

# Lecture 18: Fitting geometric models

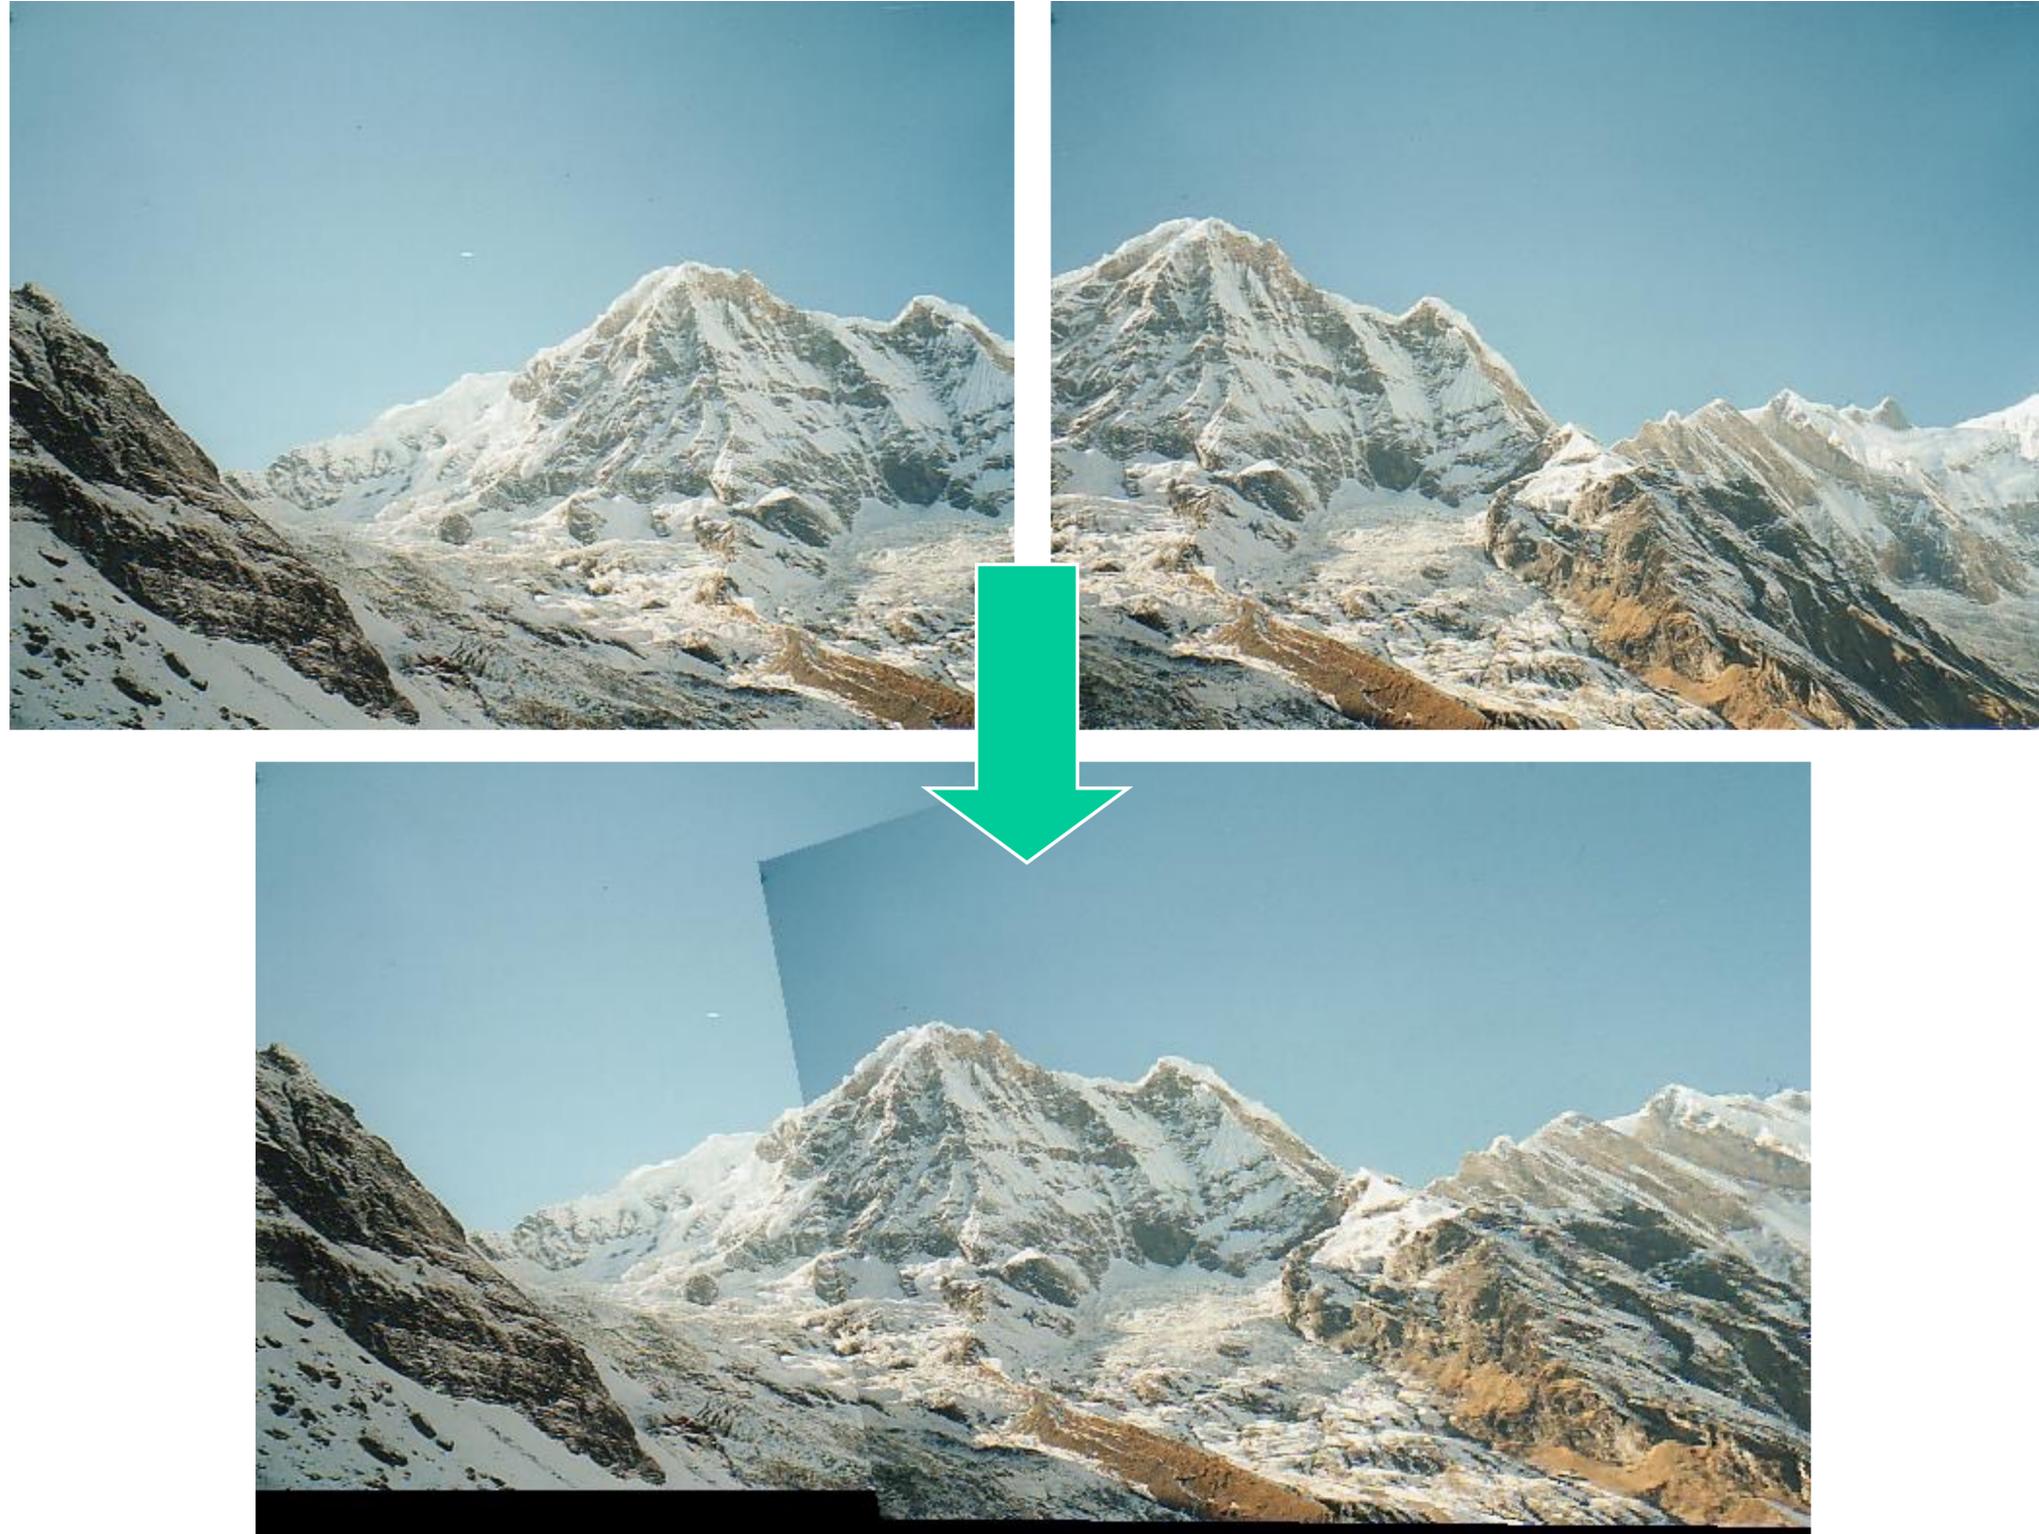
# Announcements

- Section this week: project office hours
- PS8: panorama stitching

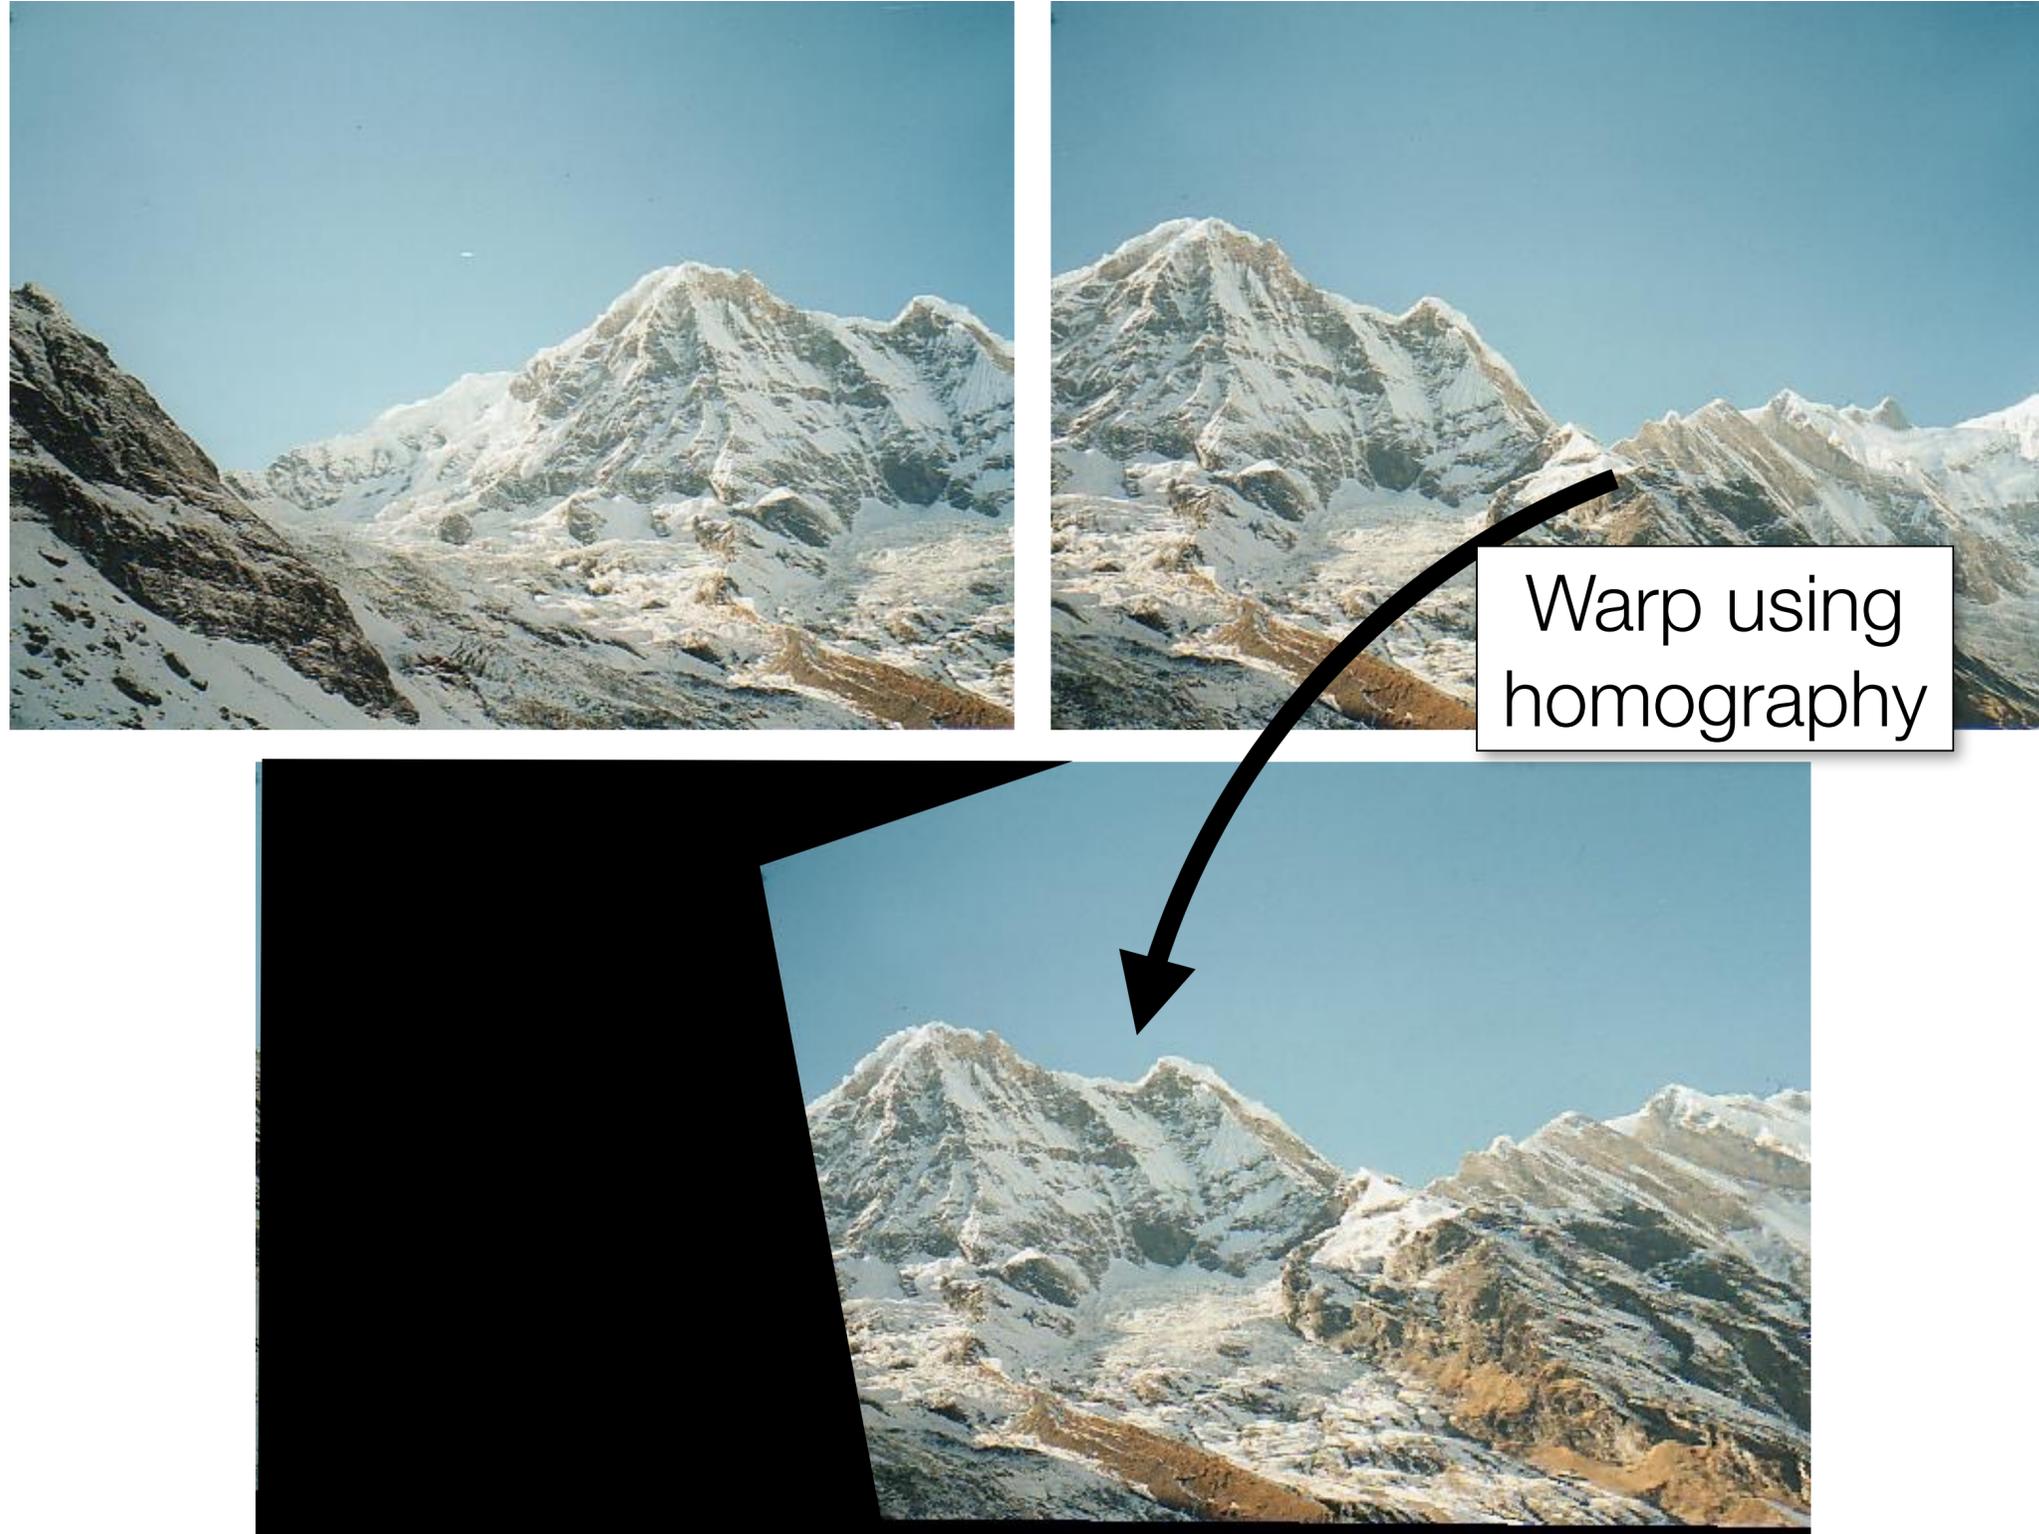
# Today

- Finding correspondences
- Fitting a homography
- RANSAC
- Triangulation

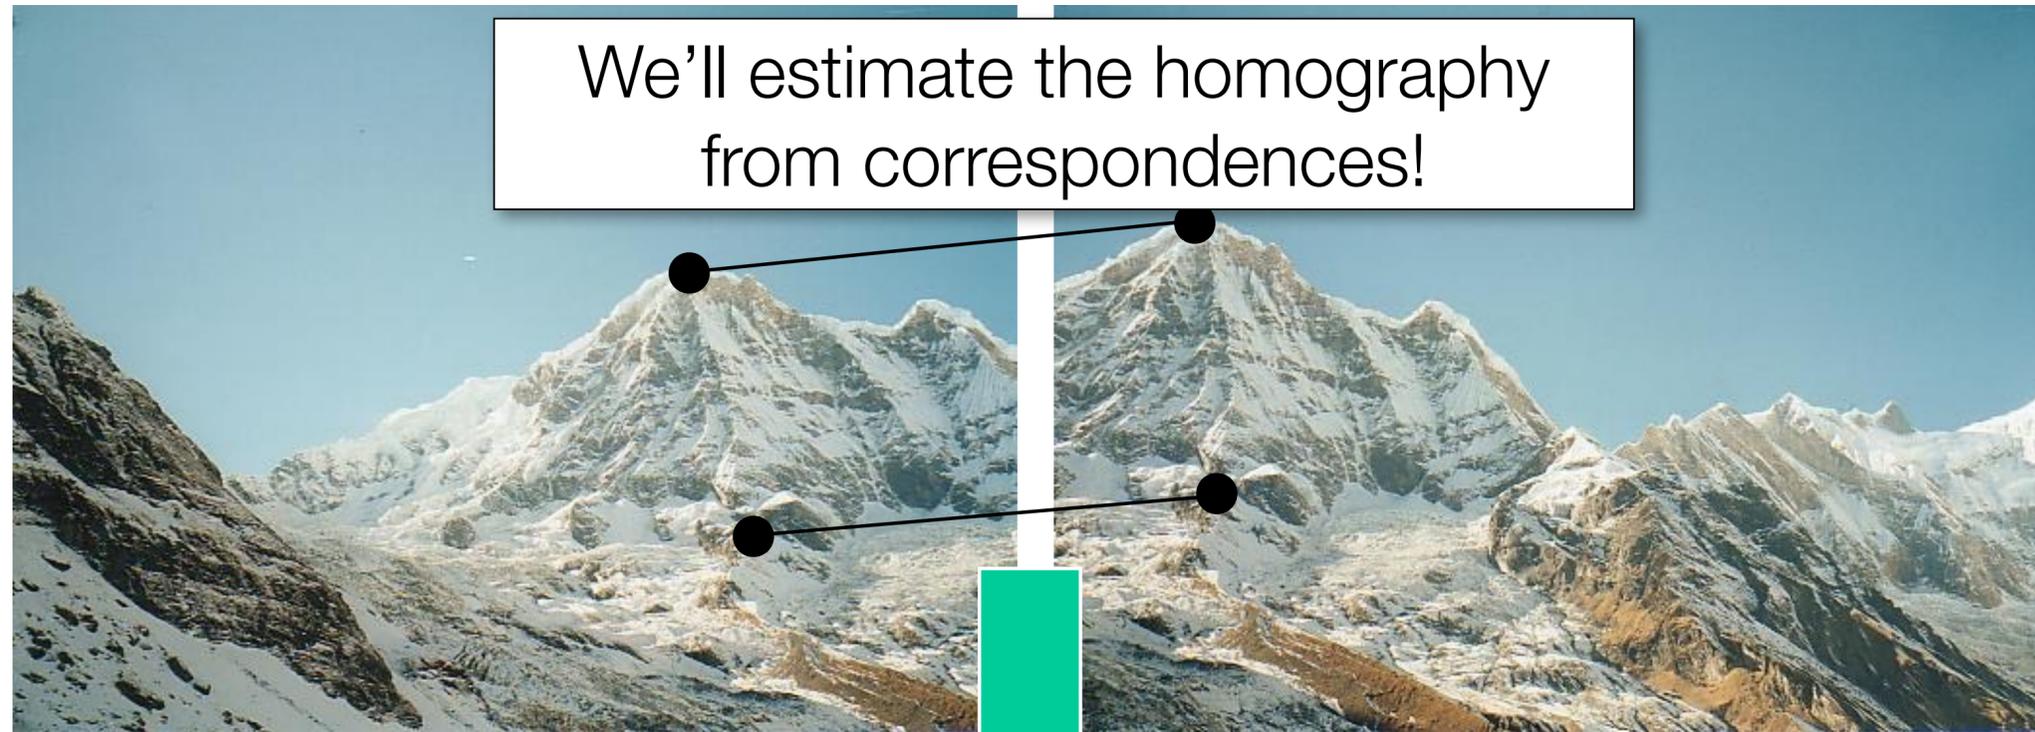
# Panorama stitching (PS9)



# Panorama stitching



# Panorama stitching



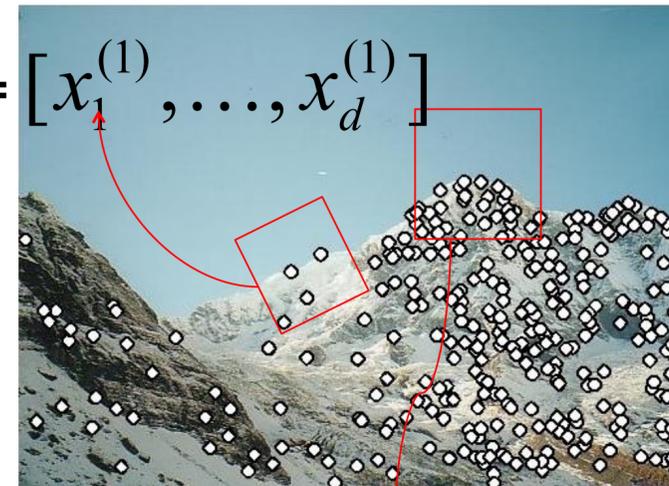
# Finding correspondences with local features

1) **Detection:** Identify the interest points (a.k.a. keypoints), the candidate points to match.



