}^{-1}$$

Depends on both intrinsic and extrinsic parameters

The 8-point algorithm solves for F given a list of corresponding points (x, x')

Assume you have M matched *image* points

$$\{\mathbf{x}_m, \mathbf{x}'_m\} \quad m = 1, \dots, M$$

Each correspondence should satisfy

$$\mathbf{x}'_m{}^\top \mathbf{F} \mathbf{x}_m = 0$$

Each corresponding set of points (x, x')
will give us **one equation**

$$\mathbf{x}'_m{}^\top \mathbf{F} \mathbf{x}_m = 0$$

$$\begin{bmatrix} x'_m & y'_m & 1 \end{bmatrix} \begin{bmatrix} f_1 & f_2 & f_3 \\ f_4 & f_5 & f_6 \\ f_7 & f_8 & f_9 \end{bmatrix} \begin{bmatrix} x_m \\ y_m \\ 1 \end{bmatrix} = 0$$

Each corresponding set of points (x, x')
will give us **one equation**

$$\mathbf{x}'_m{}^\top \mathbf{F} \mathbf{x}_m = 0$$

$$\begin{bmatrix} x'_m & y'_m & 1 \end{bmatrix} \begin{bmatrix} f_1 & f_2 & f_3 \\ f_4 & f_5 & f_6 \\ f_7 & f_8 & f_9 \end{bmatrix} \begin{bmatrix} x_m \\ y_m \\ 1 \end{bmatrix} = 0$$


$$\begin{aligned} &x_m x'_m f_1 + x_m y'_m f_2 + x_m f_3 + \\ &y_m x'_m f_4 + y_m y'_m f_5 + y_m f_6 + \\ &x'_m f_7 + y'_m f_8 + f_9 = 0 \end{aligned}$$

Like always, we can re-write it as a linear equation with M corresponding pairs of points:

$$\begin{bmatrix} x_1x'_1 & x_1y'_1 & x_1 & y_1x'_1 & y_1y'_1 & y_1 & x'_1 & y'_1 & 1 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ x_Mx'_M & x_My'_M & x_M & y_Mx'_M & y_My'_M & y_M & x'_M & y'_M & 1 \end{bmatrix} \begin{bmatrix} f_1 \\ f_2 \\ f_3 \\ f_4 \\ f_5 \\ f_6 \\ f_7 \\ f_8 \\ f_9 \end{bmatrix} = \mathbf{0}$$


Solve using SVD!

You can find correspondences using Harris + RANSAC




You can use the corresponding points to calculate F

Once you have F , you can map points to epipolar lines:



Here are a bunch of epipolar lines across these two images



Today's agenda

- Triangulation
- Epipolar geometry
- Essential matrix
- Fundamental matrix
- **Structure from motion**

Structure-from-Motion

Given many images, how can we

- a) figure out where they were all taken from?
- b) build a 3D model of the scene?



N. Snavely, S. Seitz, and R. Szeliski, [Photo tourism: Exploring photo collections in 3D](http://phototour.cs.washington.edu/), SIGGRAPH 2006.
<http://phototour.cs.washington.edu/>