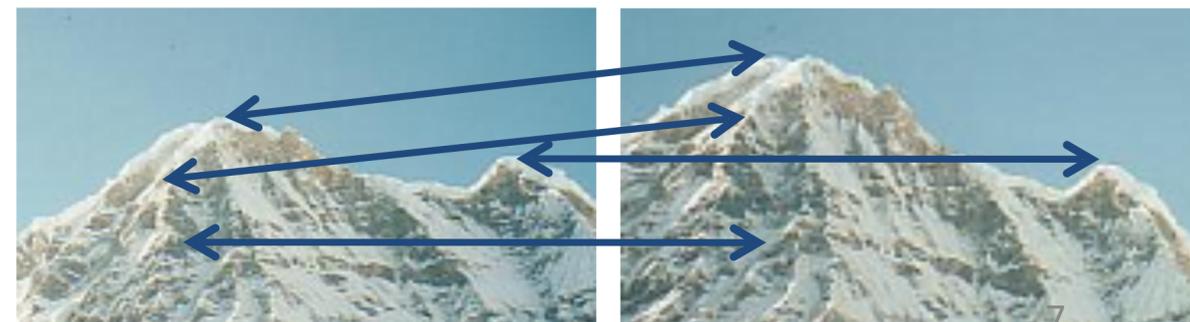
2) **Description:** Extract vector feature descriptor surrounding each interest point.

$$\mathbf{x}_1 = [x_1^{(1)}, \dots, x_d^{(1)}]$$



$$\mathbf{x}_2 = [x_1^{(2)}, \dots, x_d^{(2)}]$$

3) **Matching:** Determine correspondence between descriptors in two views

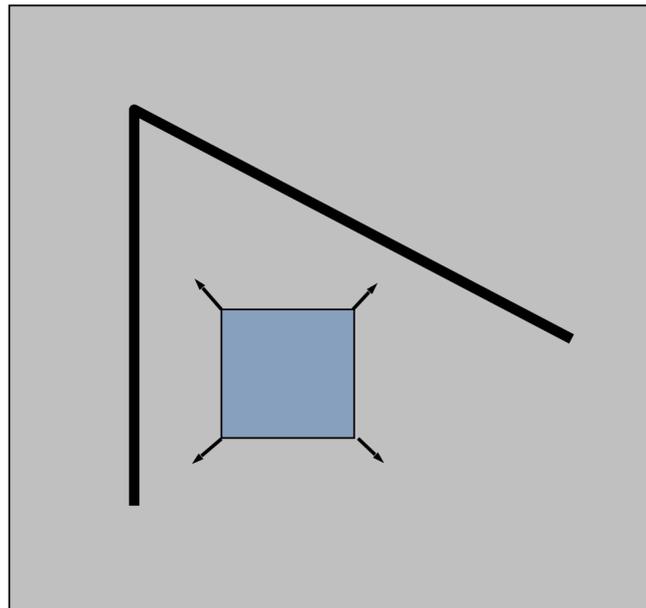


# What are good regions to match?

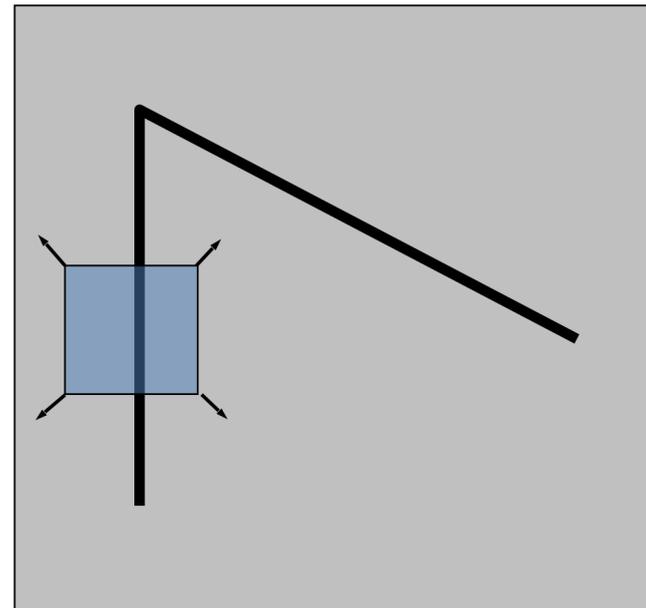


# What are good regions to match?

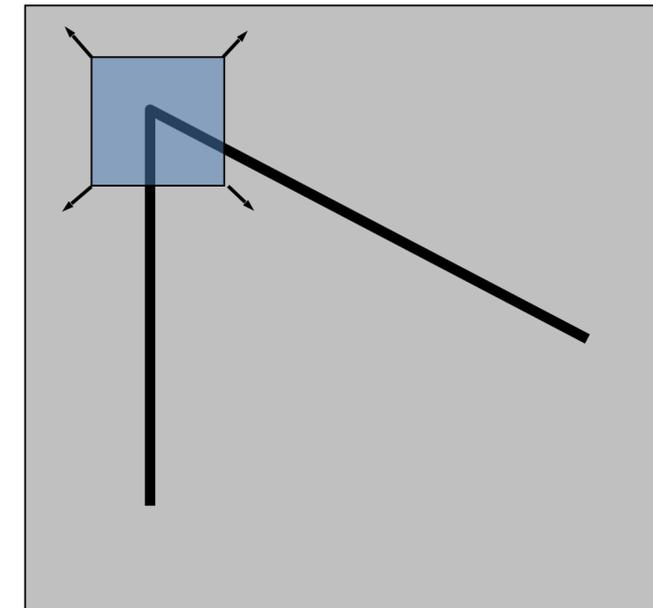
- How does the window change when you shift it?
- Shifting the window in any direction causes a big change



“flat” region:  
no change in all  
directions

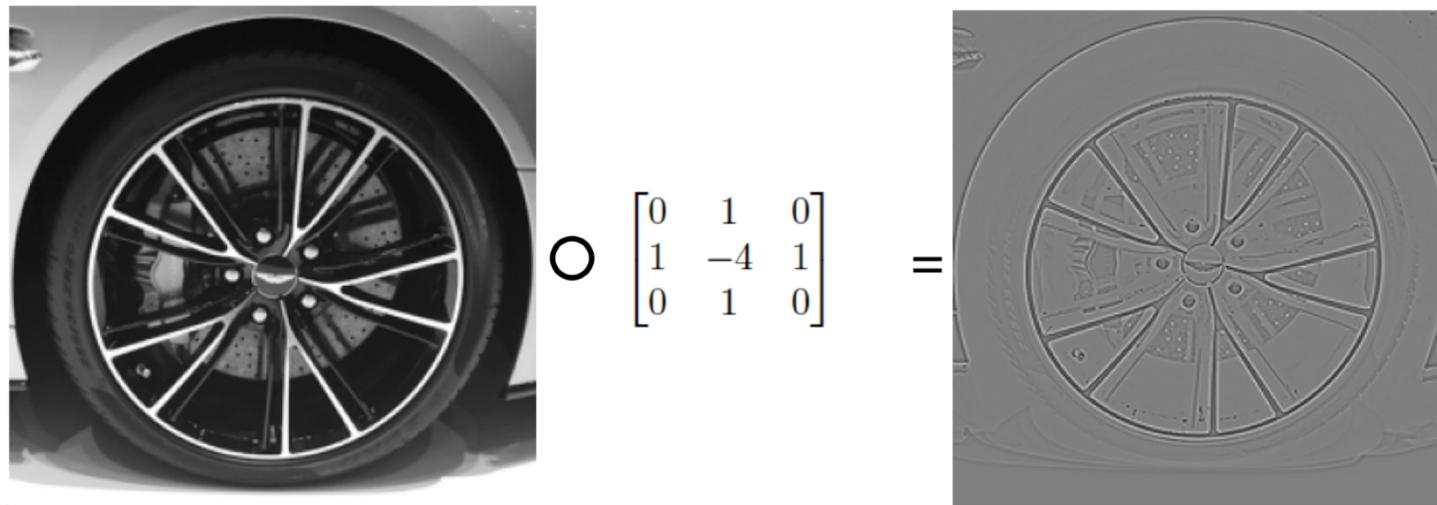


“edge”:  
no change along the  
edge direction

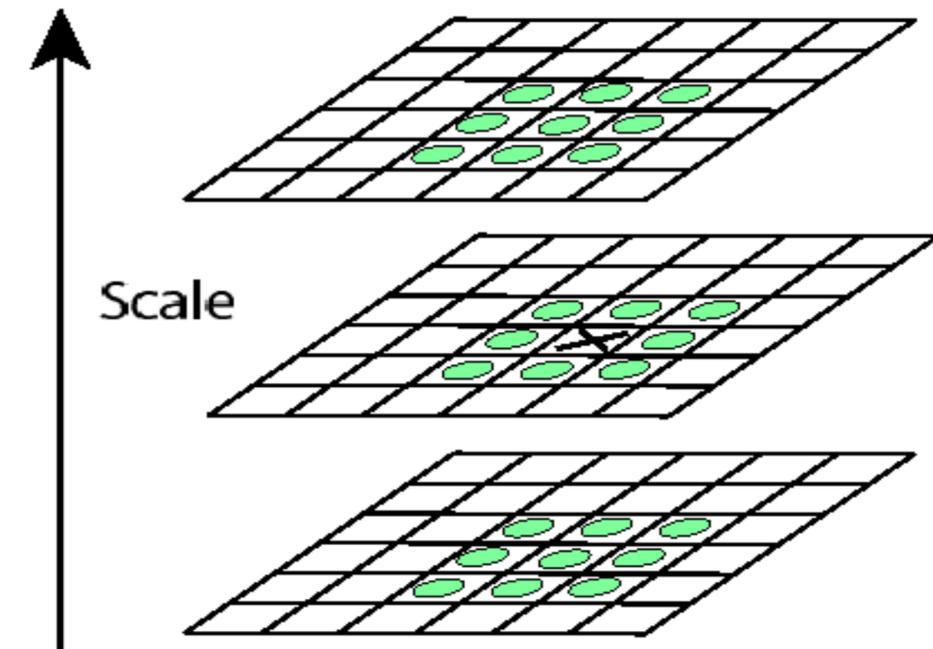


“corner”:  
significant change in  
all directions

# Finding good key points to match



Compute difference-of-Gaussians filter (approx. to Laplacian).

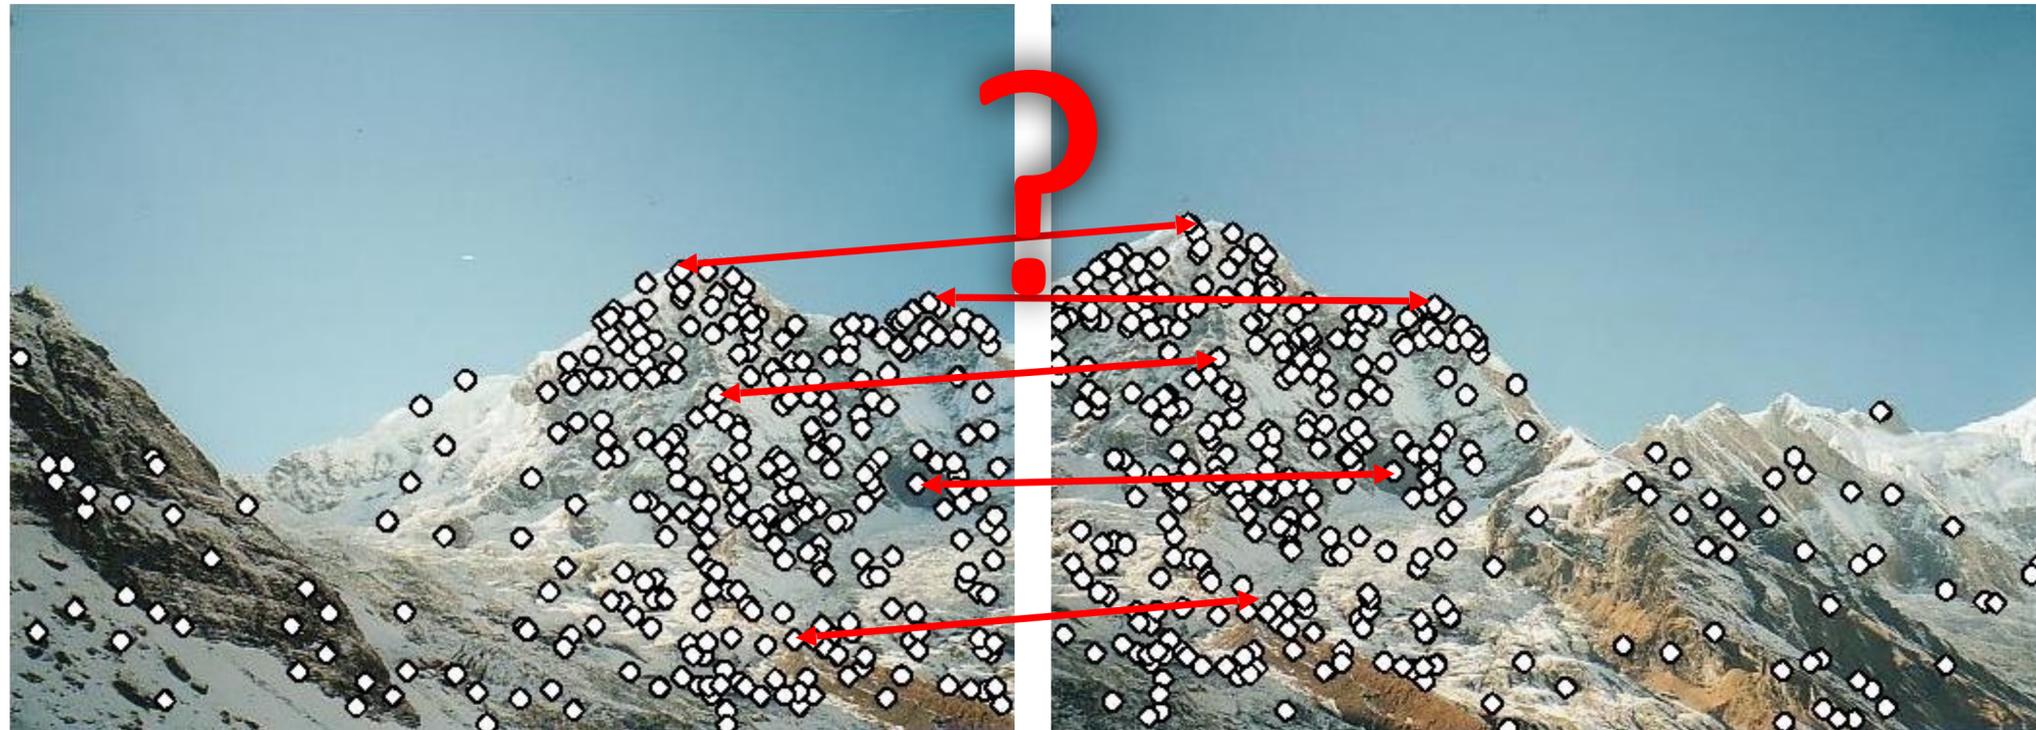


Find local optima in space and scale using Laplacian pyramid.

# Feature descriptors

We know how to detect good points

Next question: **How to match them?**



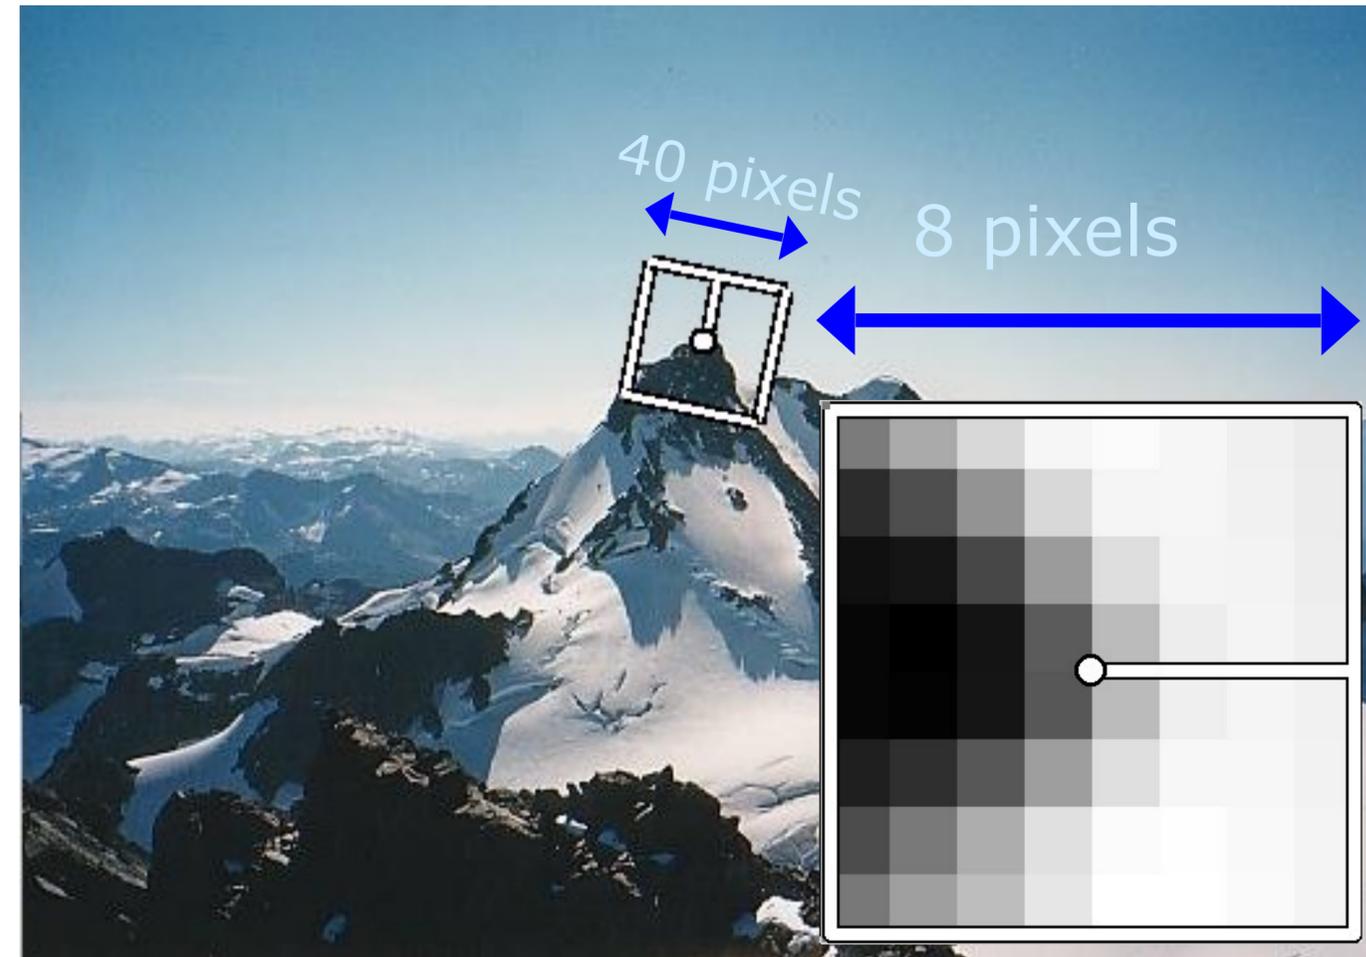
Come up with a *descriptor* (feature vector) for each point, find similar descriptors between the two images

# Simple idea: normalized image patch

We want invariance to rotation, lighting, and tiny spatial shifts.

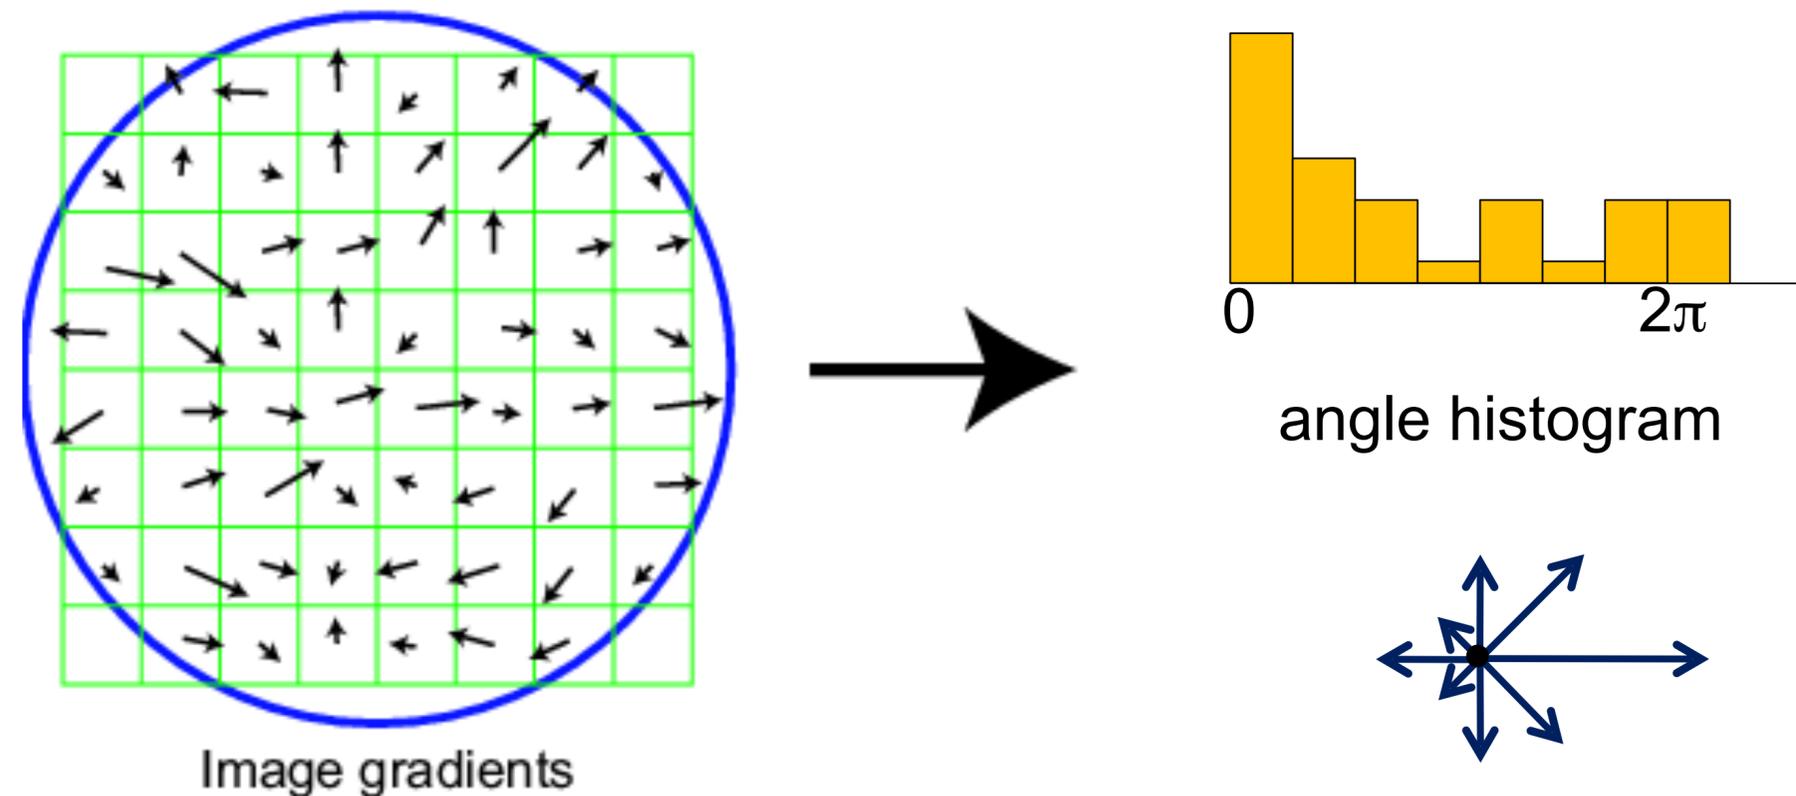
Take 40x40 window around feature

- Find dominant orientation
- Rotate to horizontal
- Downsample to 8x8
- Intensity normalize the window by subtracting the mean, dividing by the standard deviation in the window



# Scale Invariant Feature Transform (SIFT)

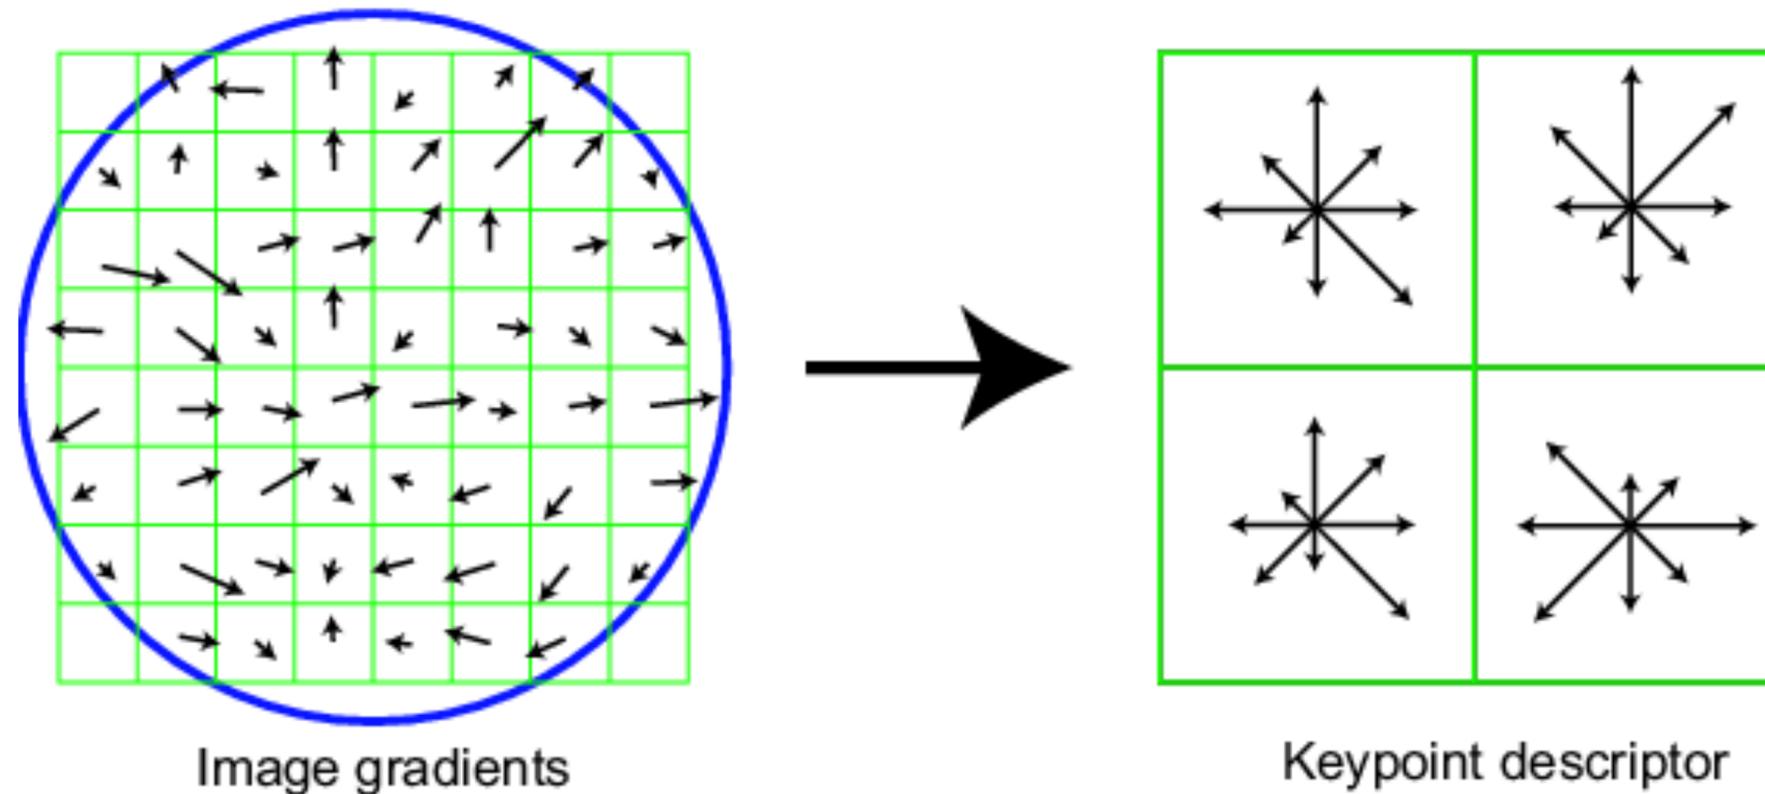
- Compute histograms of oriented gradients
- Take 16x16 square window around detected feature
- Compute edge orientation for each pixel
- Looks like a small, hand-crafted CNN



# Scale Invariant Feature Transform

Create the descriptor:

- Rotation invariance: rotate by “dominant” orientation
- Spatial invariance: spatial pool to 2x2
- Compute an orientation histogram for each cell
- $(4 \times 4)$  cells  $\times$  8 orientations = 128 dimensional descriptor



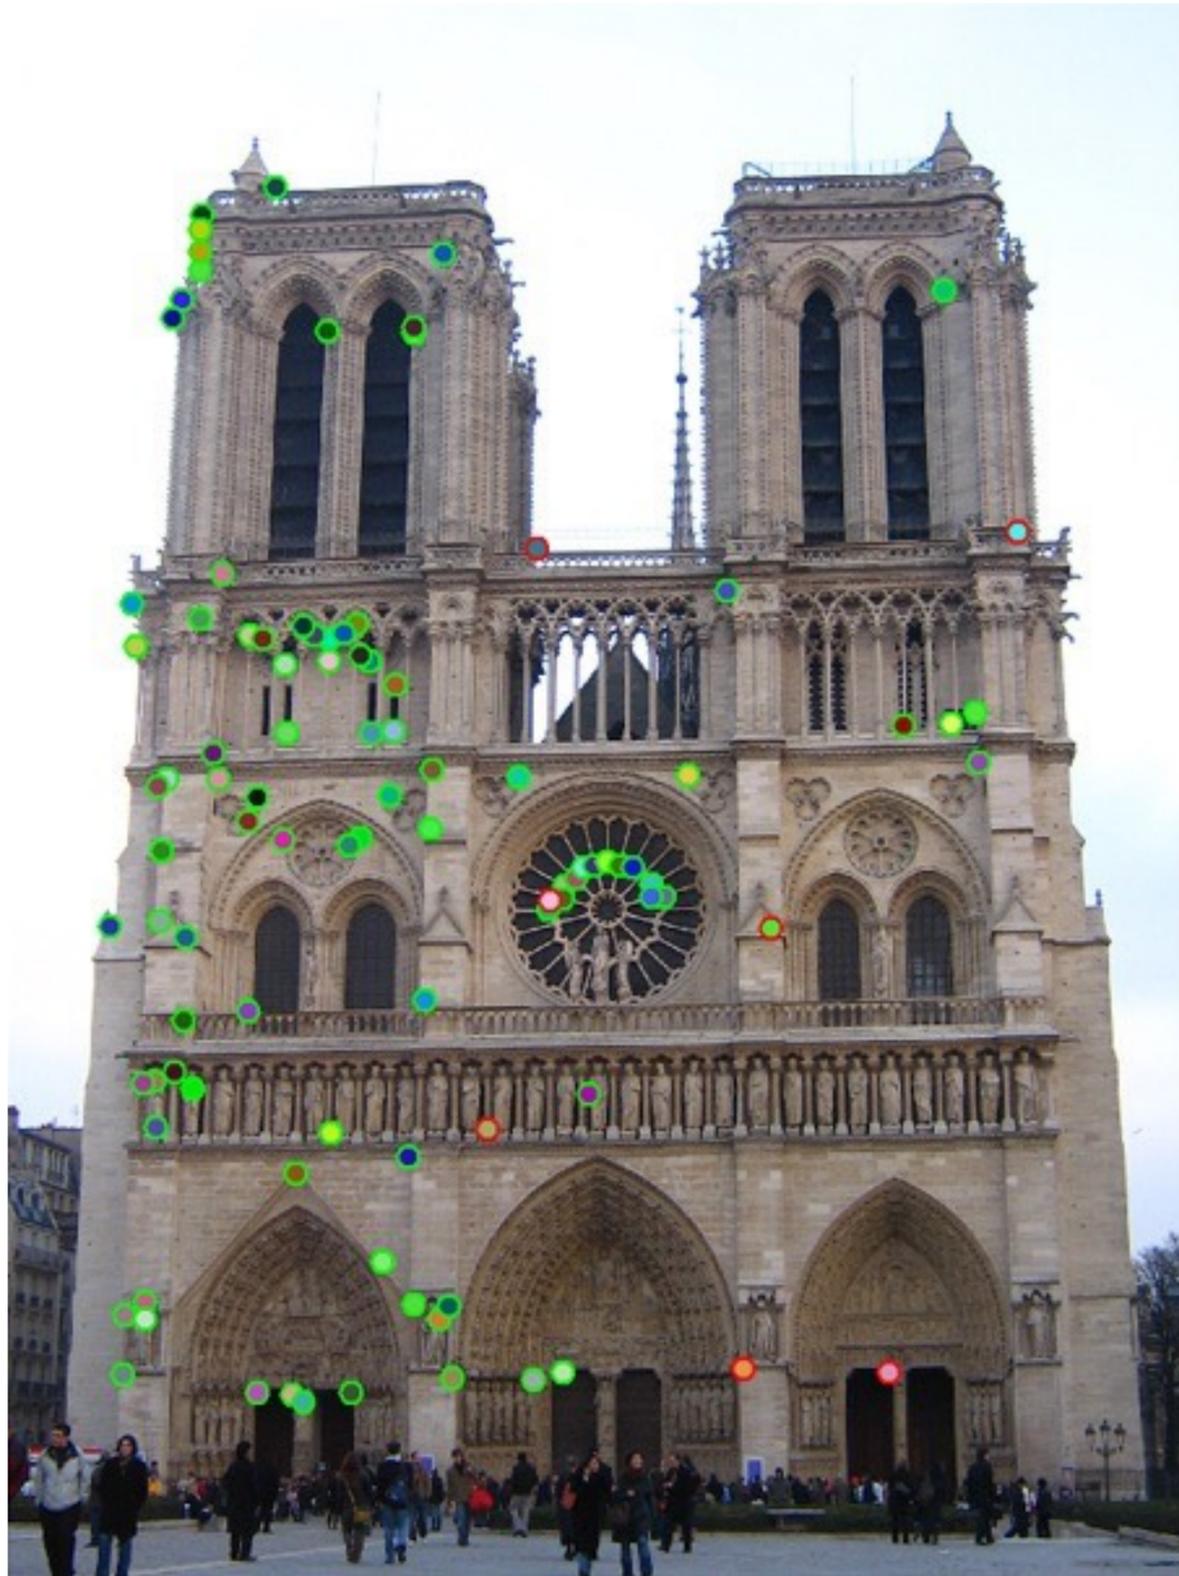
# SIFT invariances



# Today

- Finding correspondences
  - Computing local features
  - **Matching**
- Fitting a homography
- RANSAC

# How can we tell if two features match?



# Feature matching

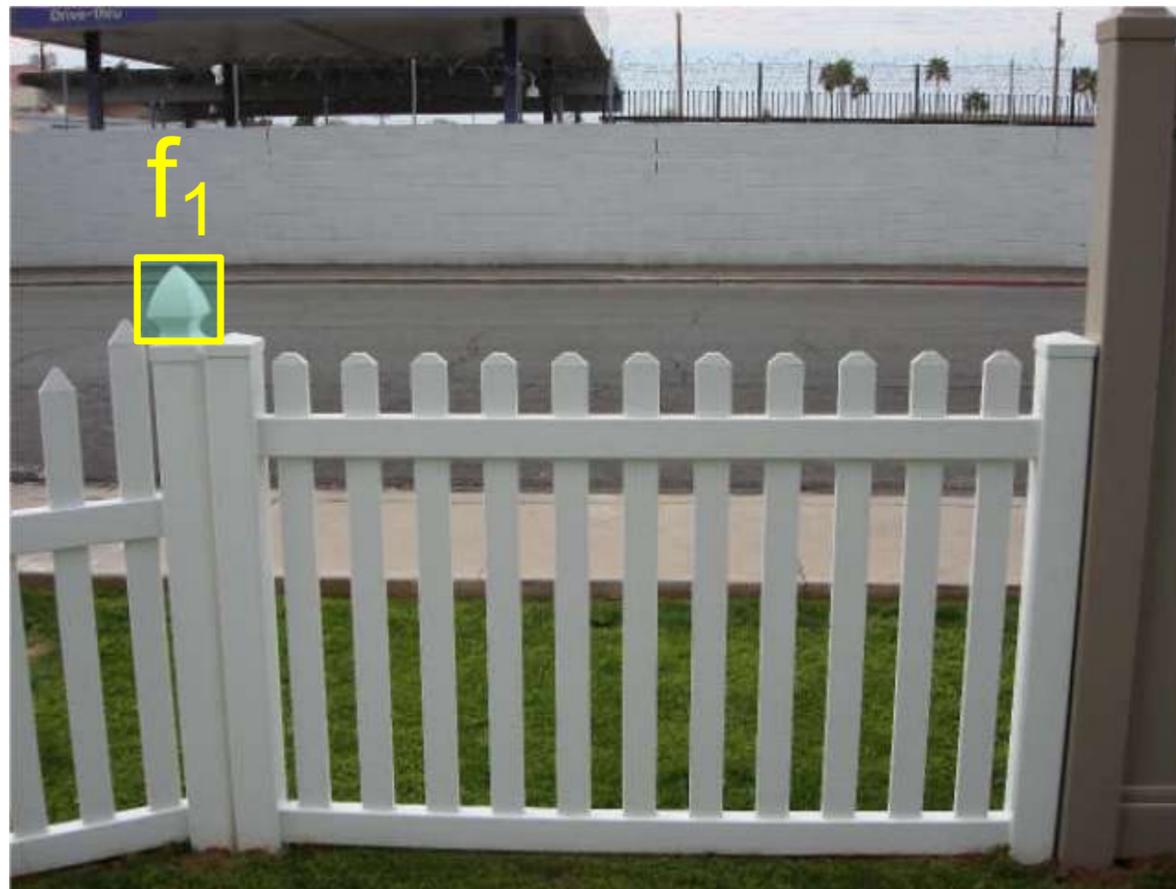
Given a feature in  $I_1$ , how do we find the best match in  $I_2$ ?

1. Define distance function that compares two descriptors
2. Test all the features in  $I_2$ , find the closest one.

# Finding matches

How do we know if two features match?

- Simple approach: are they the nearest neighbor in  $L_2$  distance,  $\|f_1 - f_2\|$ ?



$I_1$

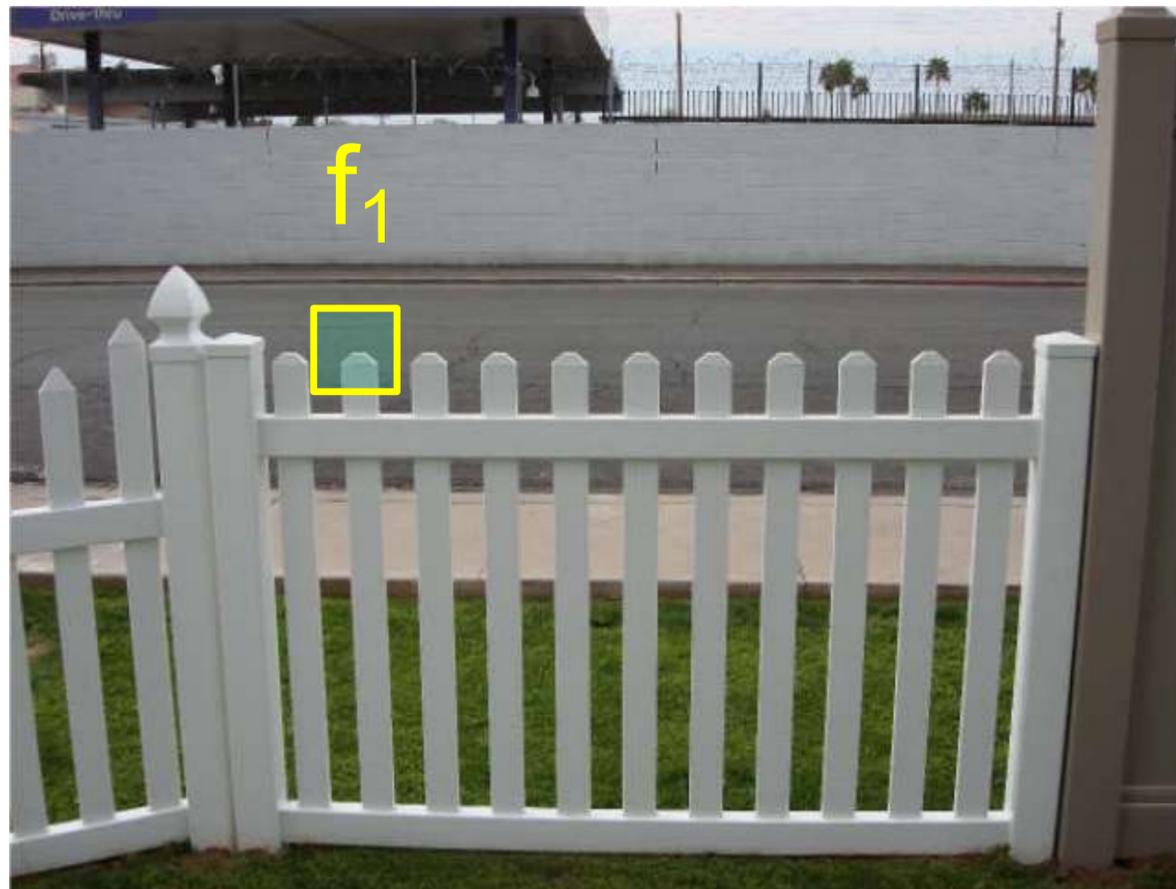


$I_2$

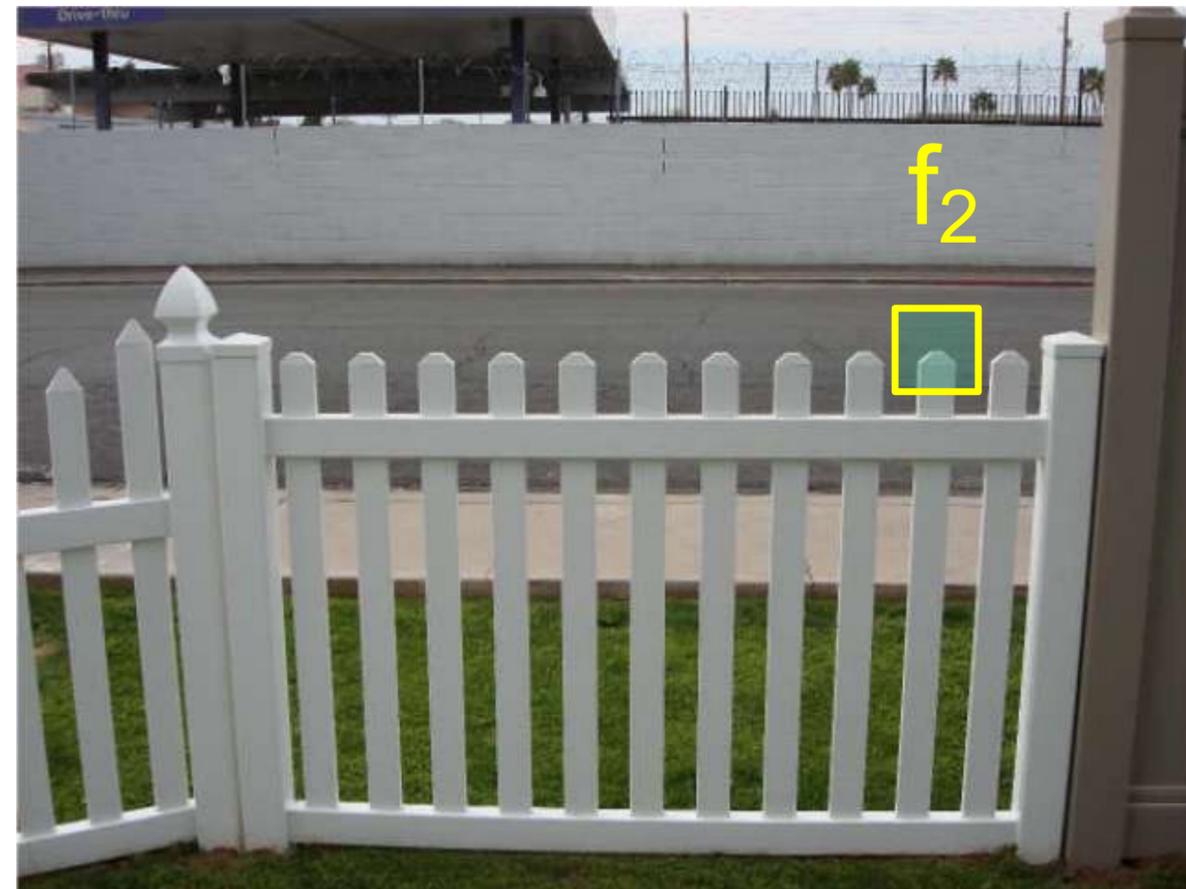
# Finding matches

How do we know if two features match?