Large-scale structure-from-motion



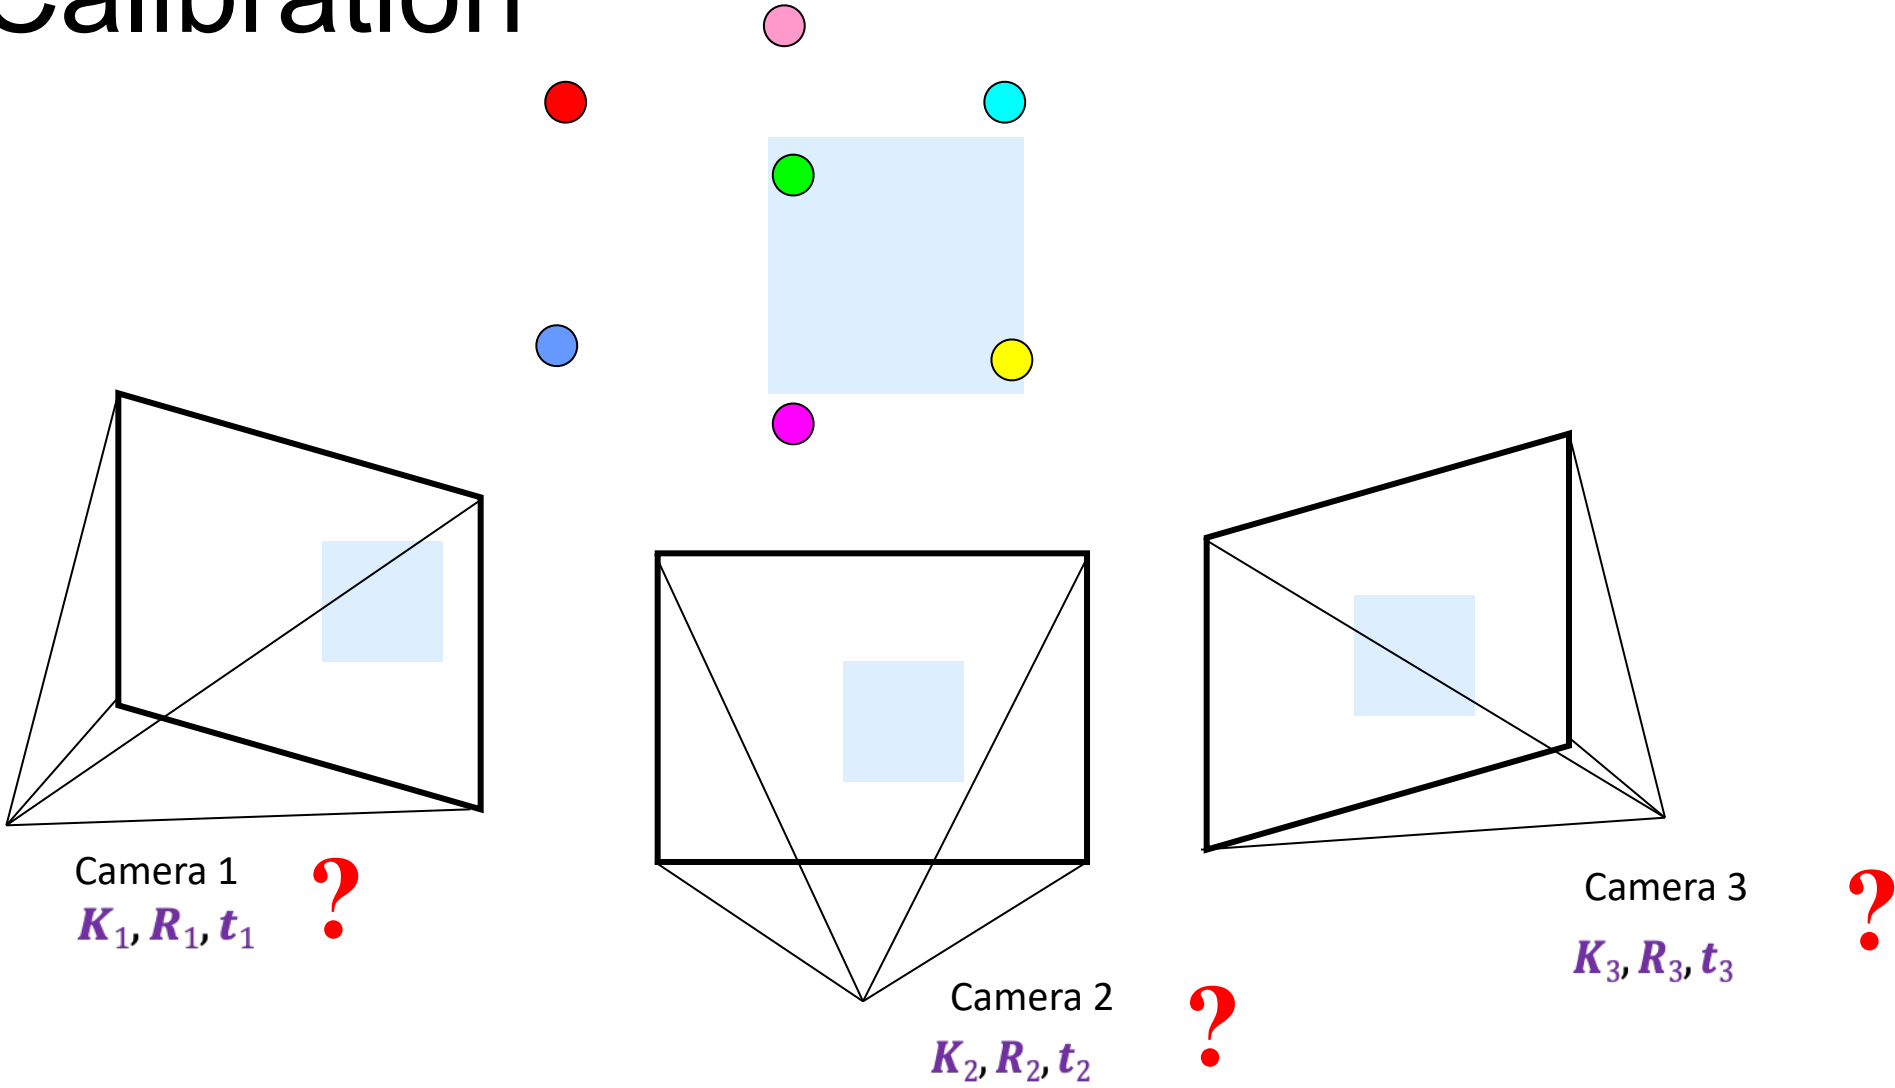
Dubrovnik, Croatia. 4,619 images (out of an initial 57,845 downloaded from Flickr). 3.5M points!