- Simple approach: are they the nearest neighbor in  $L_2$  distance,  $\|f_1 - f_2\|$ ?
- Can give good scores to ambiguous (incorrect) matches



$I_1$

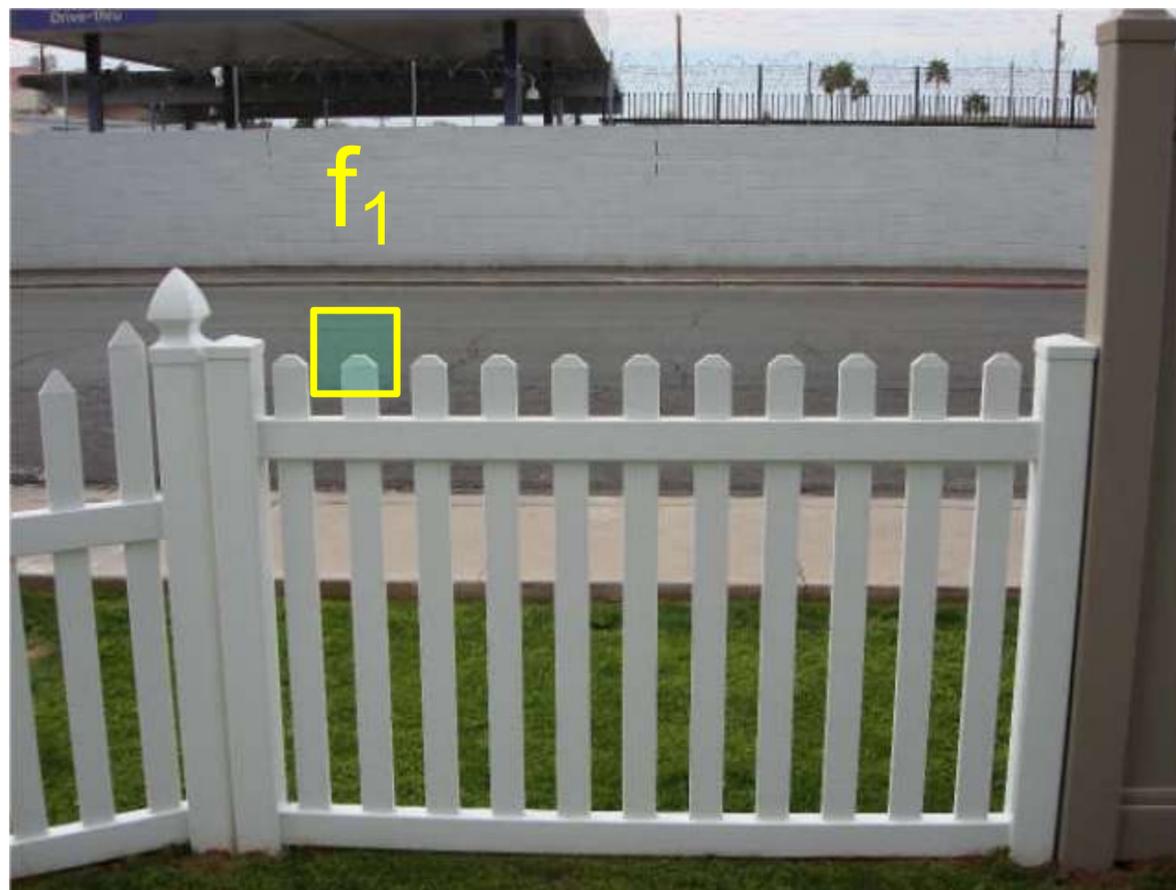


$I_2$

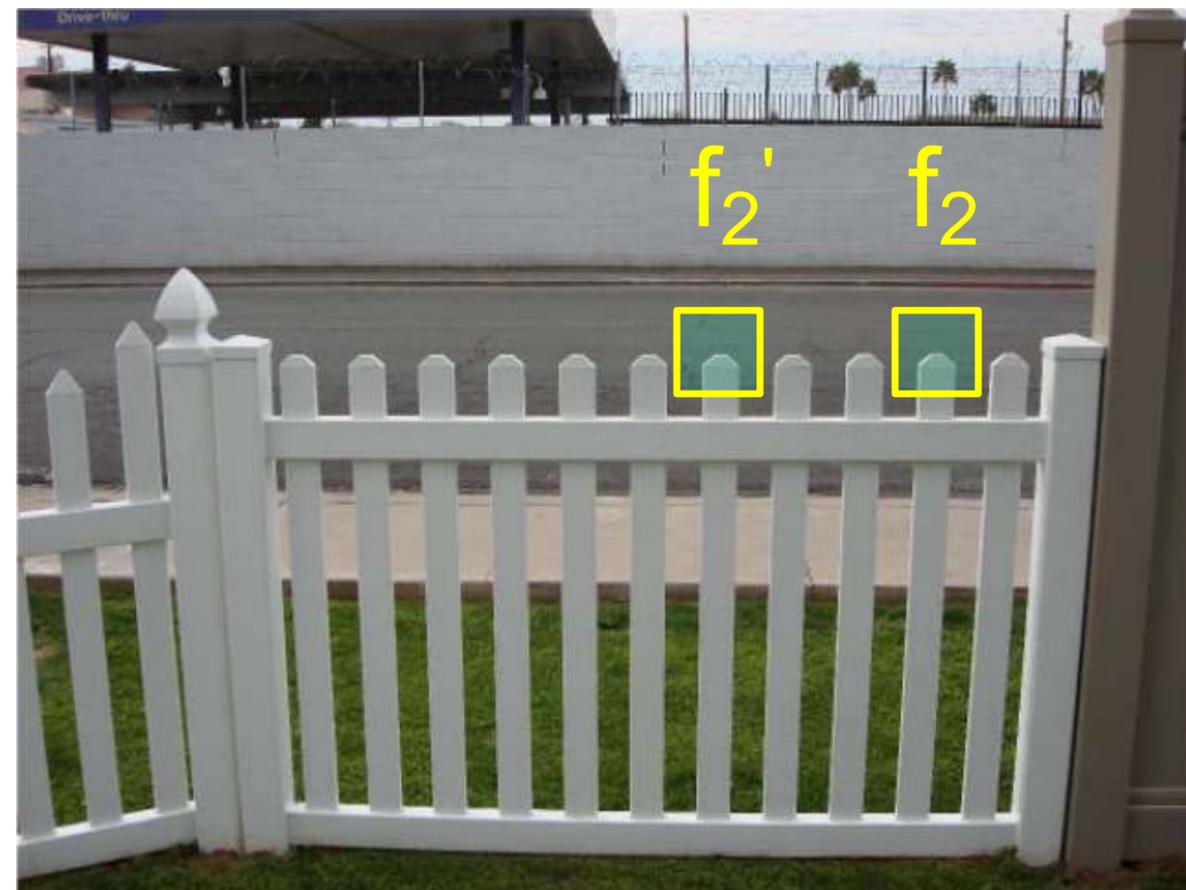
# Finding matches

Throw away matches that fail tests:

- **Ratio test:** this *by far* the best match? Compare best and 2nd-best matches.
  - Ratio distance =  $\|f_1 - f_2\| / \|f_1 - f_2'\|$
  - $f_2$  is best SSD match to  $f_1$  in  $I_2$
  - $f_2'$  is 2<sup>nd</sup> best SSD match to  $f_1$  in  $I_2$
- **Forward-backward consistency:**  $f_1$  should also be nearest neighbor of  $f_2$

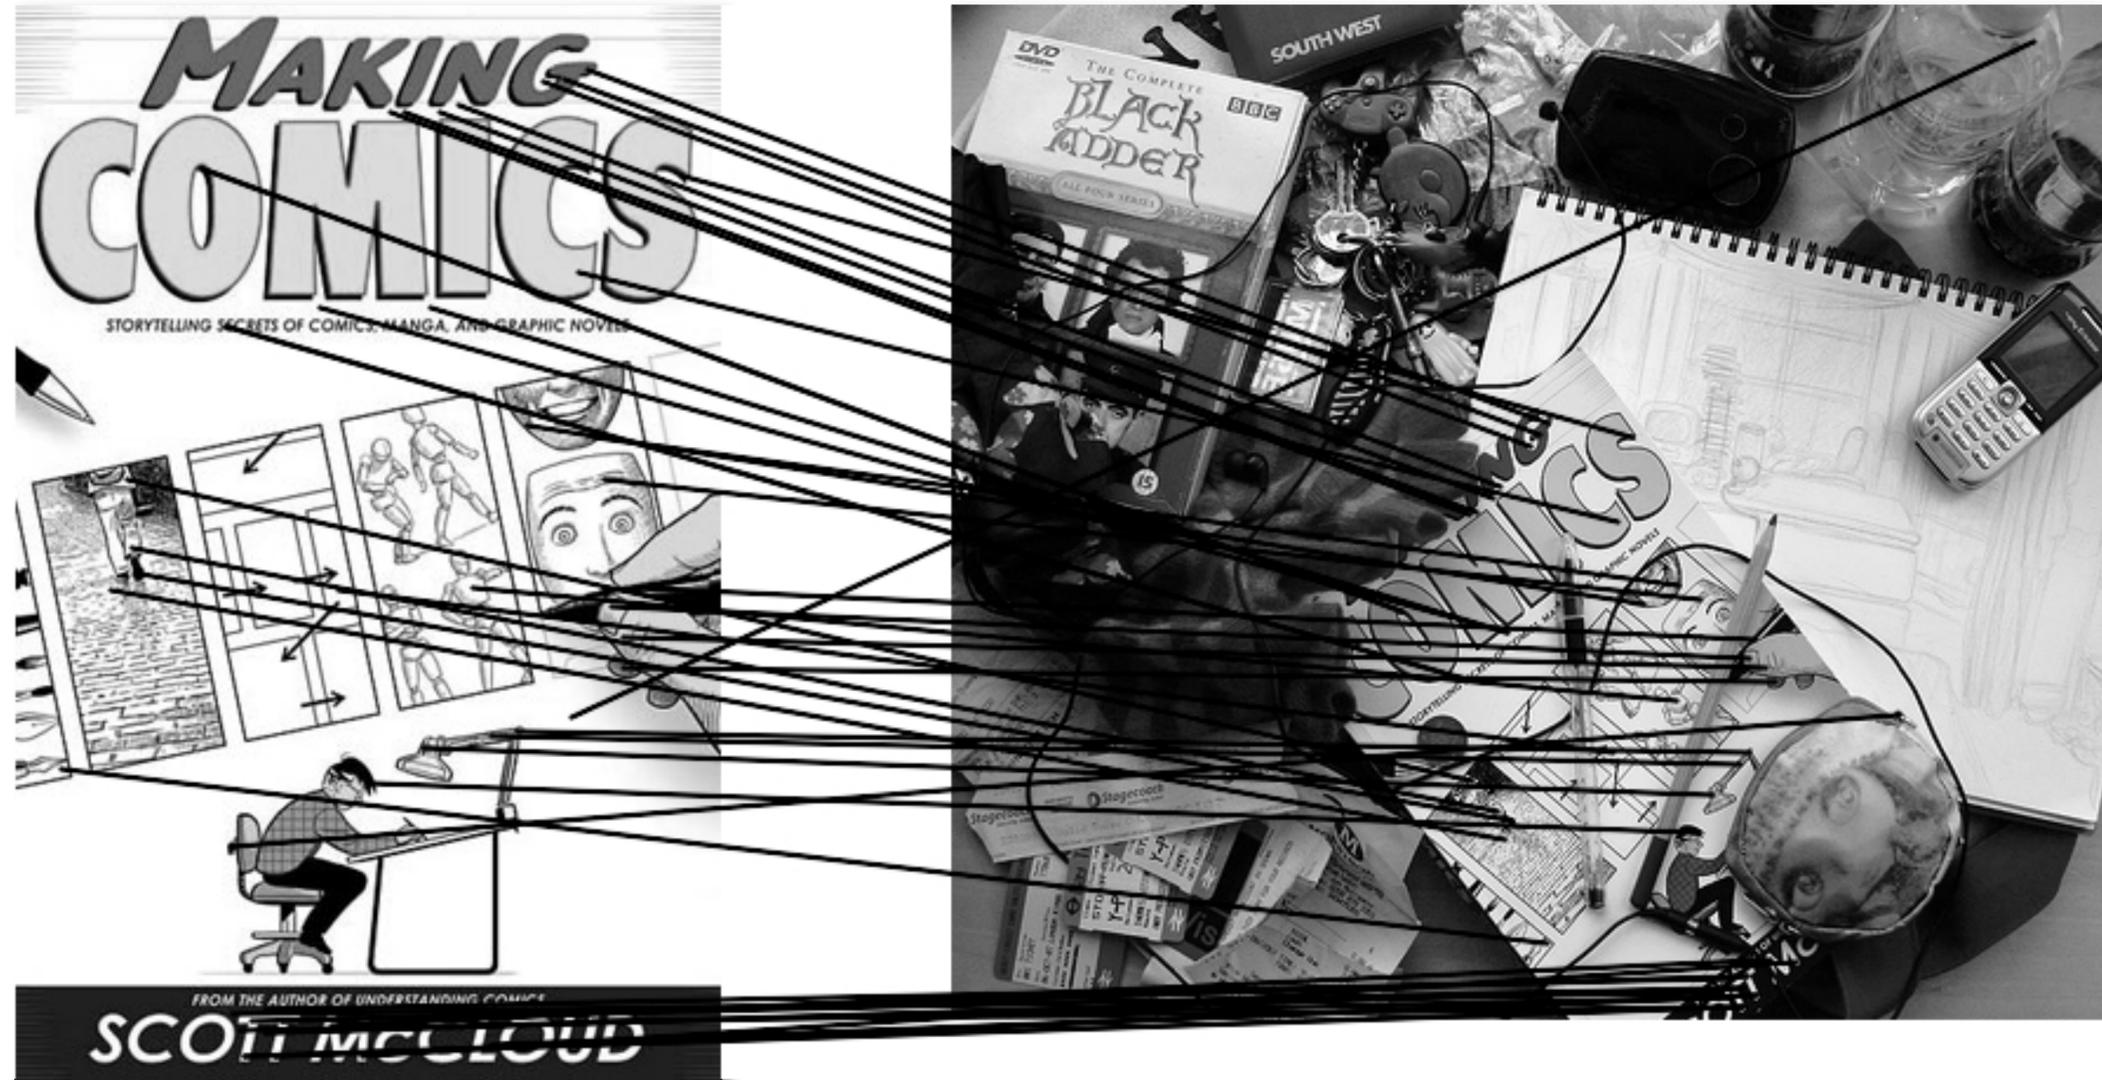


$I_1$



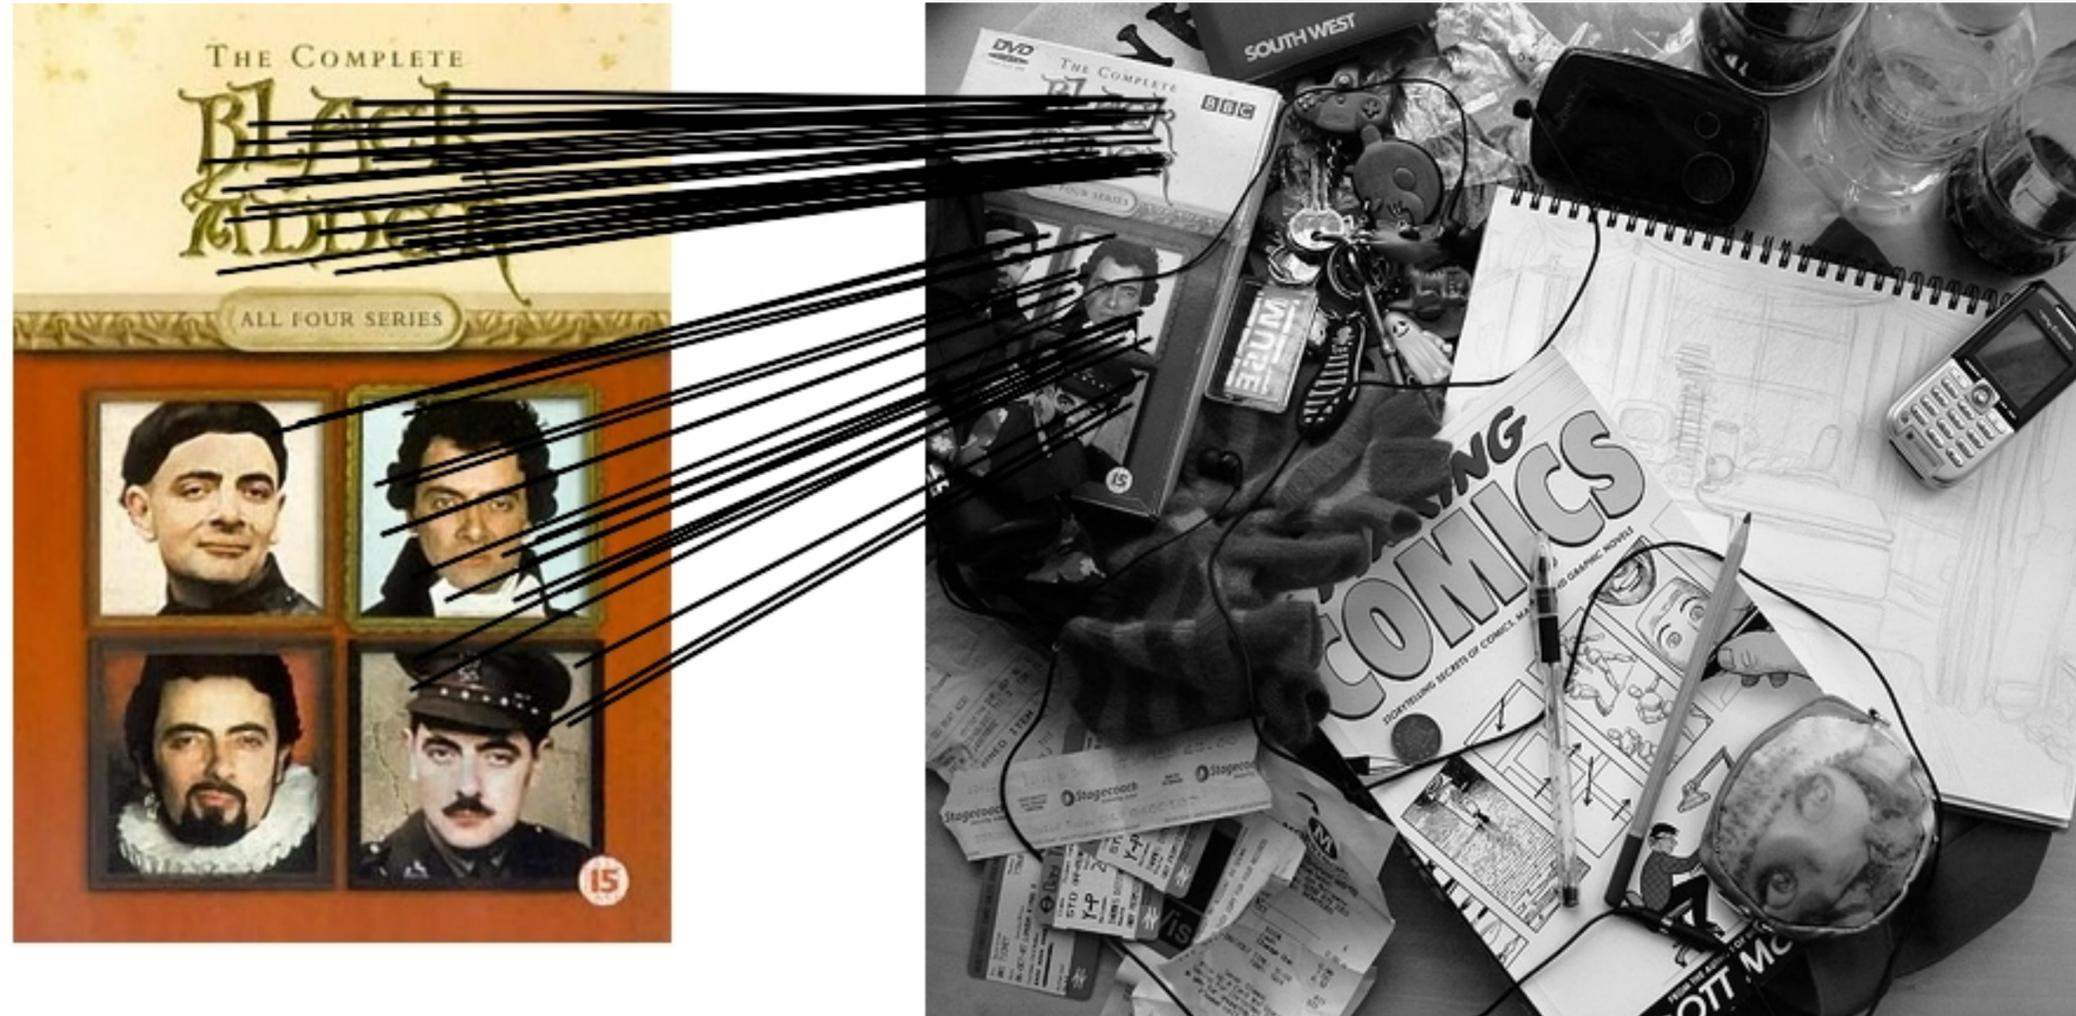
$I_2$

# Feature matching example



51 feature matches after ratio test

# Feature matching example

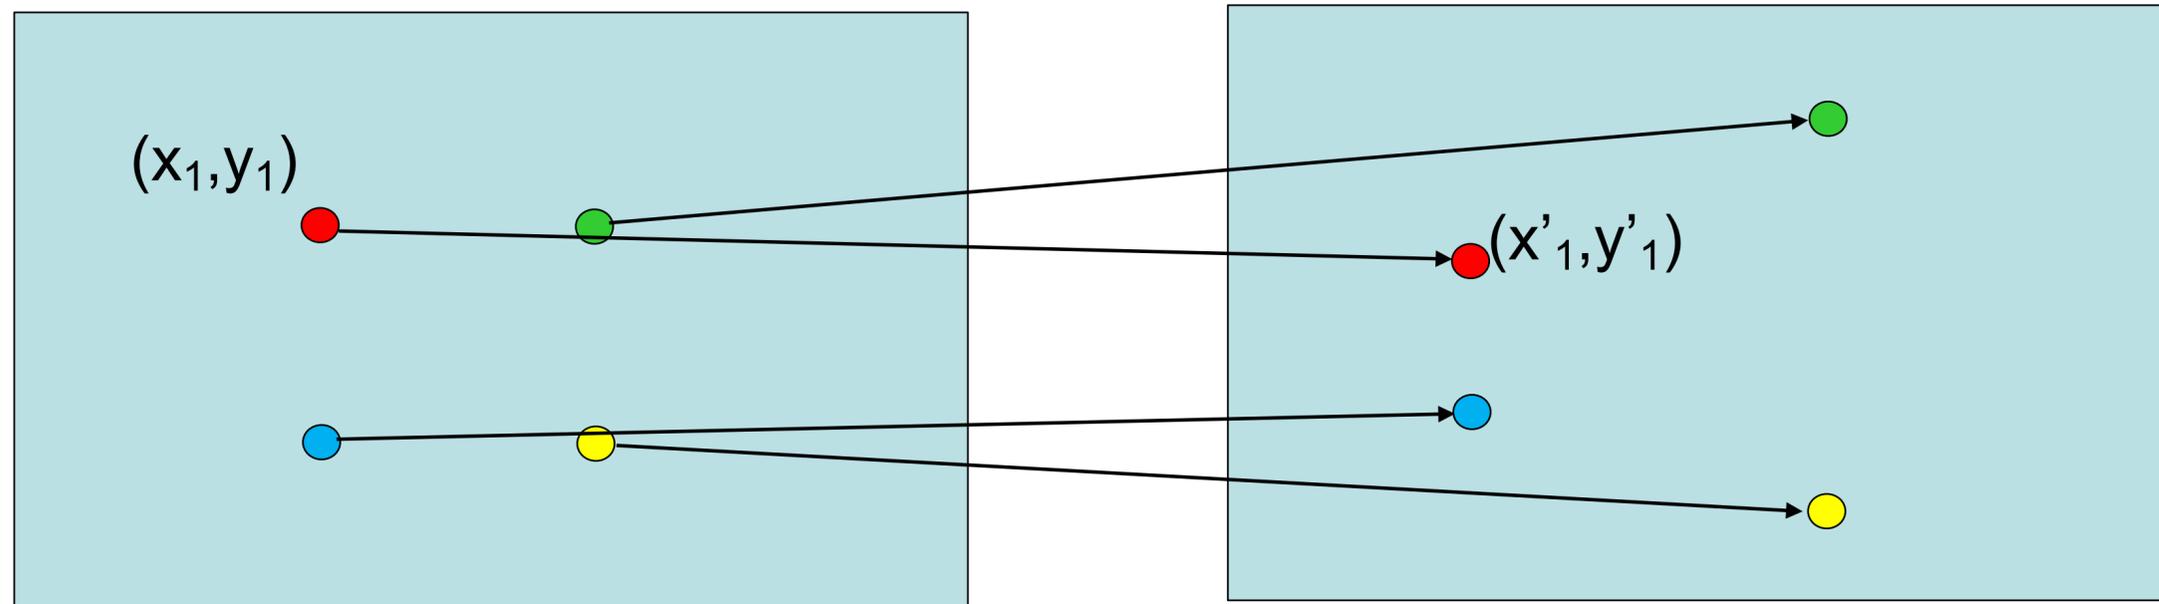


**58 feature matches after ratio test**

# Today

- Finding correspondences
  - Computing local features
  - Matching
- **Fitting a homography**
- RANSAC

# From matches to a homography



$$\begin{bmatrix} x'_1 \\ y'_1 \\ w_1 \end{bmatrix} = \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & i \end{bmatrix} \cdot \begin{bmatrix} x_1 \\ y_1 \\ 1 \end{bmatrix}$$

# From matches to a homography

Point in 1st image

Matched point in 2nd

$$\text{minimize } J(H) = \sum_i \|f_H(p_i) - p'_i\|^2$$

where  $f_H(p_i) = Hp_i / (H_3^T p_i)$  applies homography

Remember: homogenous coordinates.  
 $H_3$  is the third row of  $H$

# Option #1: Direct linear transform

$$\begin{bmatrix} x_1' \\ y_1' \\ w_1 \end{bmatrix} = \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & i \end{bmatrix} \cdot \begin{bmatrix} x_1 \\ y_1 \\ 1 \end{bmatrix}$$

Leaving homogeneous coordinates:

$$x_1' = \frac{ax_1 + by_1 + c}{gx_1 + hy_1 + i}$$

$$y_1' = \frac{dx_1 + ey_1 + f}{gx_1 + hy_1 + i}$$

Re-arranging the terms:

$$gx_1x_1' + hy_1x_1' + ix_1' = ax_1 + by_1 + c$$

$$gx_1y_1' + hy_1y_1' + ix_1' = dx_1 + ey_1 + f$$

# Option #1: Direct linear transform

$$gx_1x'_1 + hy_1x'_1 + ix'_1 = ax_1 + by_1 + c$$

$$gx_1y'_1 + hy_1y'_1 + iy'_1 = dx_1 + ey_1 + f$$

More rearranging:

$$gx_1x'_1 + hy_1x'_1 + ix'_1 - ax_1 - by_1 - c = 0$$

$$gx_1y'_1 + hy_1y'_1 + iy'_1 - dx_1 - ey_1 - f = 0$$

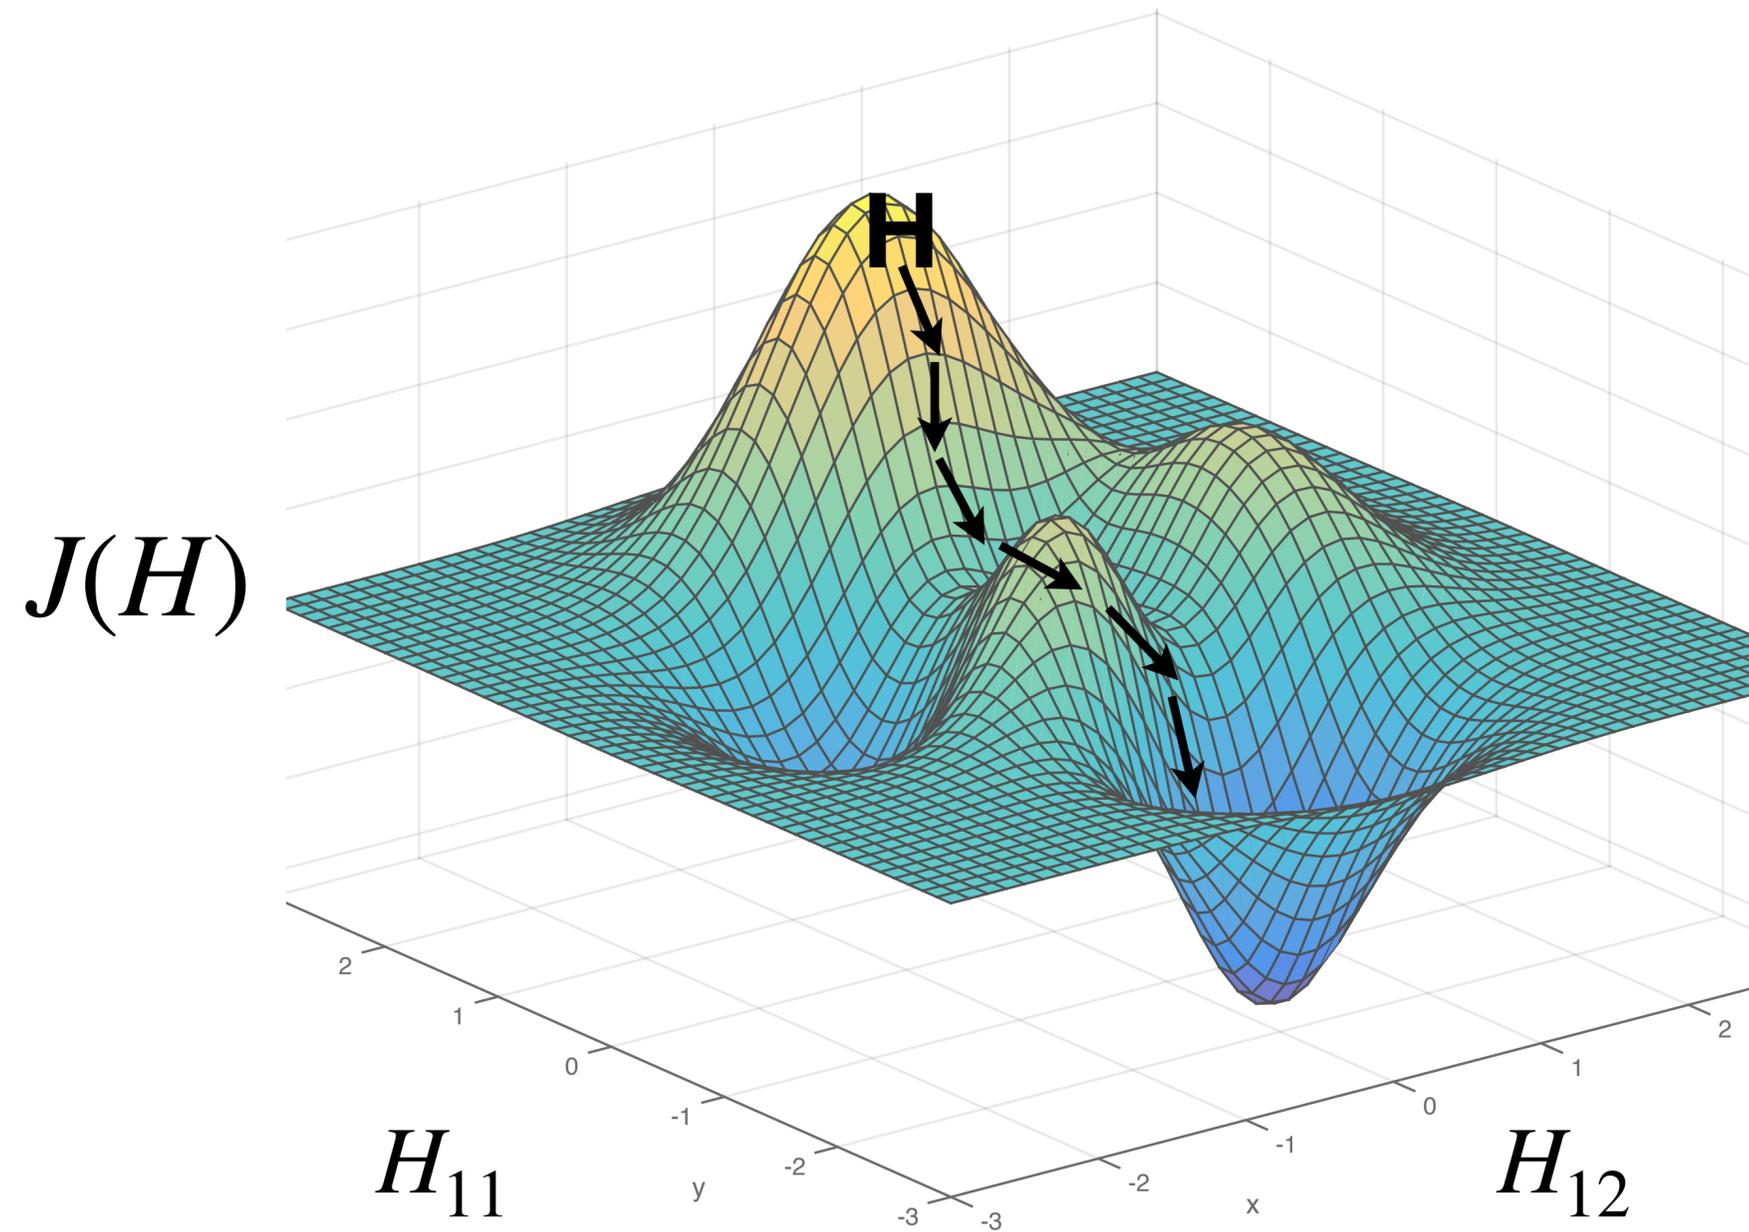
In matrix form:

$$\begin{bmatrix} -x_1 & -y_1 & -1 & 0 & 0 & 0 & x_1x'_1 & y_1x'_1 & x'_1 \\ 0 & 0 & 0 & -x_1 & -y_1 & -1 & x_1y'_1 & y_1y'_1 & y'_1 \end{bmatrix} \begin{bmatrix} a \\ b \\ c \\ d \\ e \\ f \\ g \\ h \\ i \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$$

Can solve using Singular Value Decomposition (SVD).

Fast to solve (but not using “right” loss function). Uses an algebraic trick.  
Often used in practice for initial solutions!

# Option #2: Optimization



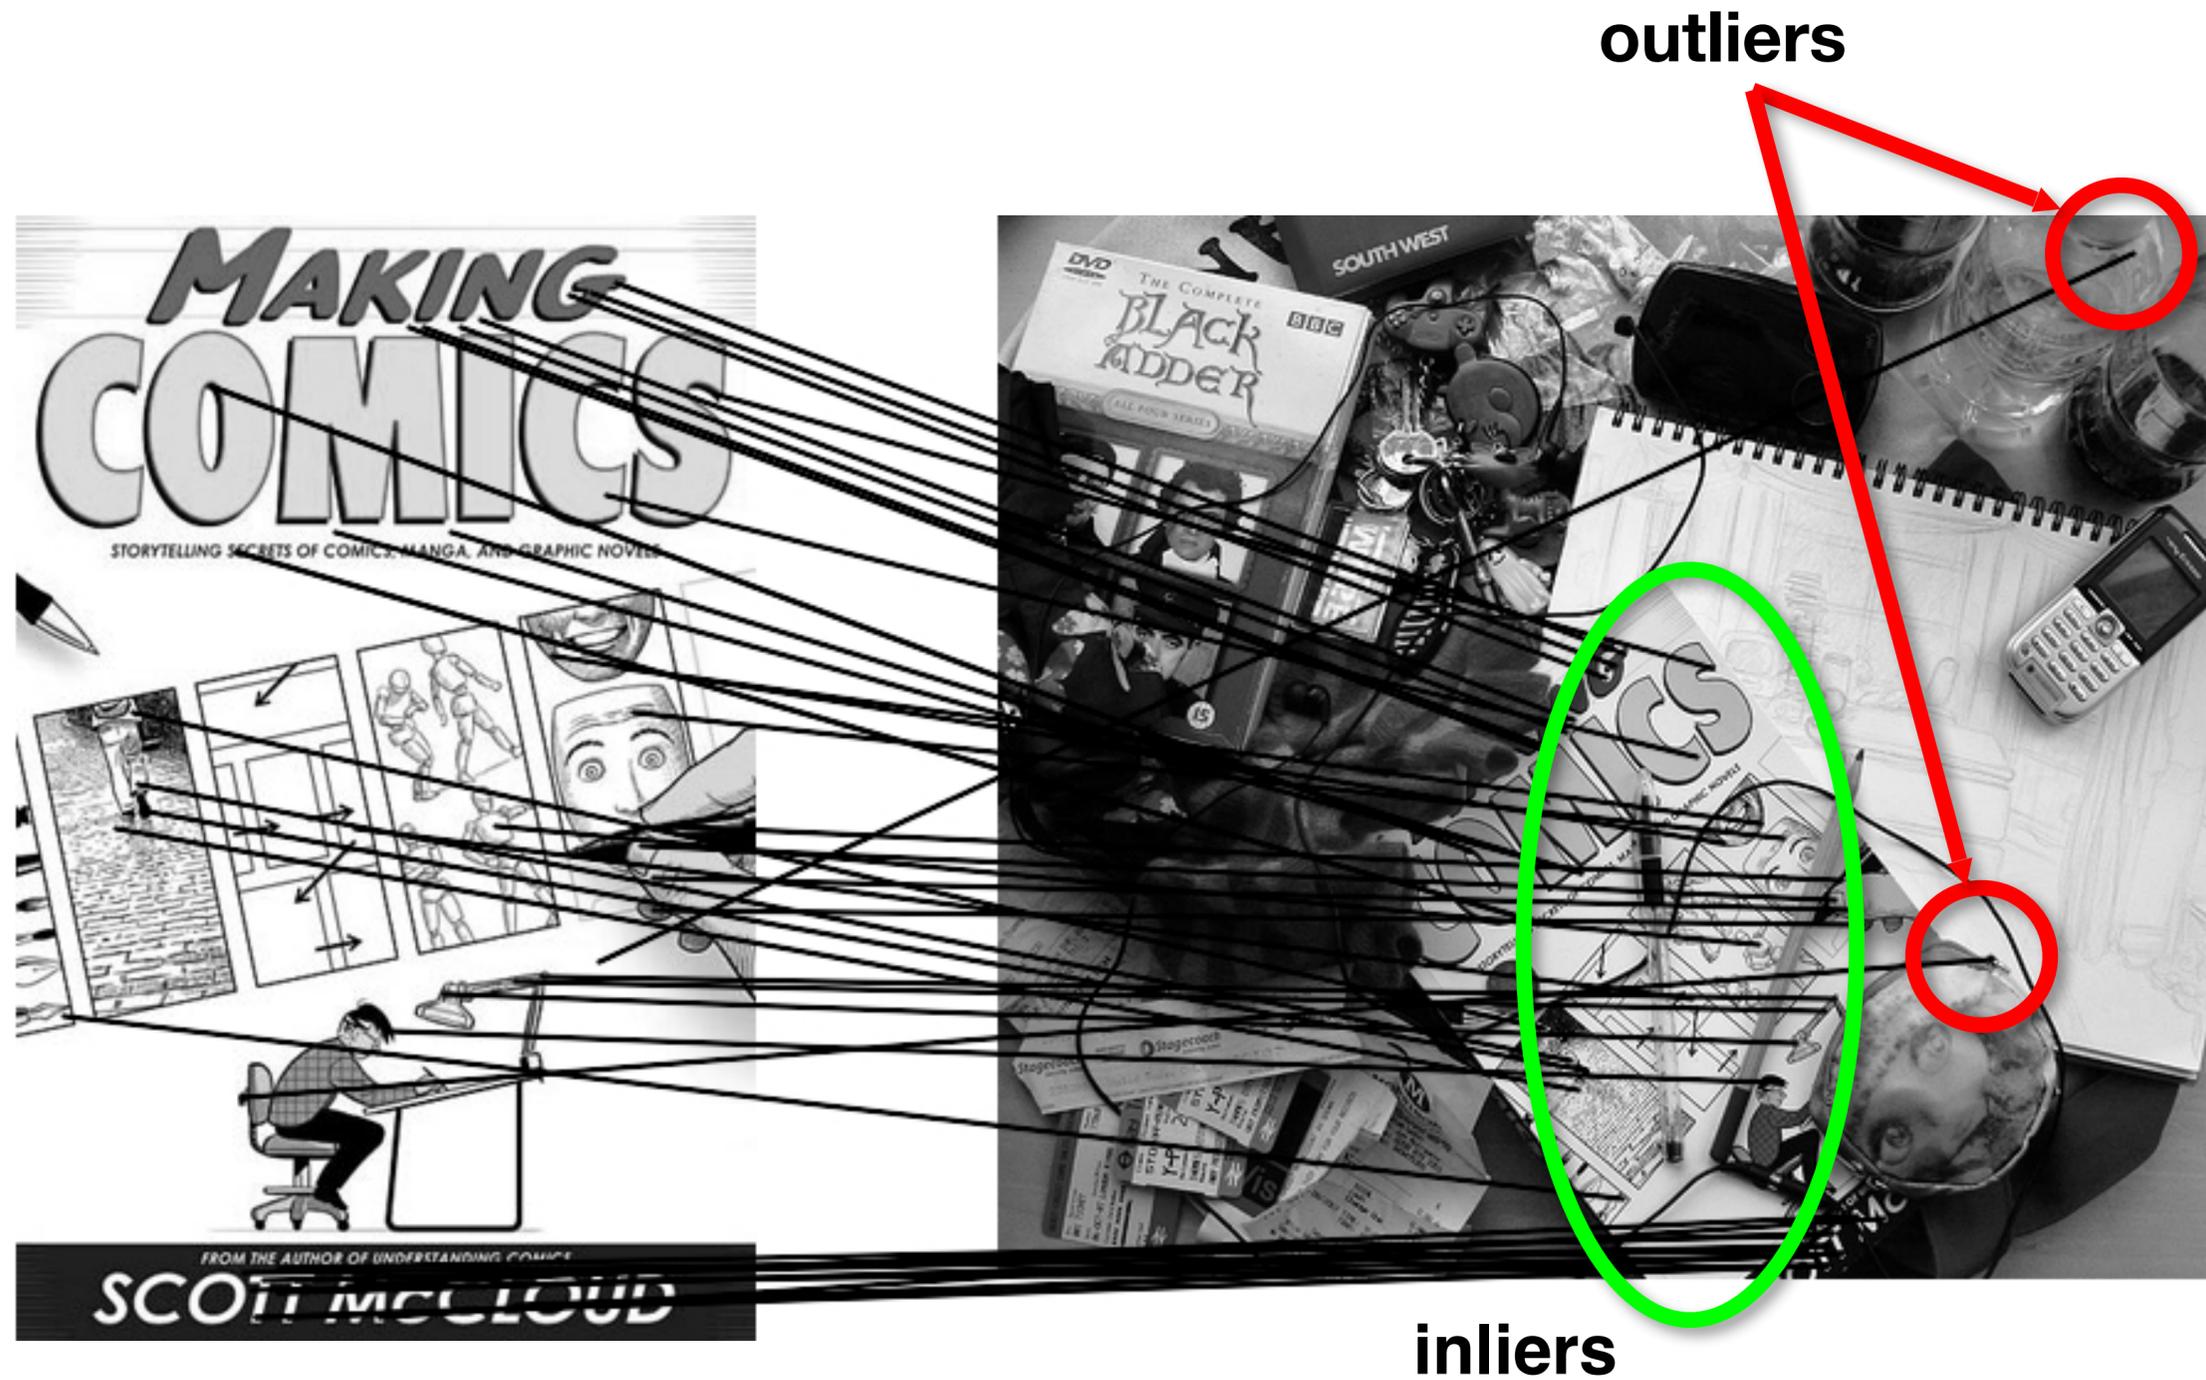
$$\text{minimize } J(H) = \sum_i ||f_H(p_i) - p'_i||^2$$

# Optimization

$$\text{minimize } J(H) = \sum_i \|f_H(p_i) - p'_i\|^2$$

- Can use gradient descent, just like when learning neural nets
- These problems are **smaller scale** than deep learning problems but have **more local optima**:
  - Use 2nd derivatives to improve optimization
  - Can use finite differences or autodiff
- Can use special-purpose **nonlinear least squares** methods.
  - Exploits structure in the problem for a sum-of-squares loss.

# Problem: outliers



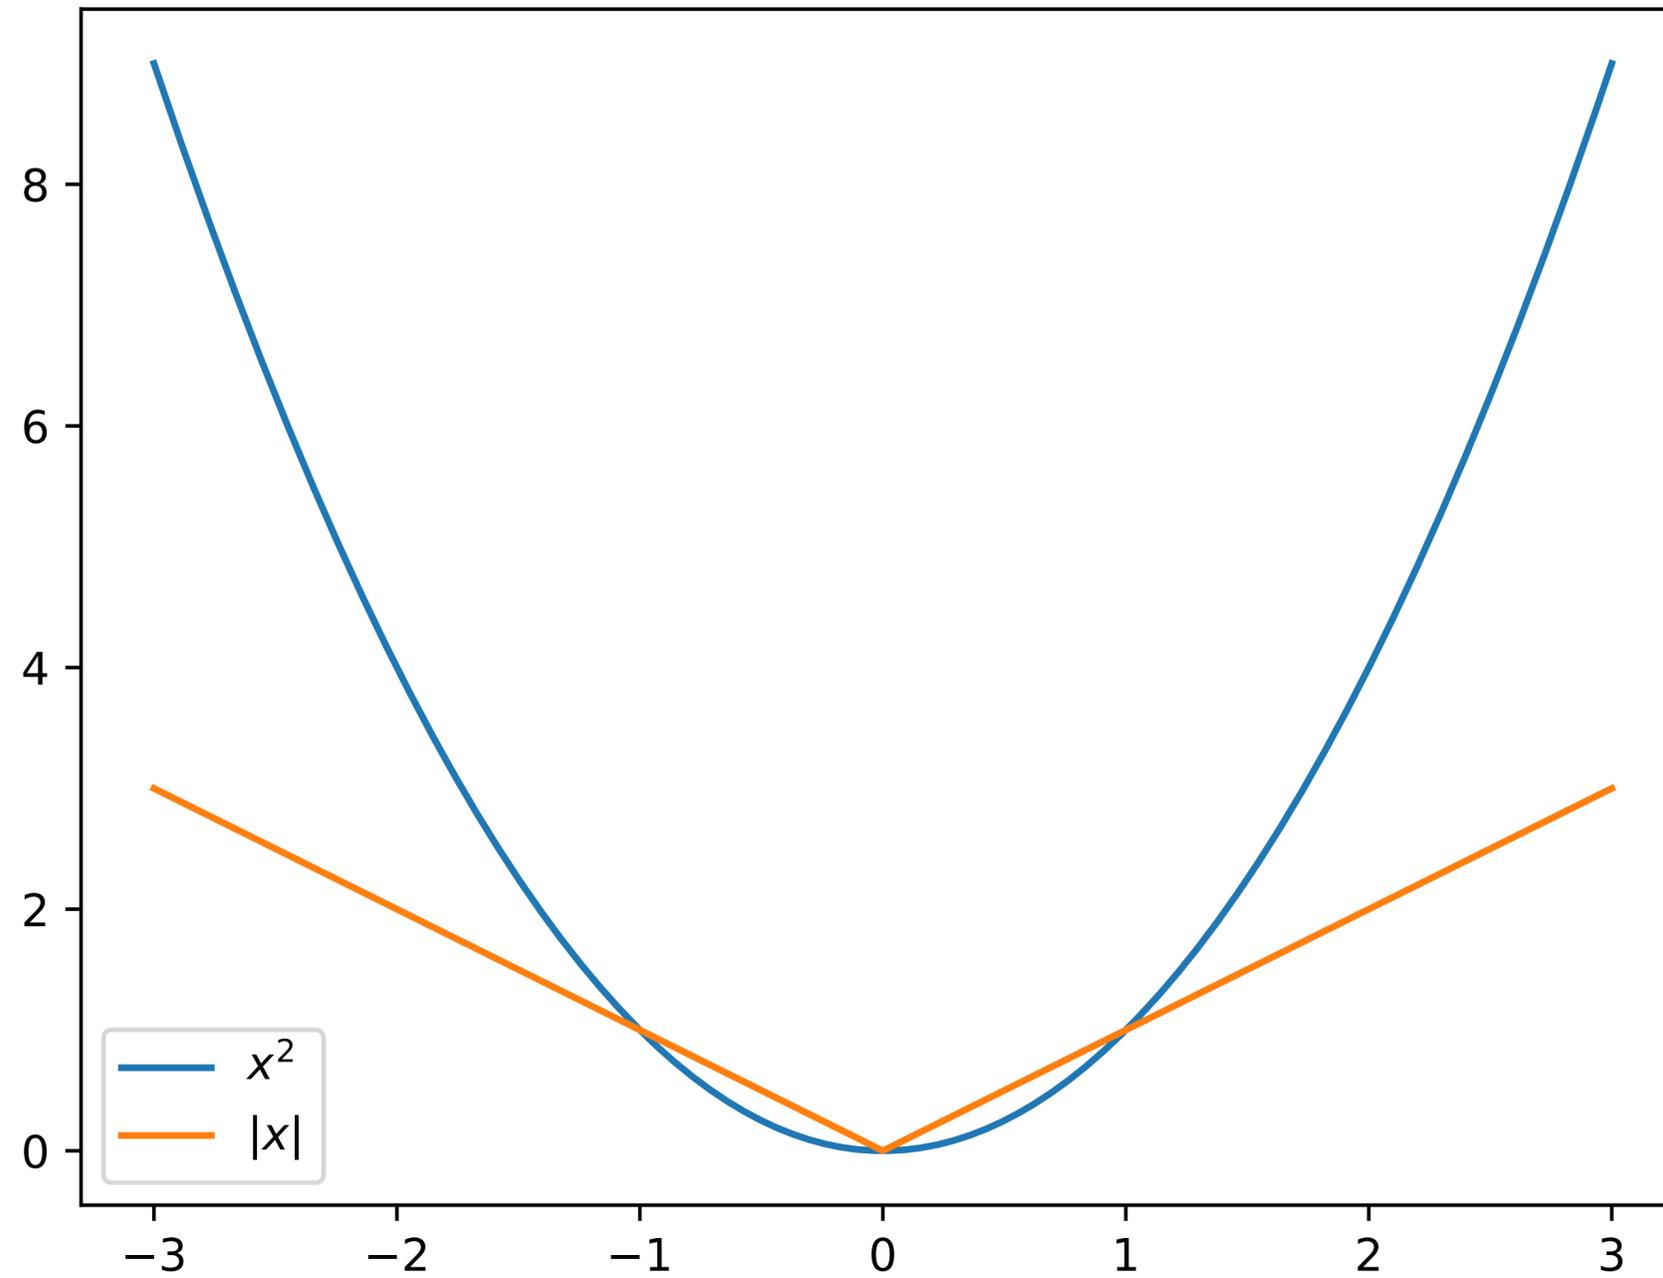
# One idea: robust loss functions

minimize  $J(H) = \sum_{i=1}^N \sum_{j=1}^2 \rho(f_H(p_{ij}) - p'_{ij})$

where  $\rho(x)$  is a **robust** loss.