Total reconstruction time: 17.5 hours on 352 cores

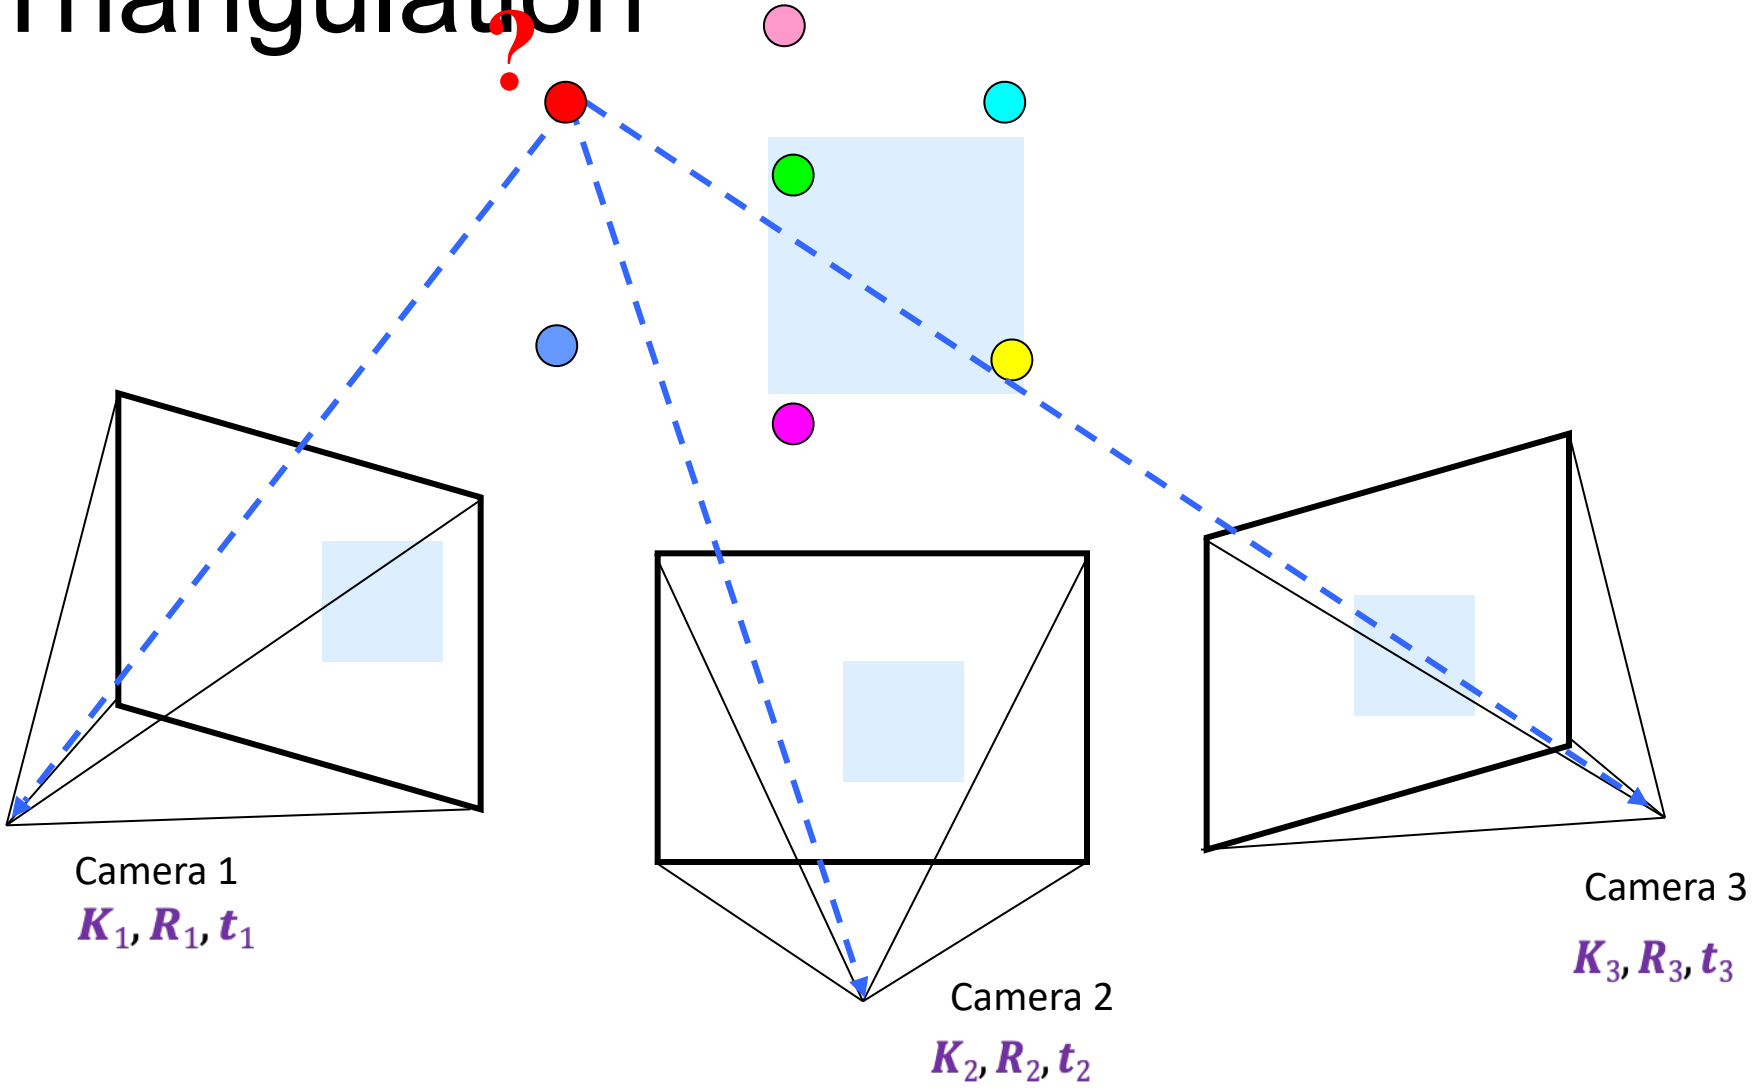
Building Rome in a Day, Agarwal et al, ICCV'09