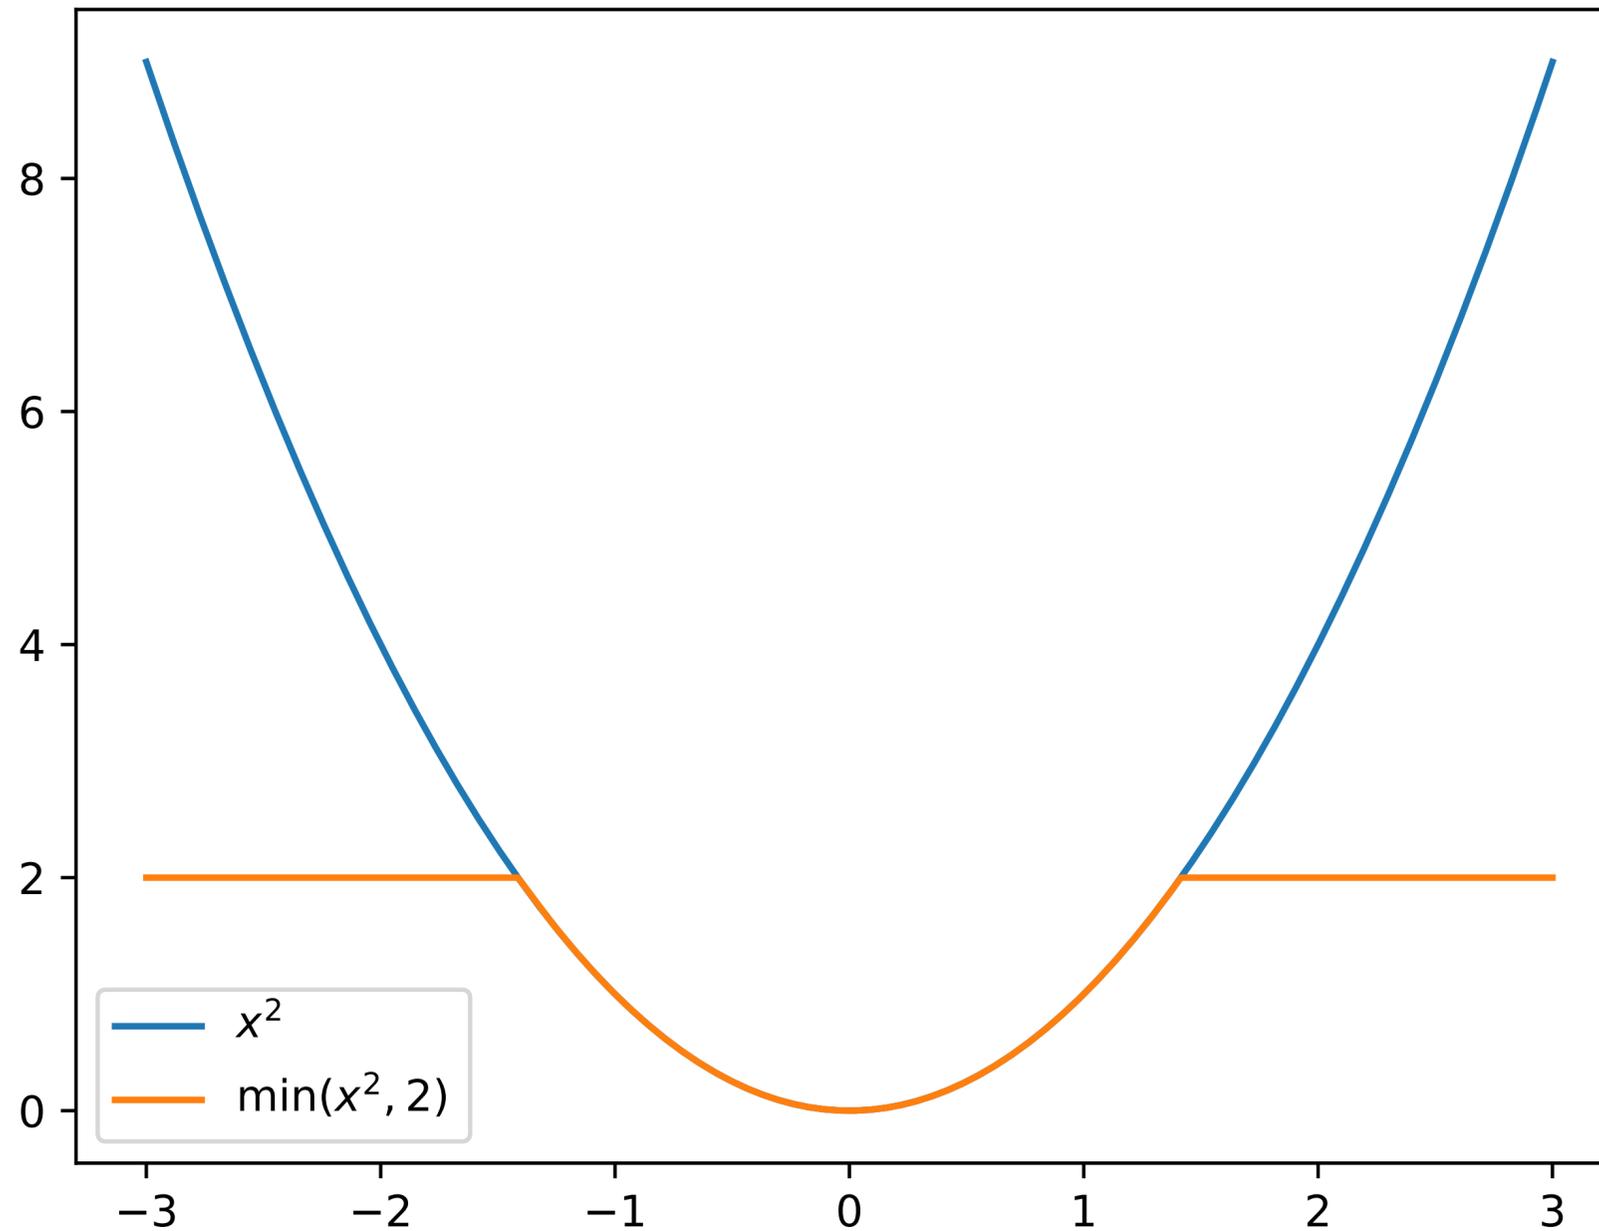
Special case:  $\rho(x) = x^2$  is L2 loss (same as before)

# Robust loss functions



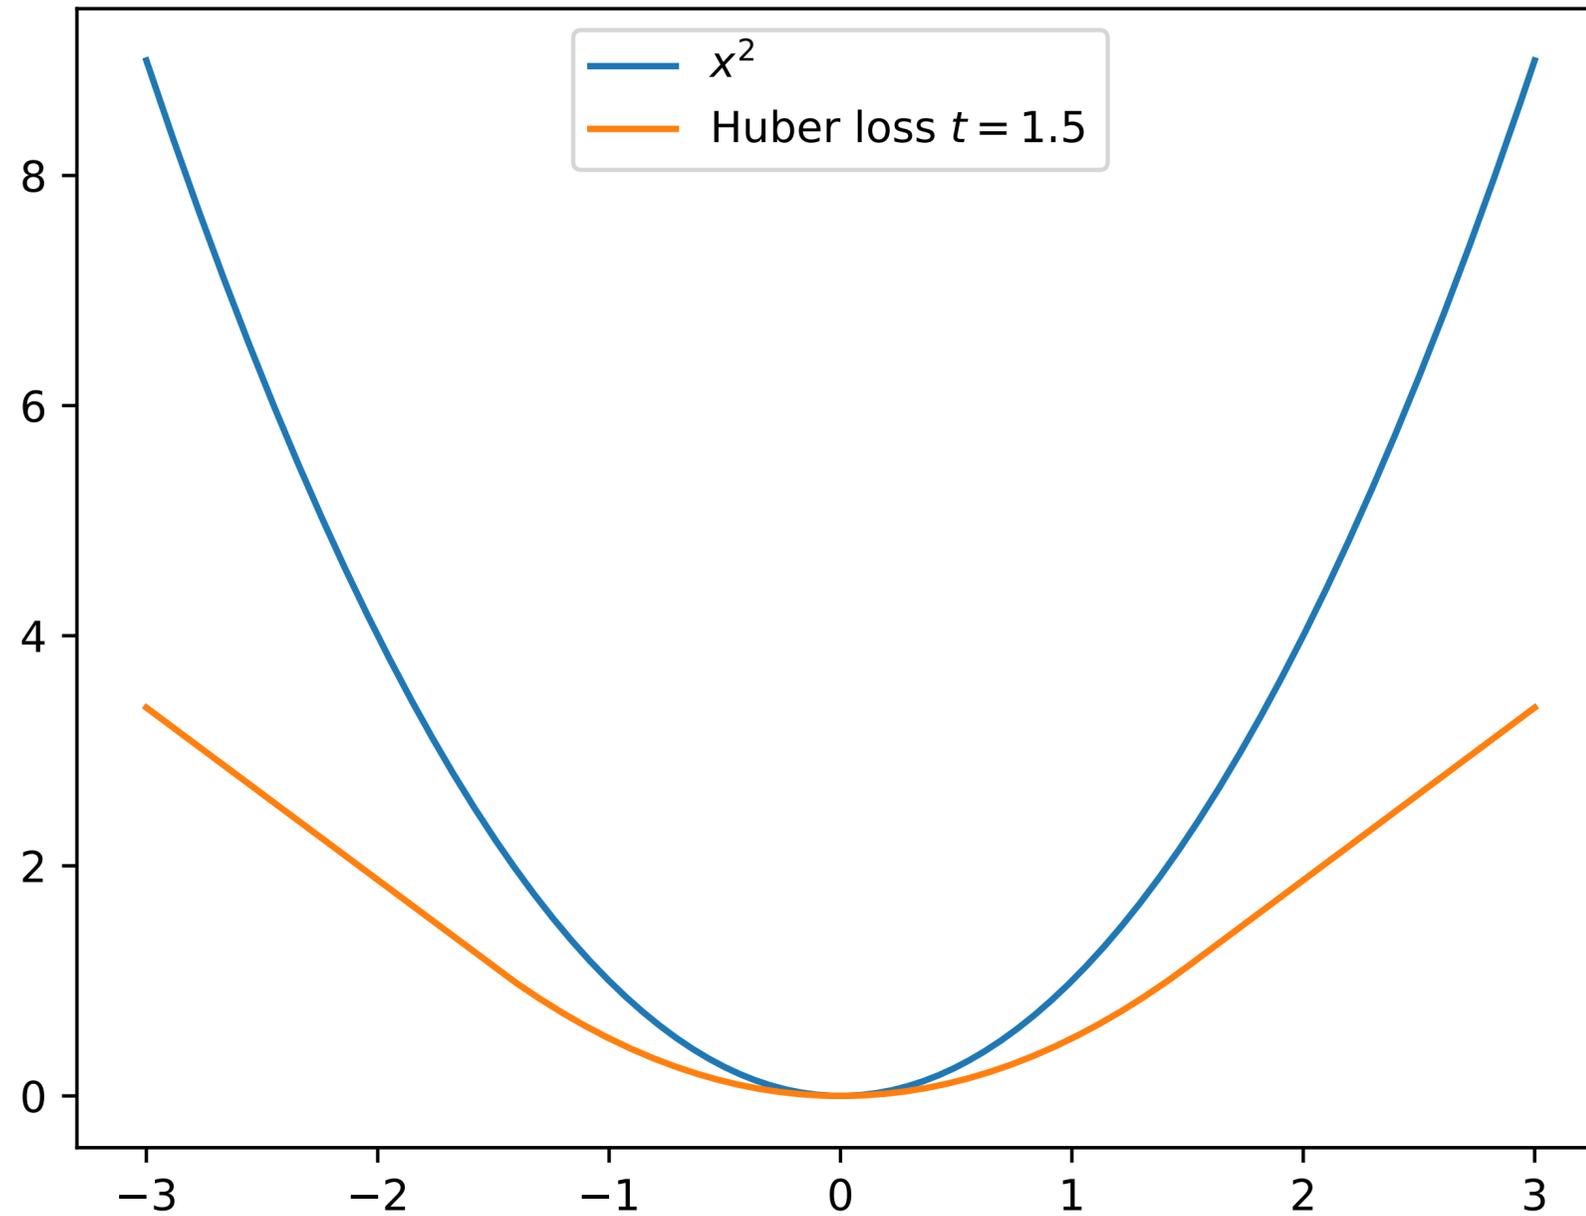
L1 loss:  $\rho(x) = |x|$

# Robust loss functions



Truncated quadratic:  $\rho(x) = \min(x^2, \tau)$

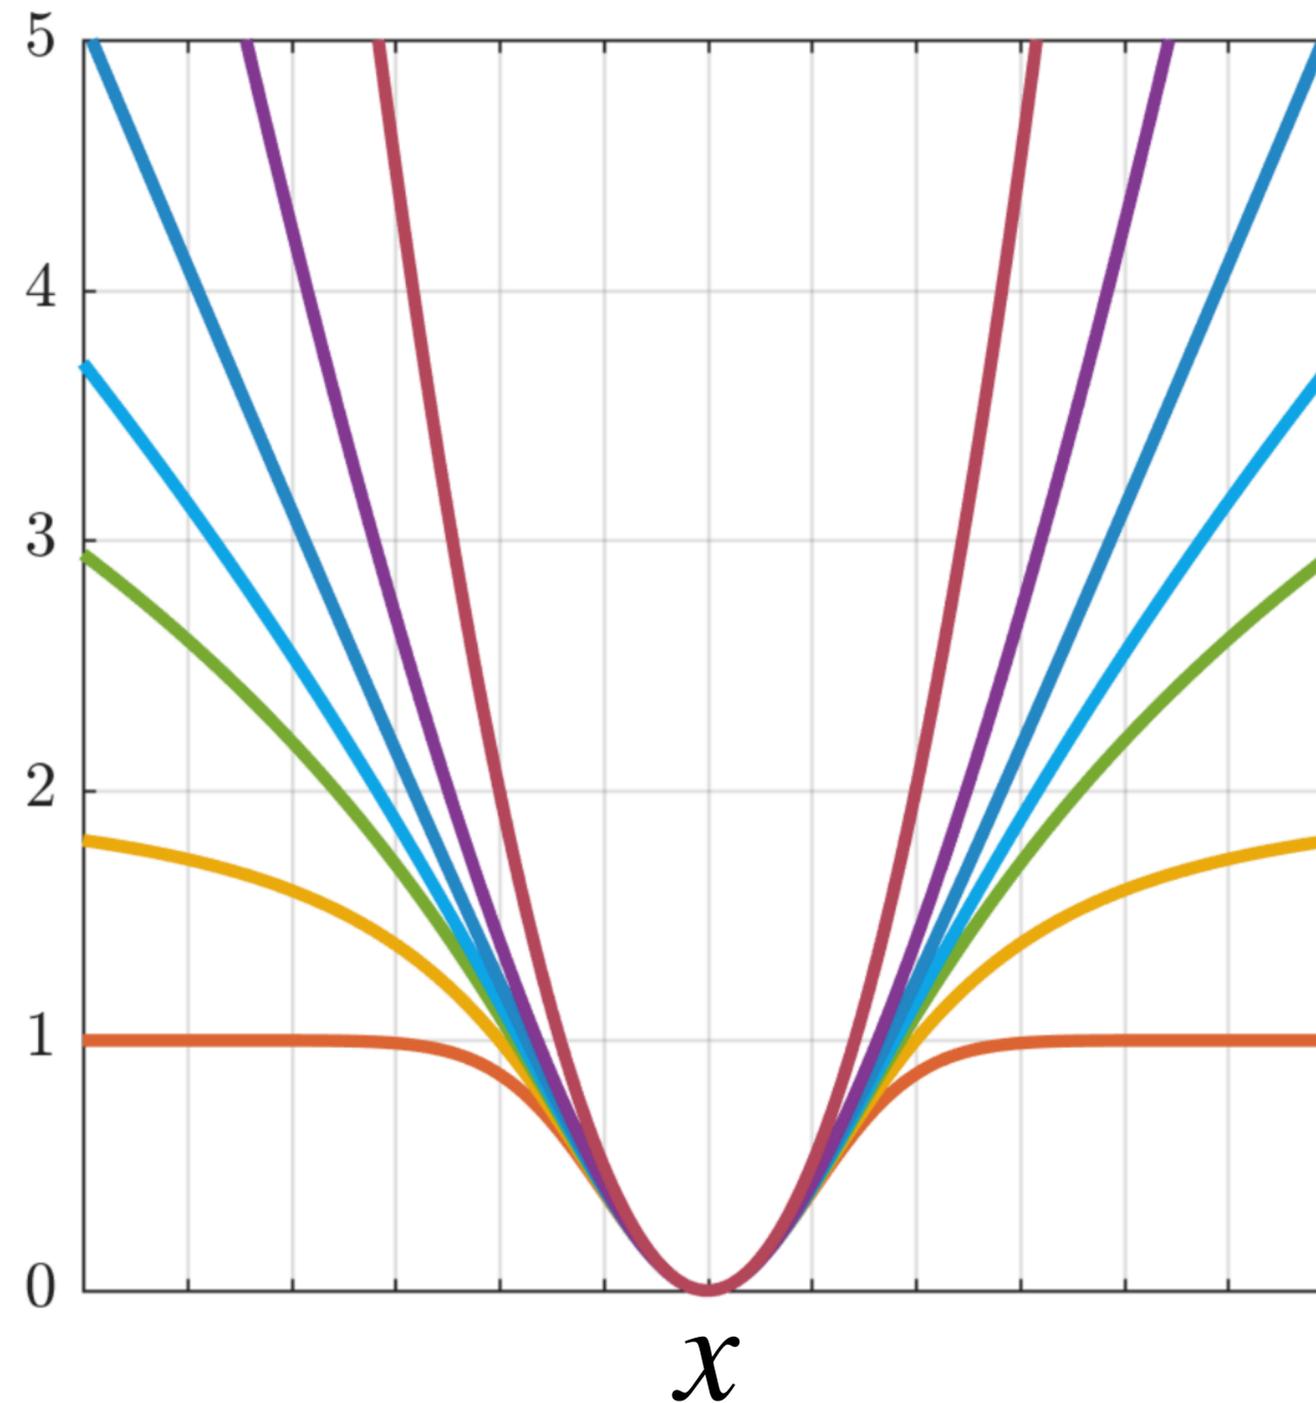
# Robust loss functions



Huber loss:

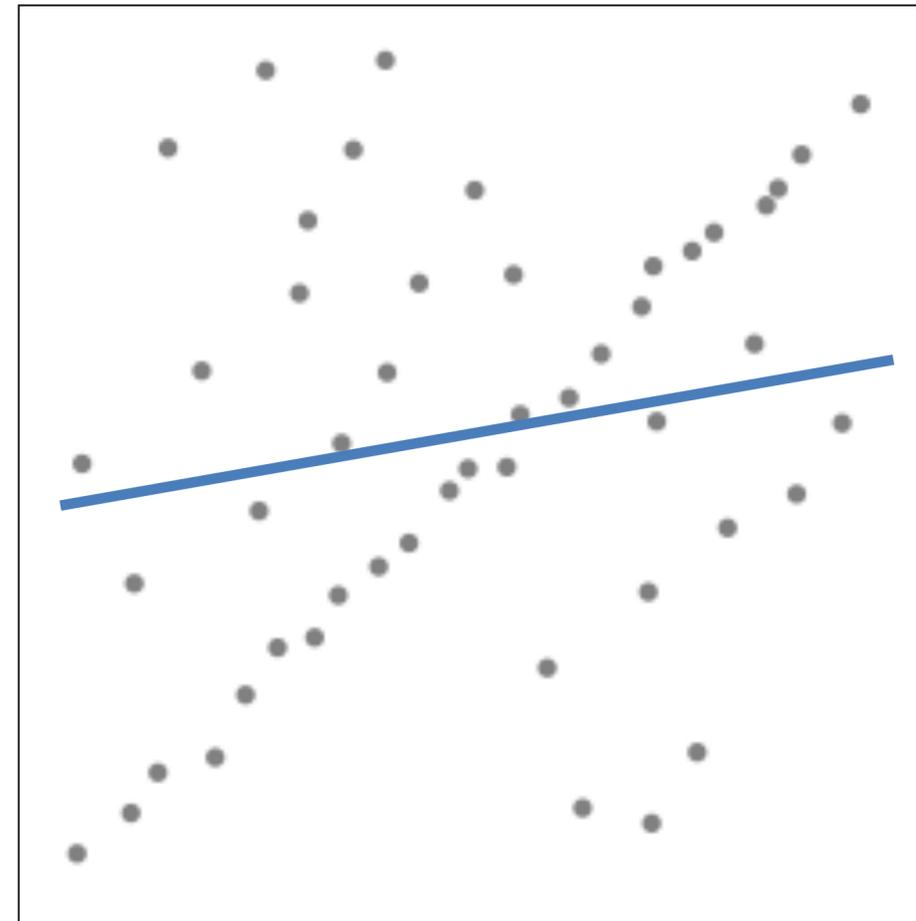
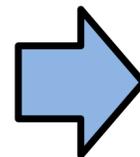
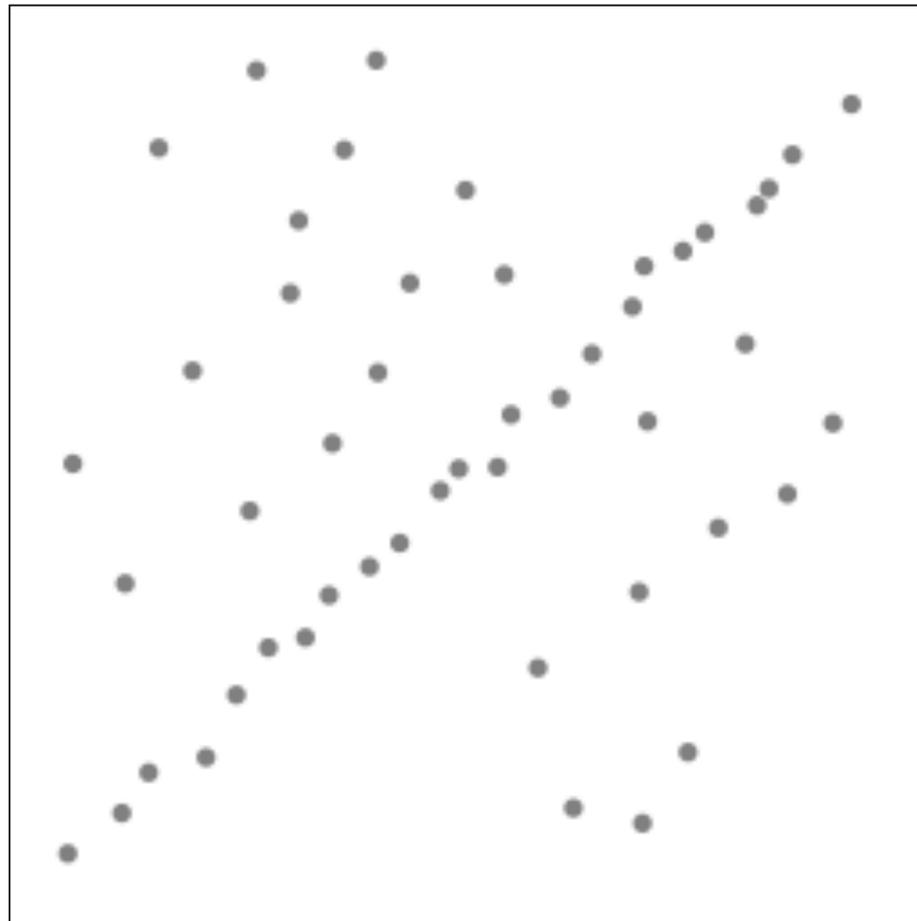
$$\rho(x) = \begin{cases} \frac{1}{2}x^2 & \text{if } |x| \leq \tau, \\ \tau(|x| - \frac{1}{2}\tau), & \text{else} \end{cases}$$