<http://grail.cs.washington.edu/rome/>

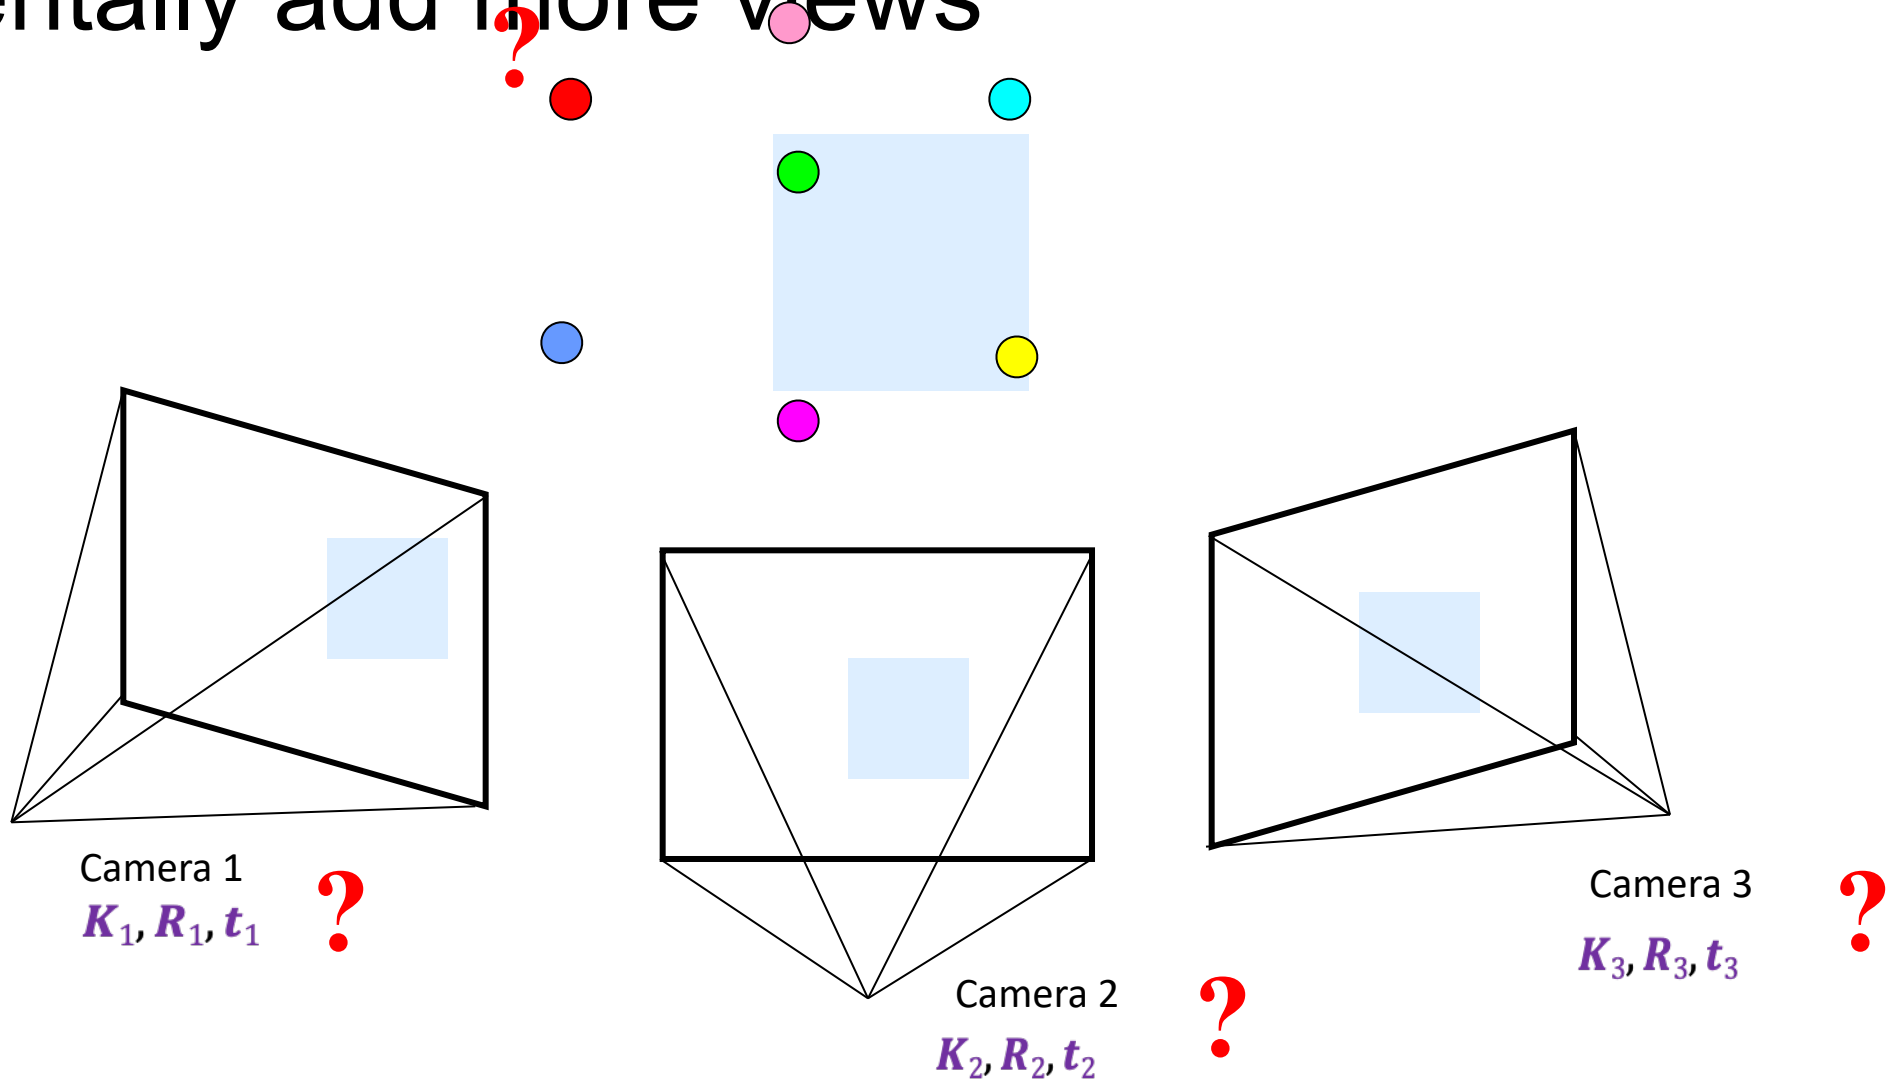
First: Calibration




Next: Triangulation




That's it: you get structure-from-motion as you incrementally add more views



Structure-from-Motion




Reconstruction
(side)



(top)

Incremental Structure-from-Motion

- Estimate **motion** between two images by calculating the fundamental matrix
- Estimate 3D **structure** by triangulation
- For each additional view:
 - Determine **motion** of new camera using all the known 3D points that have correspondence in the new image
 - Add new **structure** by estimating the new points in the new image