# Robust loss functions



# Handling outliers

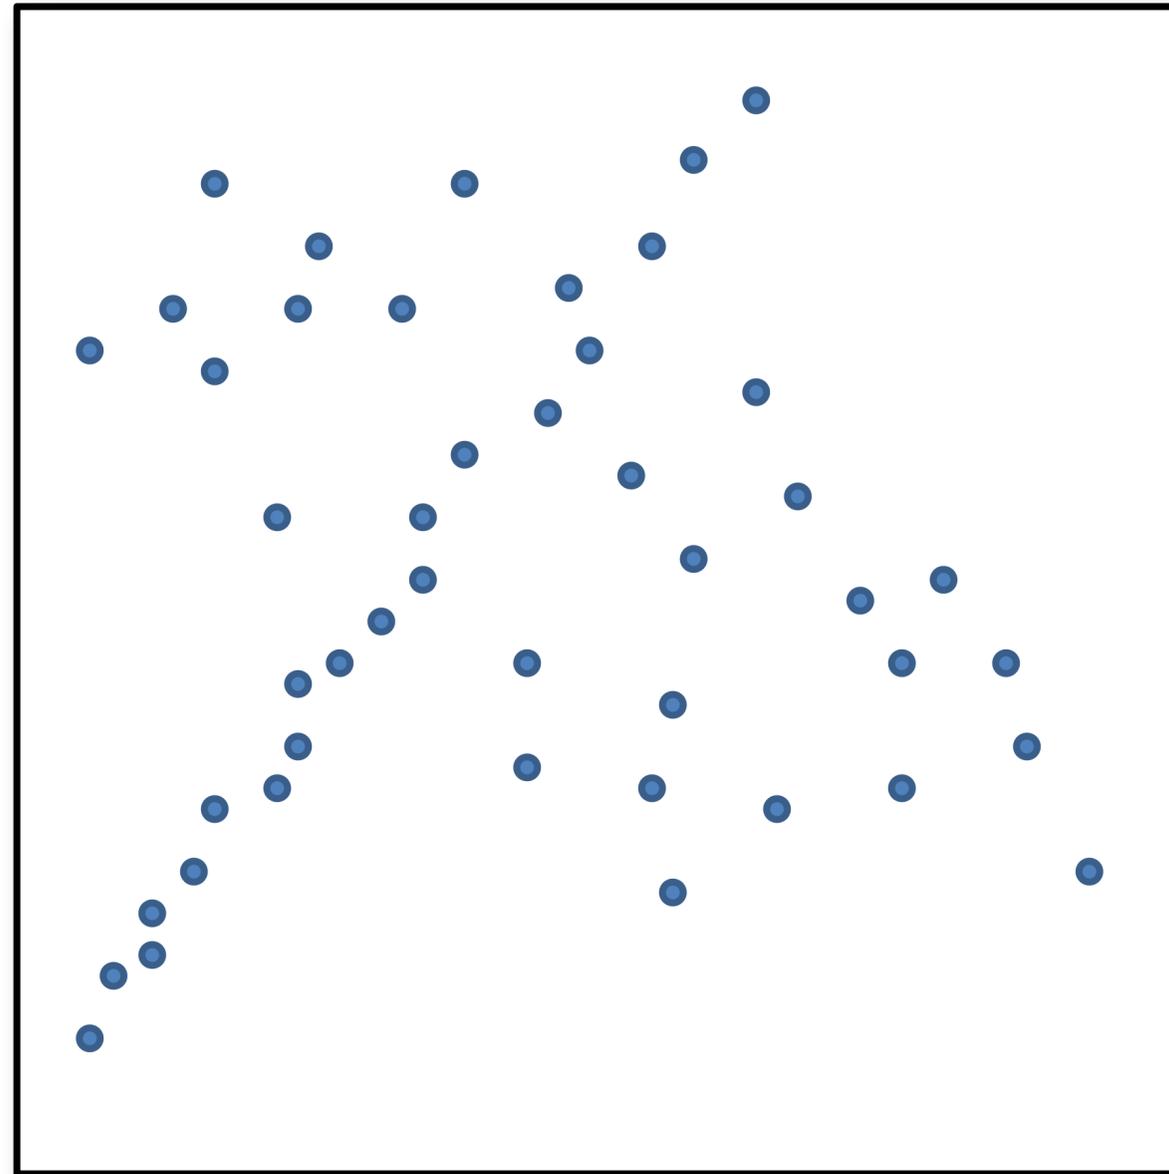
- Can be hard to fit a robust loss, e.g., due to local minima
- Another idea: trial and error!
- Let's consider the problem of linear regression



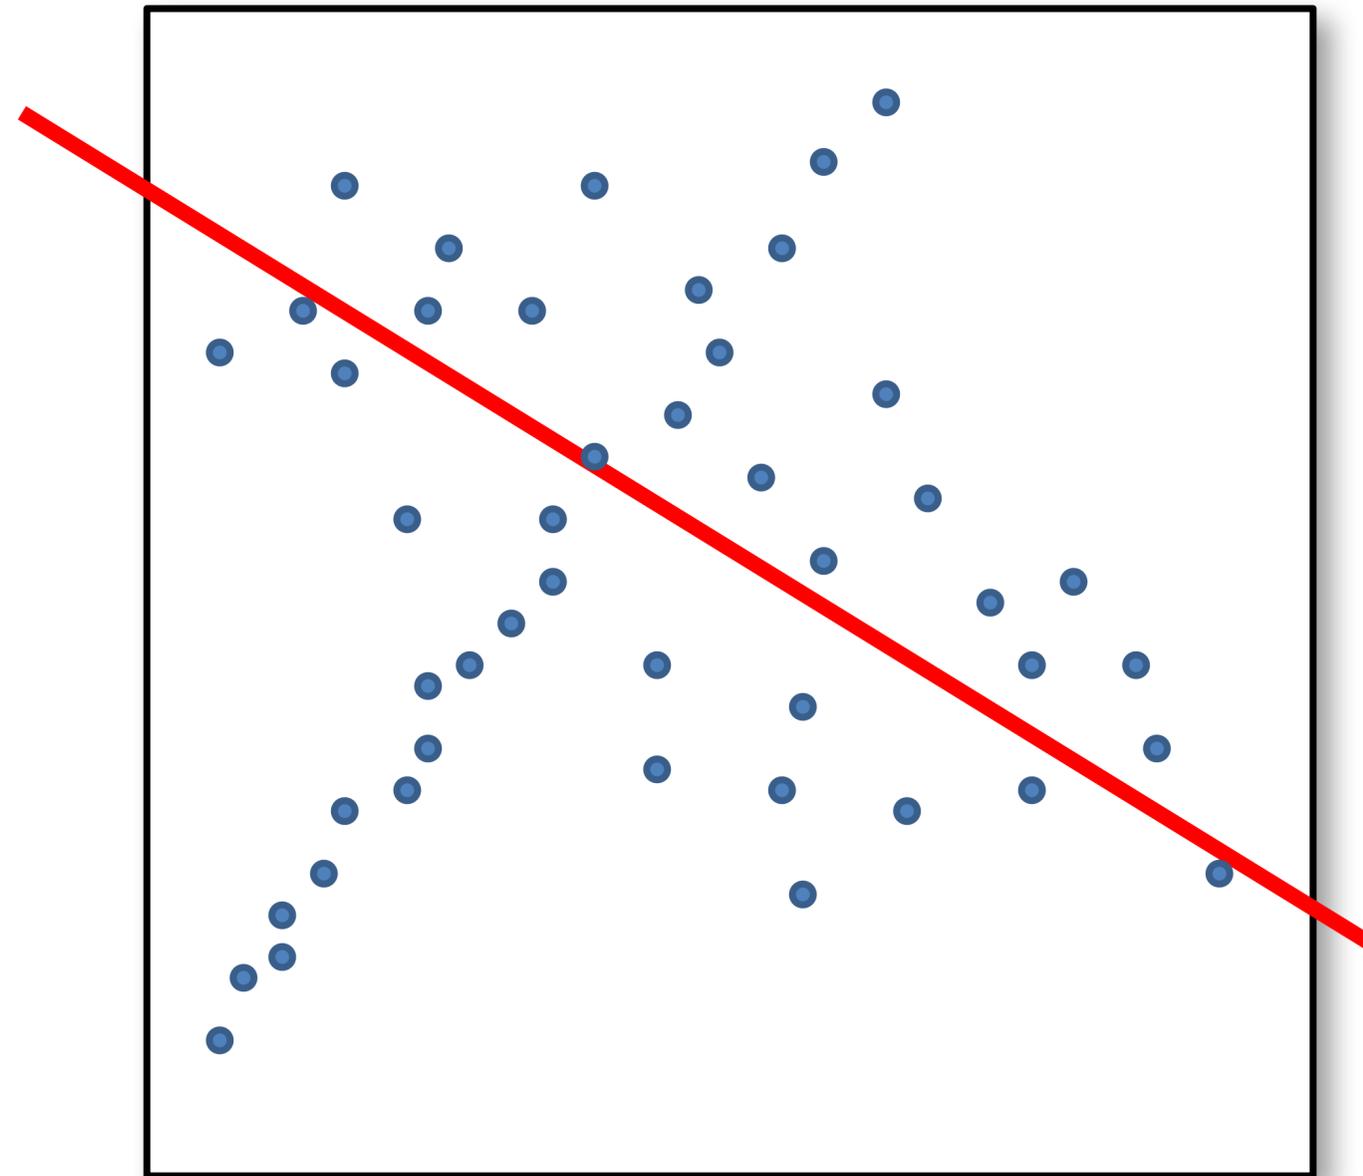
Problem: Fit a line to these data points

Least squares fit

# Counting inliers

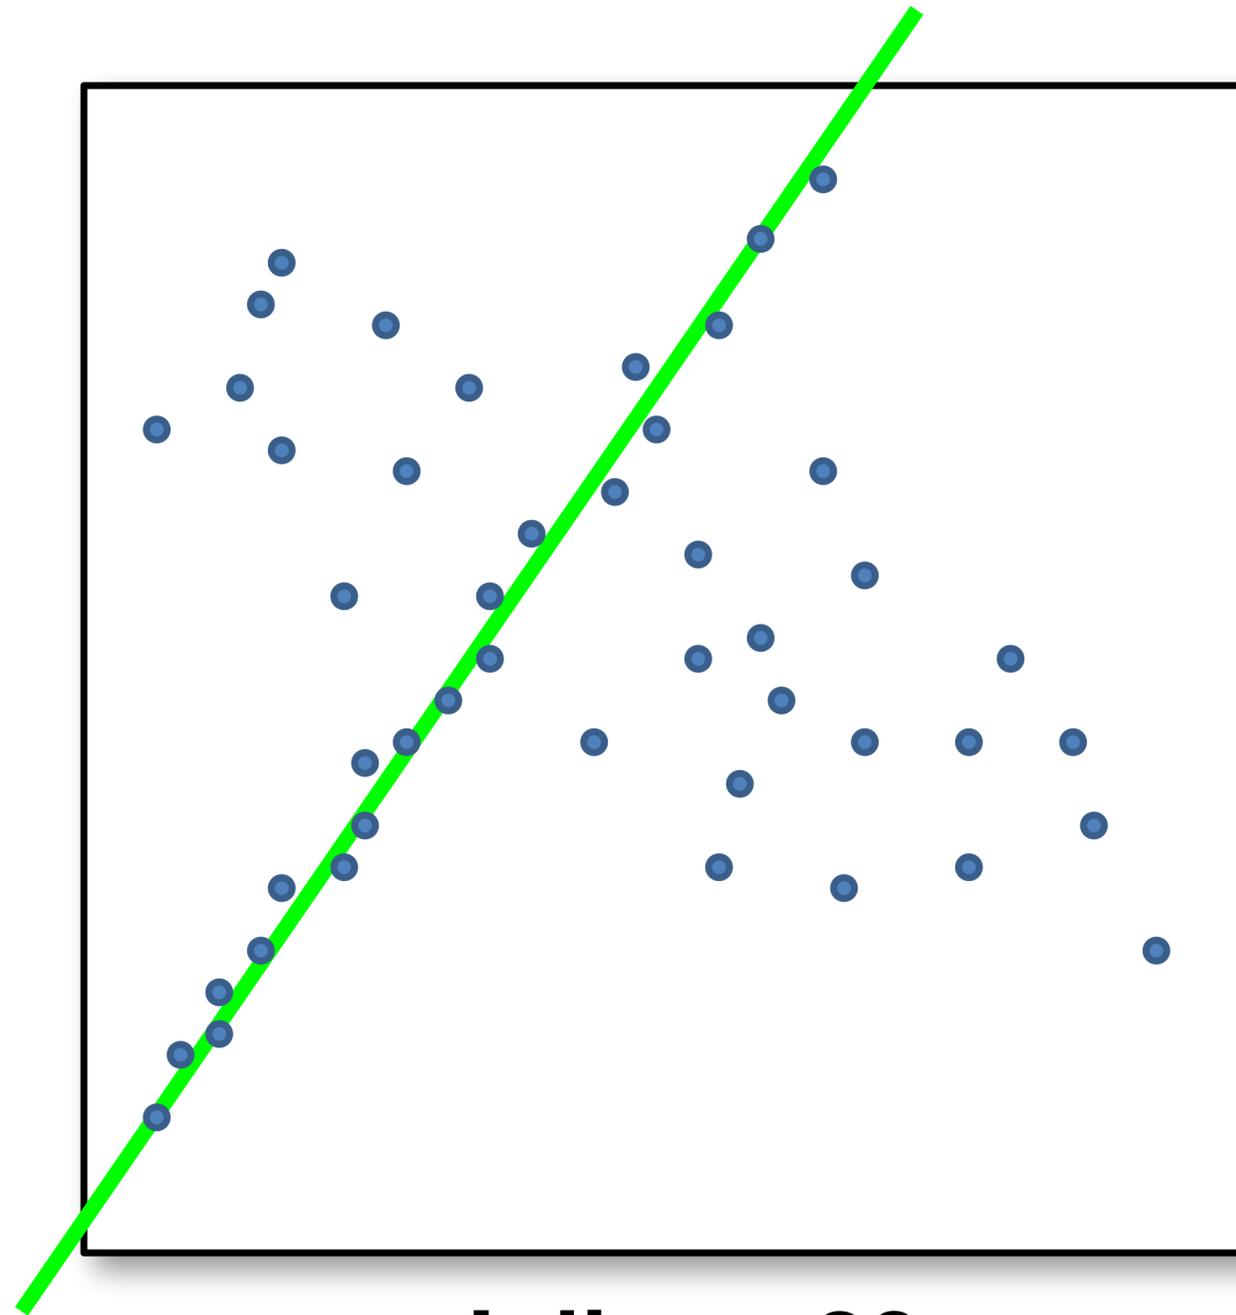


# Counting inliers



**Inliers: 3**

# Counting inliers



**Inliers: 20**

# RANSAC

- Idea:
  - All the inliers will agree with each other on the solution; the (hopefully small) number of outliers will (hopefully) disagree with each other
    - RANSAC only has guarantees if there are  $< 50\%$  outliers
  - “All good matches are alike; every bad match is bad in its own way.”
    - Tolstoy via Alyosha Efros

# RANSAC: random sample consensus

RANSAC loop (for N iterations):

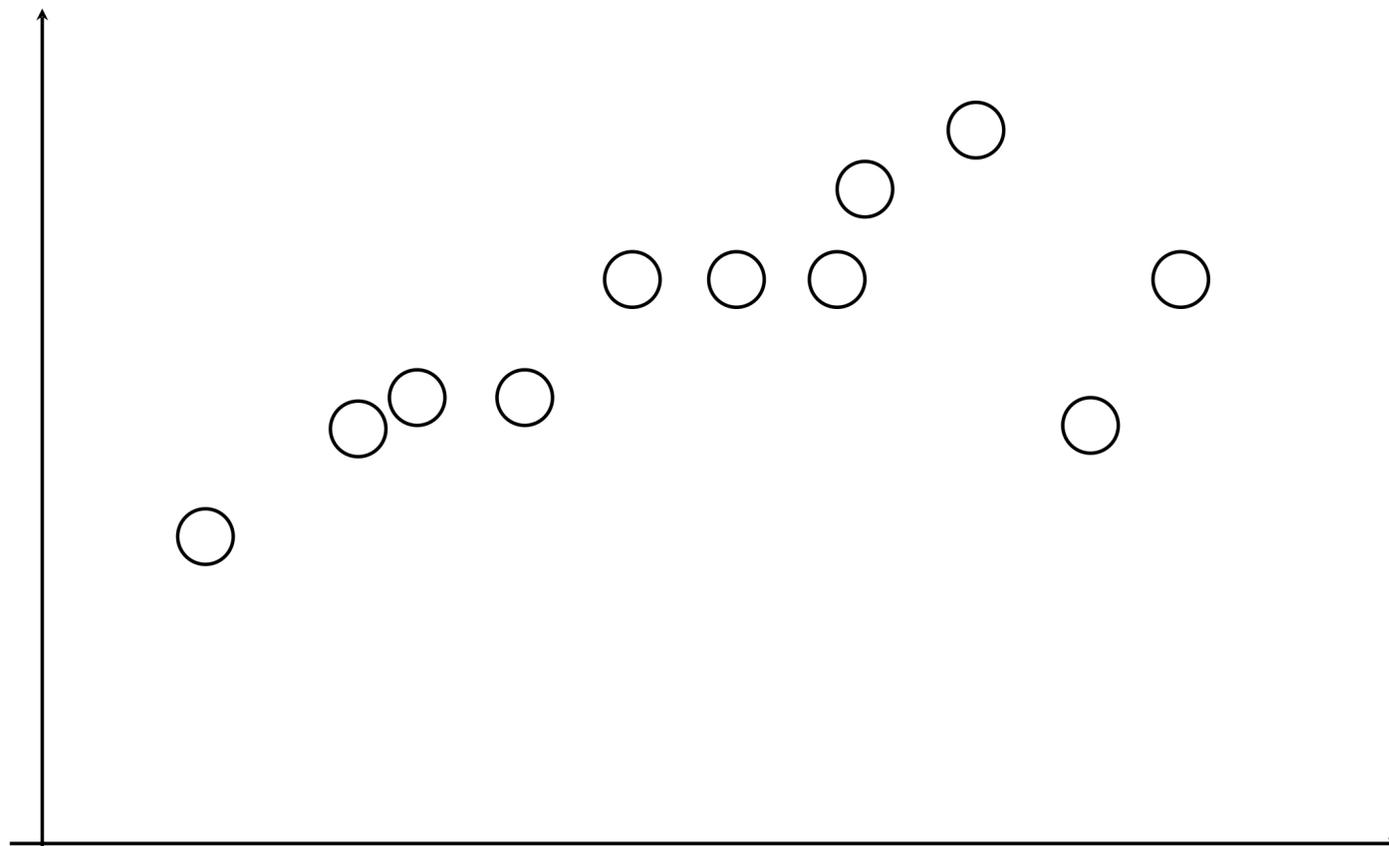
- Select four feature pairs (at random)
- Compute homography  $H$
- Count **inliers** where  $\|p_i' - f_H(p_i)\| < \varepsilon$

Afterwards:

- Choose  $H$  with largest set of inliers
- Recompute  $H$  using only those inliers (often using high-quality nonlinear least squares)

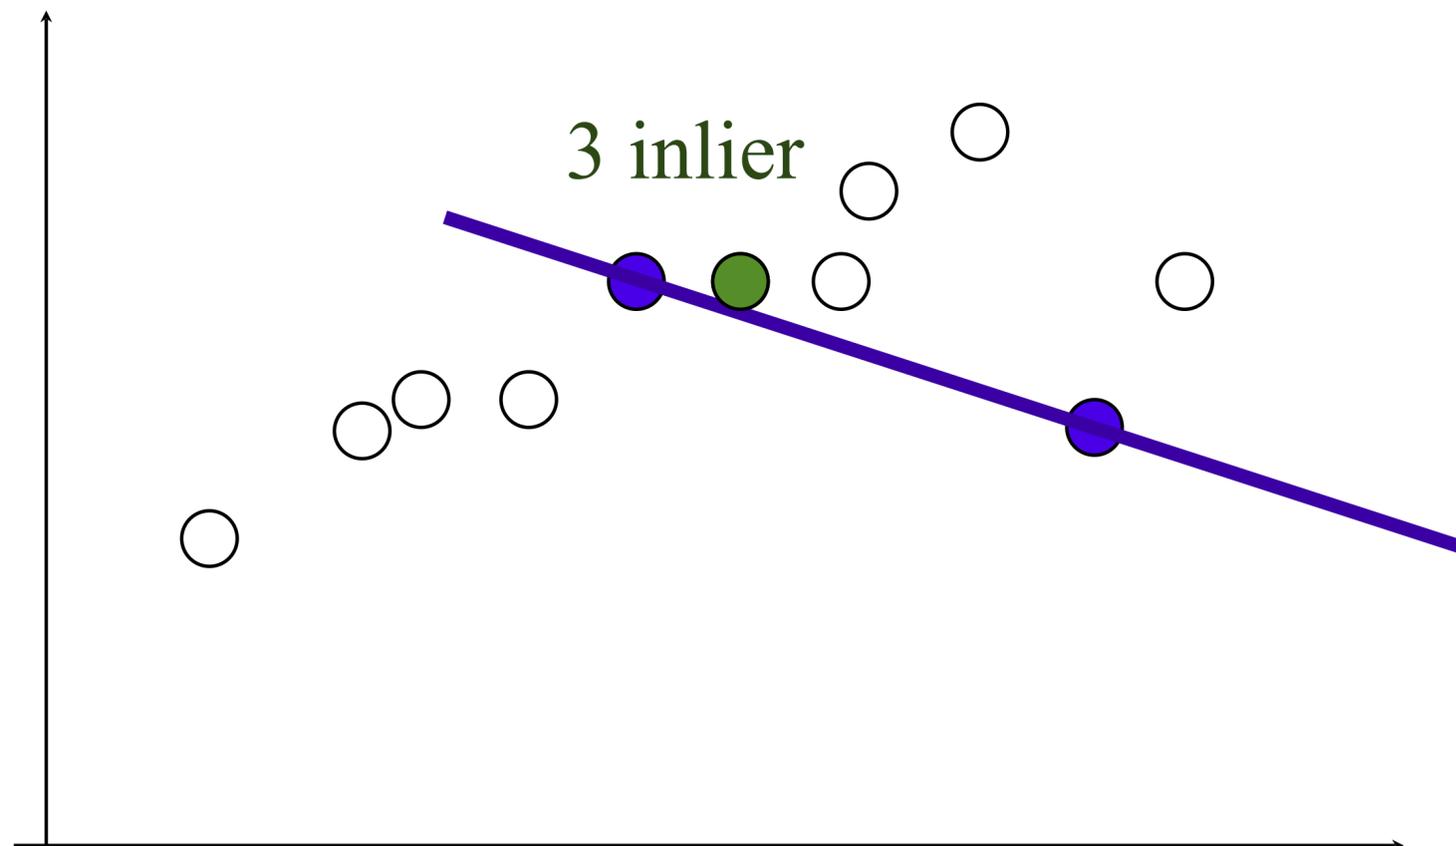
# Simple example: fit a line

- Rather than homography  $H$  (8 numbers) fit  $y=ax+b$  (2 numbers  $a, b$ ) to 2D pairs