Incremental structure from motion



Time-lapse reconstruction of Dubrovnik, Croatia, viewed from above

COLMAP

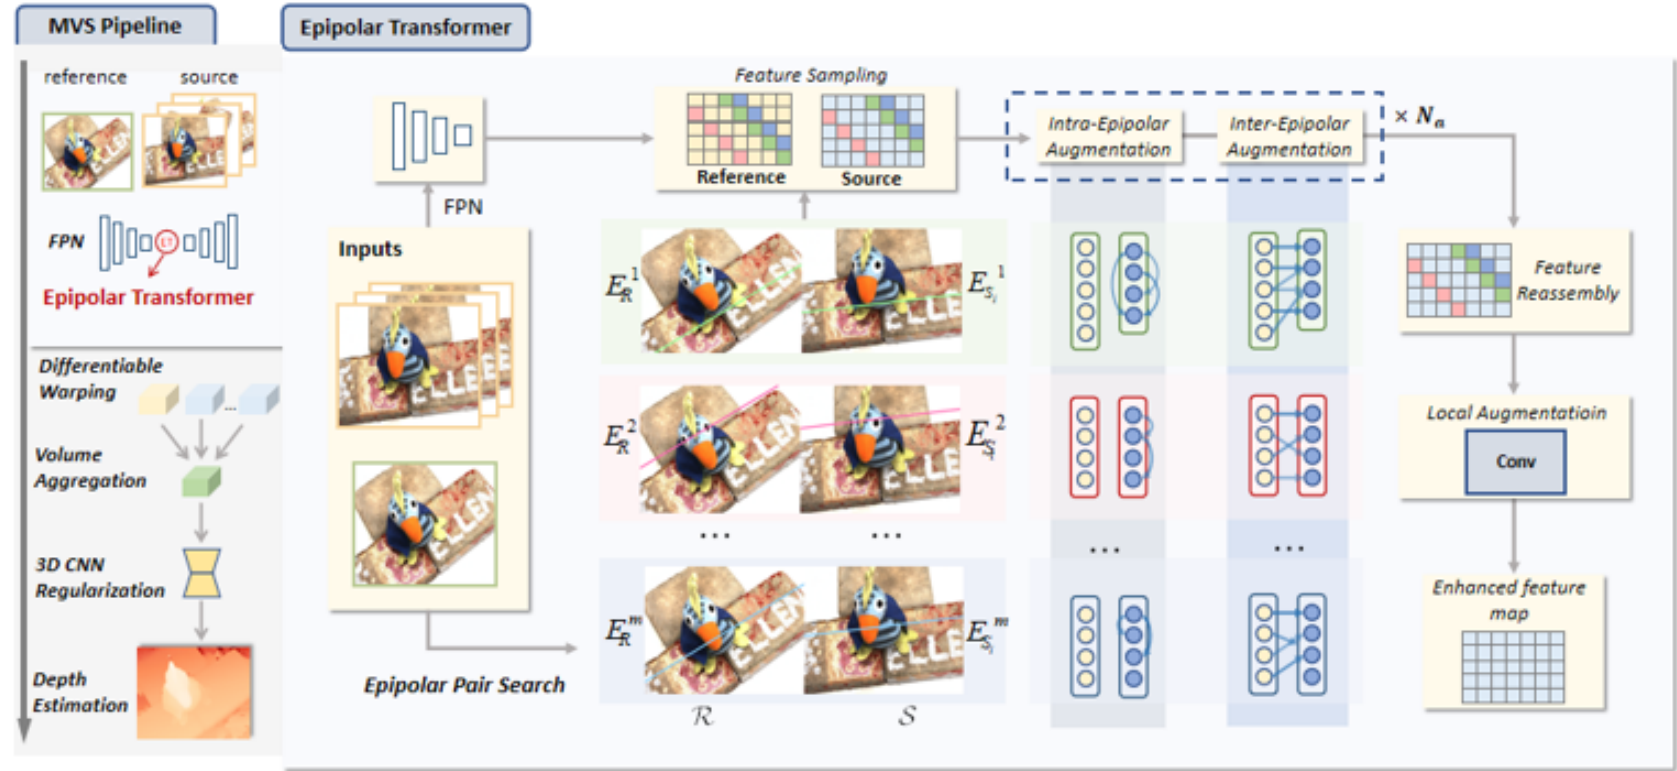
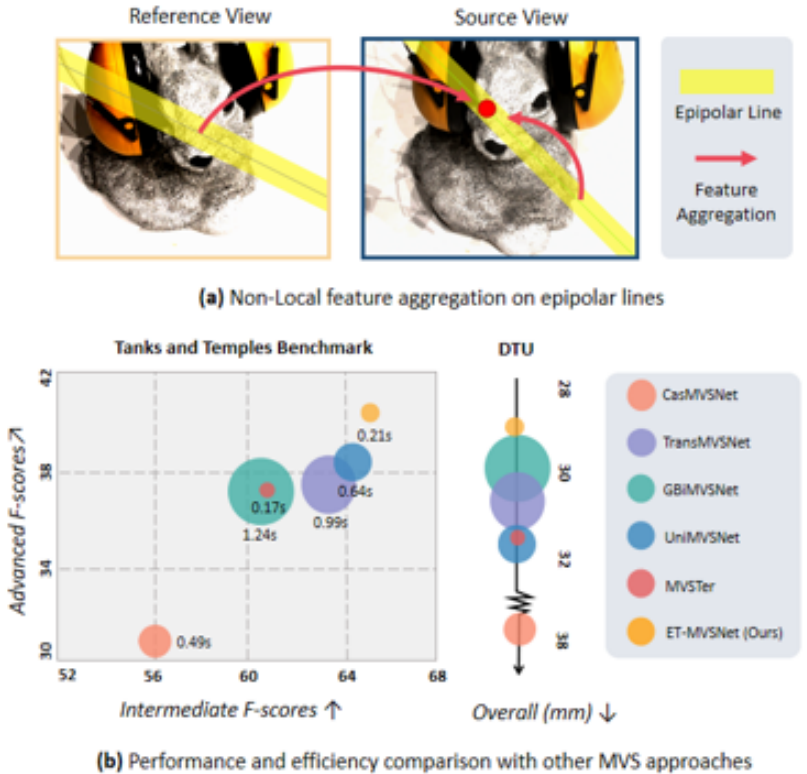


Sparse model of central Rome using 21K photos produced by COLMAP's SfM pipeline.



Dense models of several landmarks produced by COLMAP's MVS pipeline.

SfM in the age of Deep Learning



[ET-MVSNet: When Epipolar Constraint Meets Non-local Operators in Multi-View Stereo \(ICCV'23\)](#)

See also [MVSFormer: Multi-View Stereo by Learning Robust Image Features and Temperature-based Depth](#)

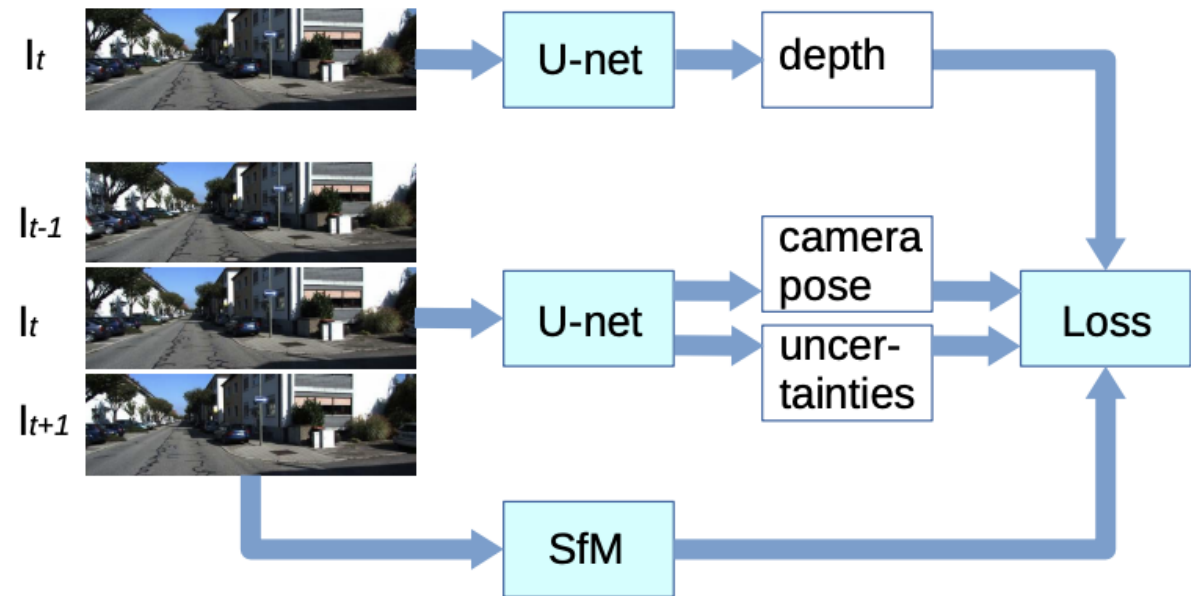
(TMLR'23)

Supervising the new with the old: learning SFM from SFM

Maria Klodt^[0000-0003-3015-9584] and Andrea Vedaldi^[0000-0003-1374-2858]

Visual Geometry Group, University of Oxford
{klodt,vedaldi}@robots.ox.ac.uk

Abstract. Recent work has demonstrated that it is possible to learn deep neural networks for monocular depth and ego-motion estimation from unlabelled video sequences, an interesting theoretical development with numerous advantages in applications. In this paper, we propose a number of improvements to these approaches. First, since such self-supervised approaches are based on the brightness constancy assumption, which is valid only for a subset of pixels, we propose a probabilistic learning formulation where the network predicts distributions over variables rather than specific values. As these distributions are conditioned on the observed image, the network can learn which scene and object types are likely to violate the model assumptions, resulting in more robust learning. We also propose to build on dozens of years of experience in developing handcrafted structure-from-motion (SfM) algorithms. We do so by using an off-the-shelf SfM system to generate a supervisory signal for the deep neural network. While this signal is also noisy, we show that our probabilistic formulation can learn and account for the defects of SfM, helping to integrate different sources of information and boosting the overall performance of the network.



(b) proposed network architecture:
the depth and pose-uncertainty networks
are supervised by traditional SfM.

https://openaccess.thecvf.com/content_ECCV_2018/papers/Maria_Klodt_Supervising_the_new_ECCV_2018_paper.pdf

Today's agenda

- Triangulation
- Epipolar geometry
- Essential matrix
- Fundamental matrix
- Structure from motion

Next lecture

The frontiers of Computer Vision