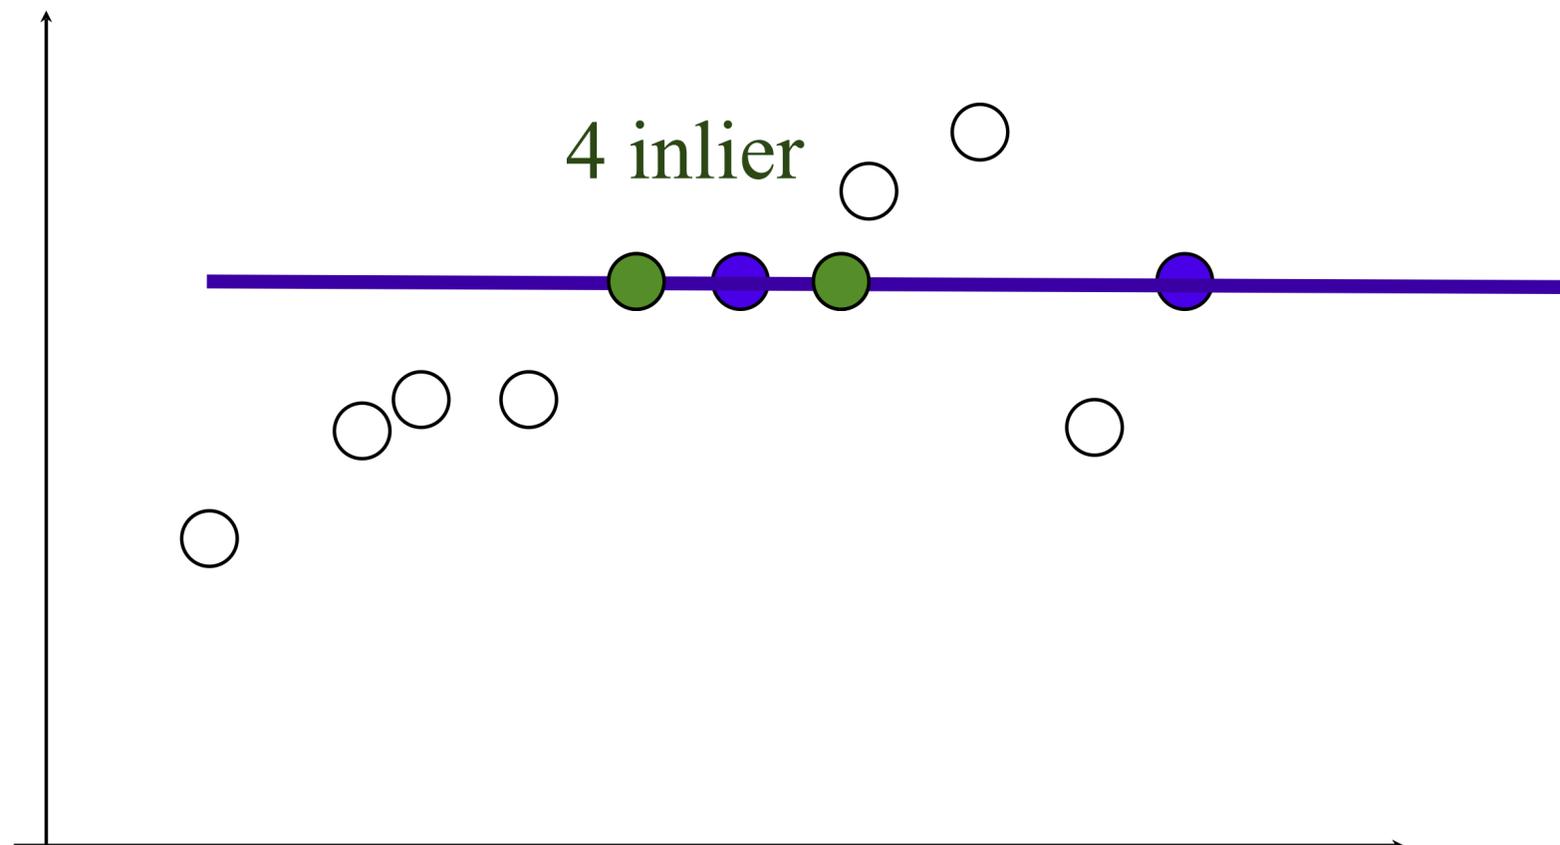
# Simple example: fit a line

- Pick 2 points
- Fit line
- Count inliers



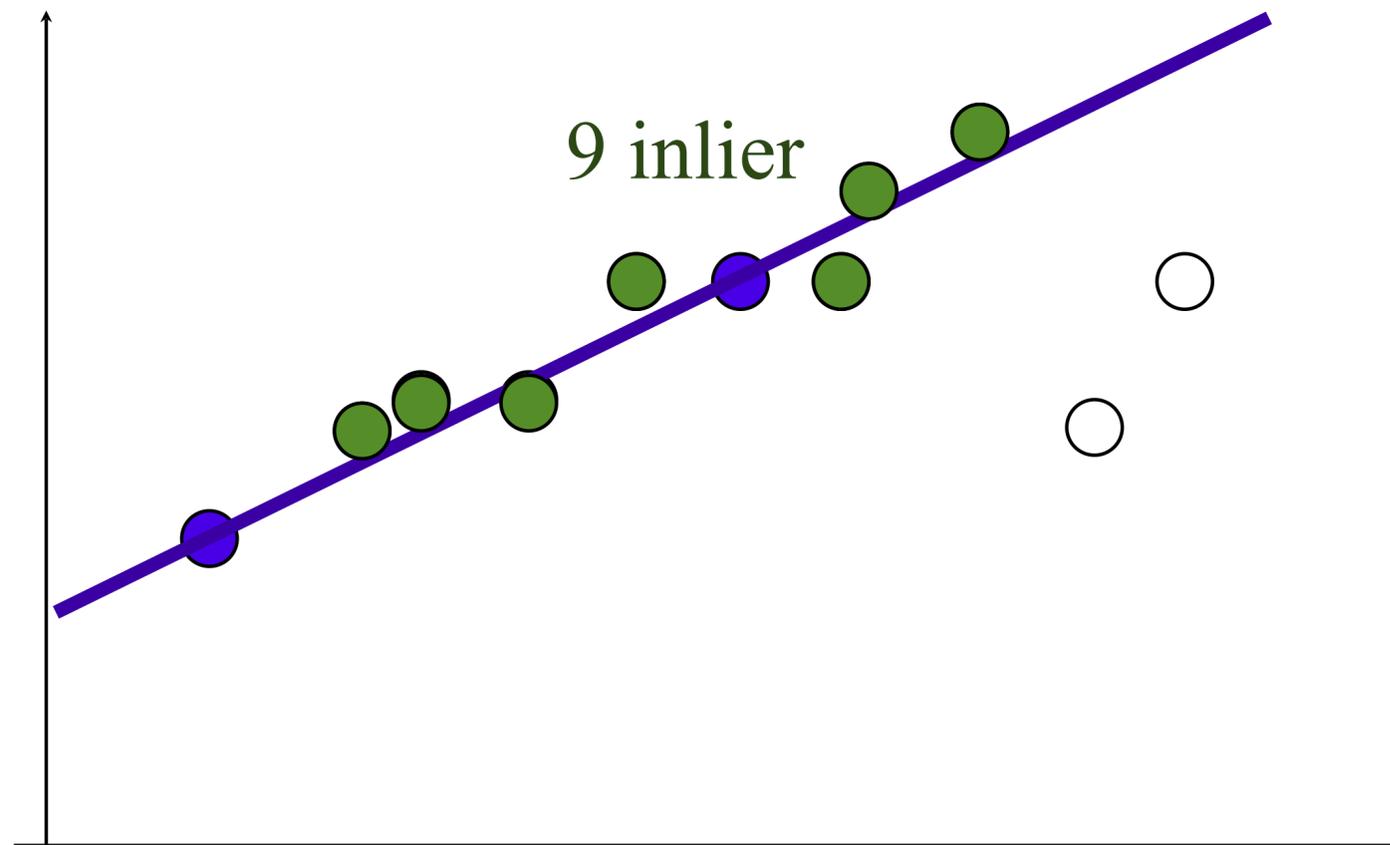
# Simple example: fit a line

- Pick 2 points
- Fit line
- Count inliers



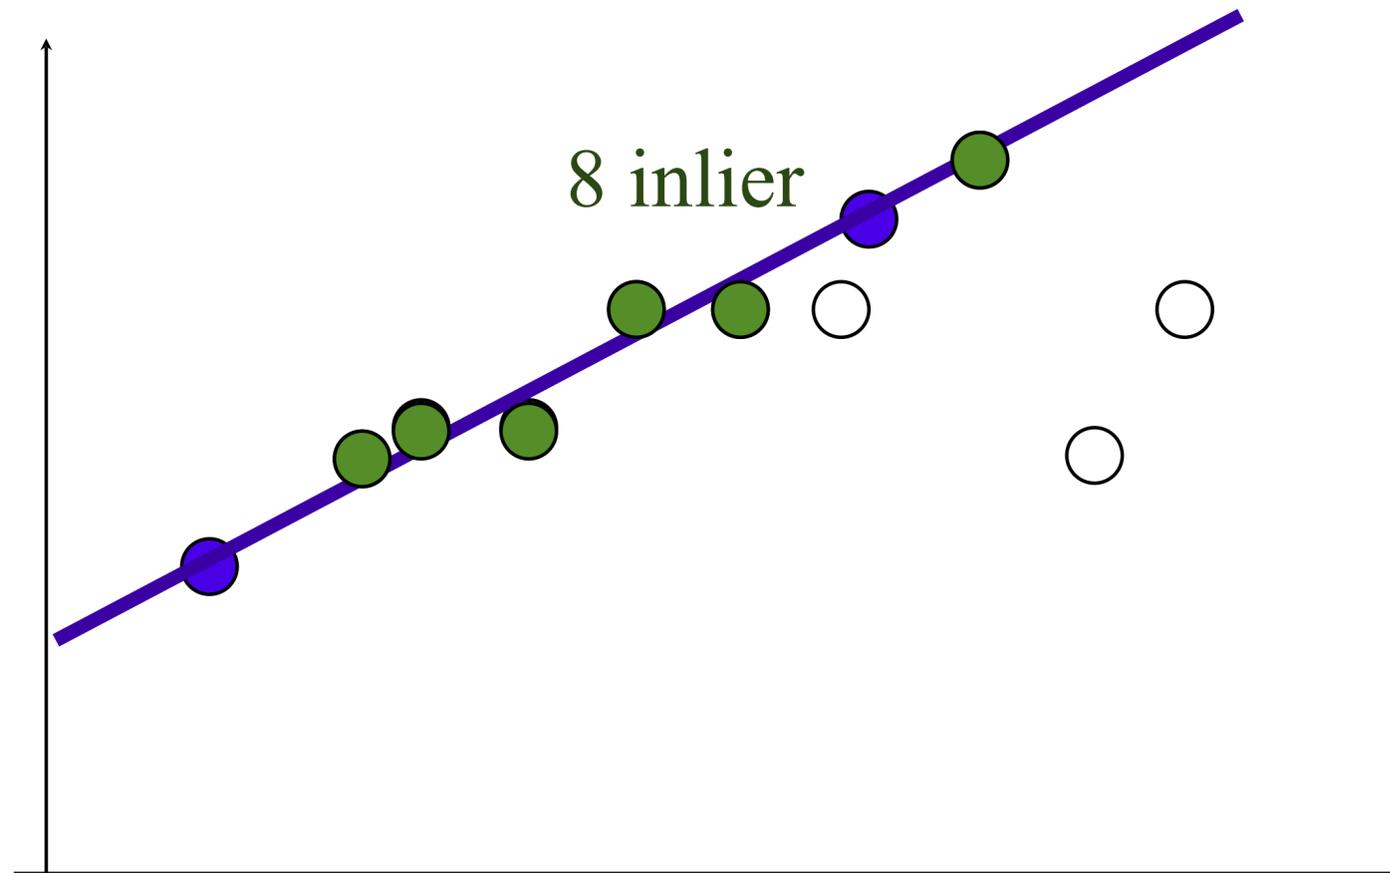
# Simple example: fit a line

- Pick 2 points
- Fit line
- Count inliers



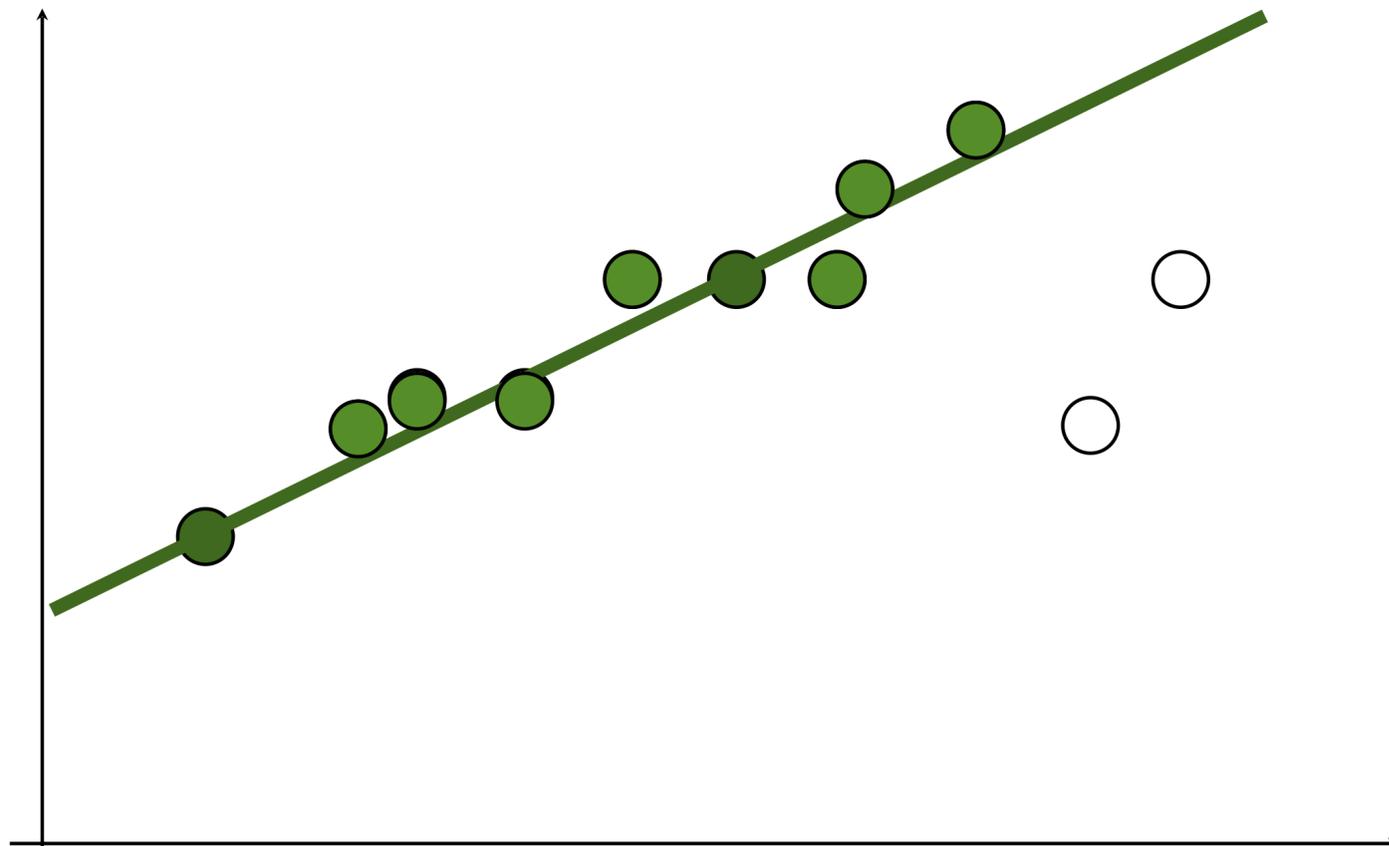
# Simple example: fit a line

- Pick 2 points
- Fit line
- Count inliers

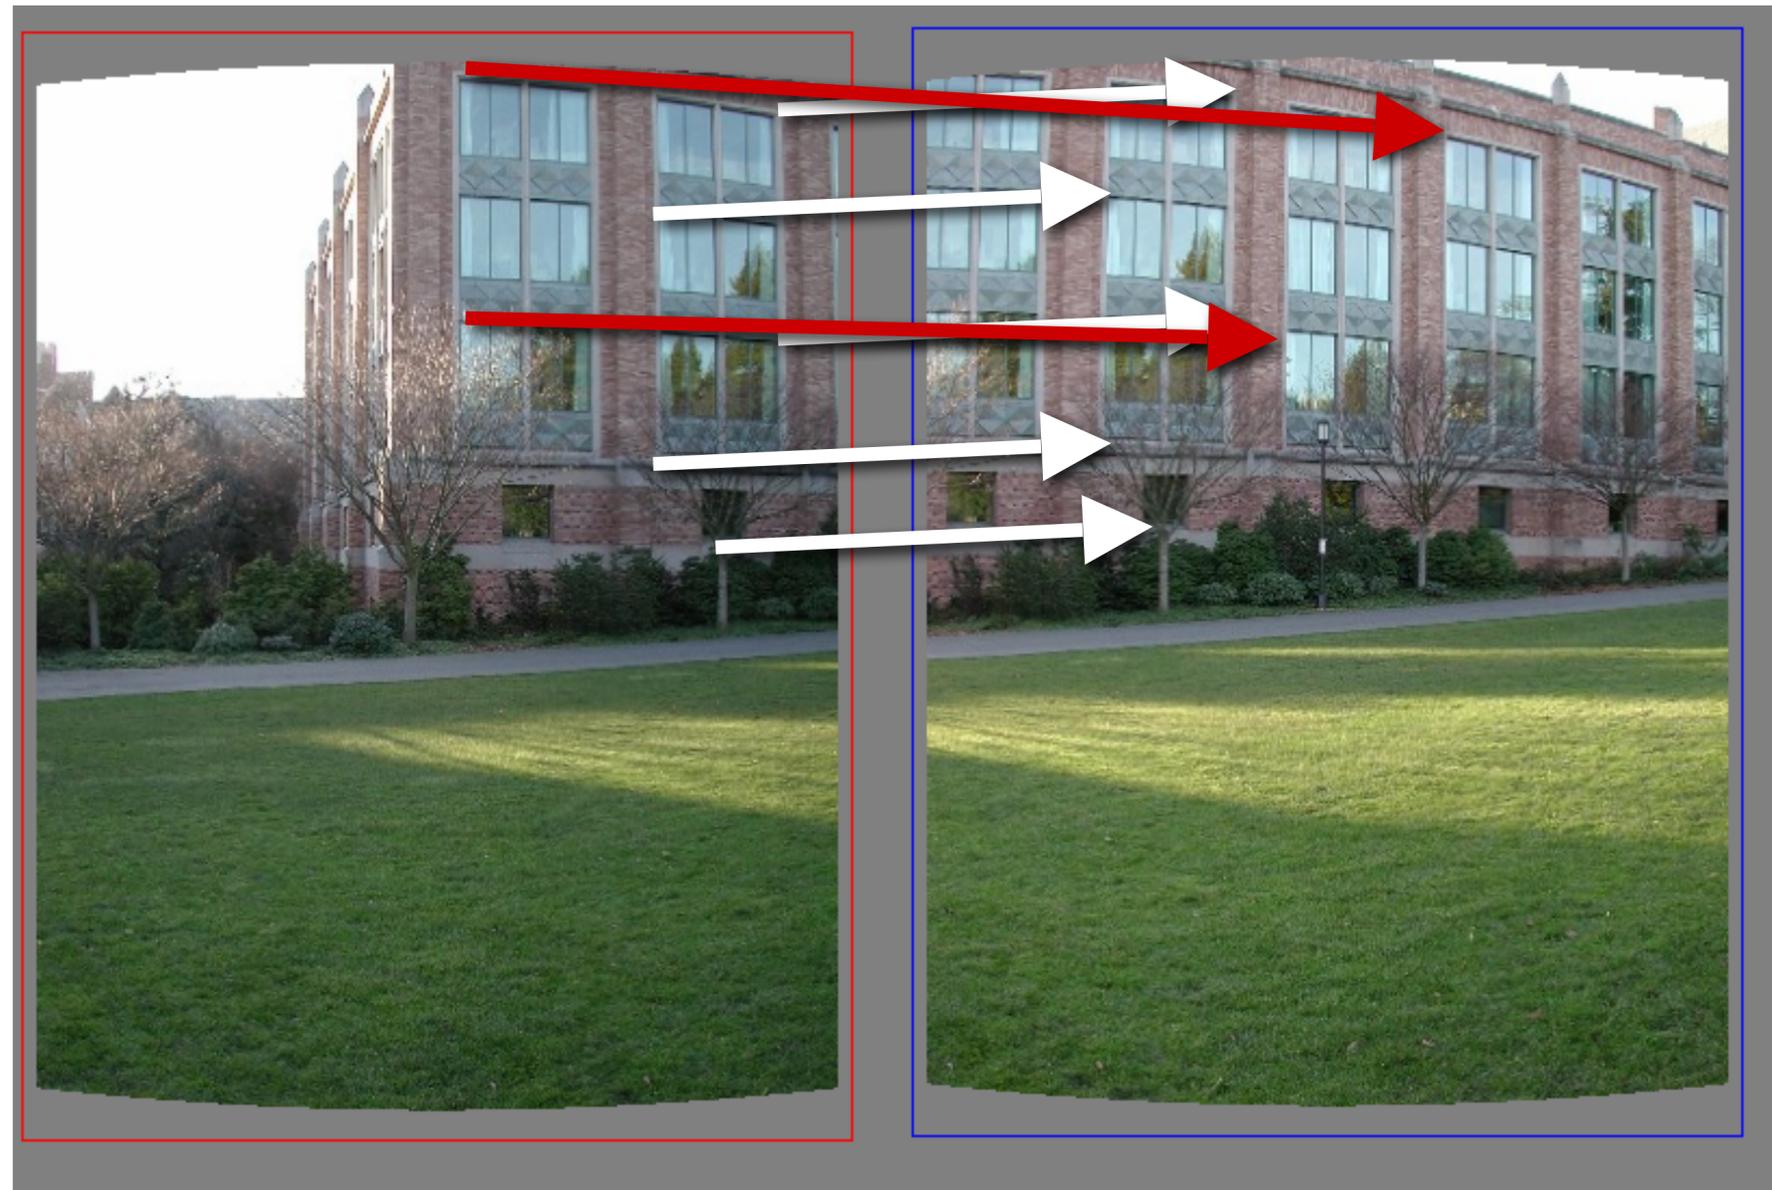


# Simple example: fit a line

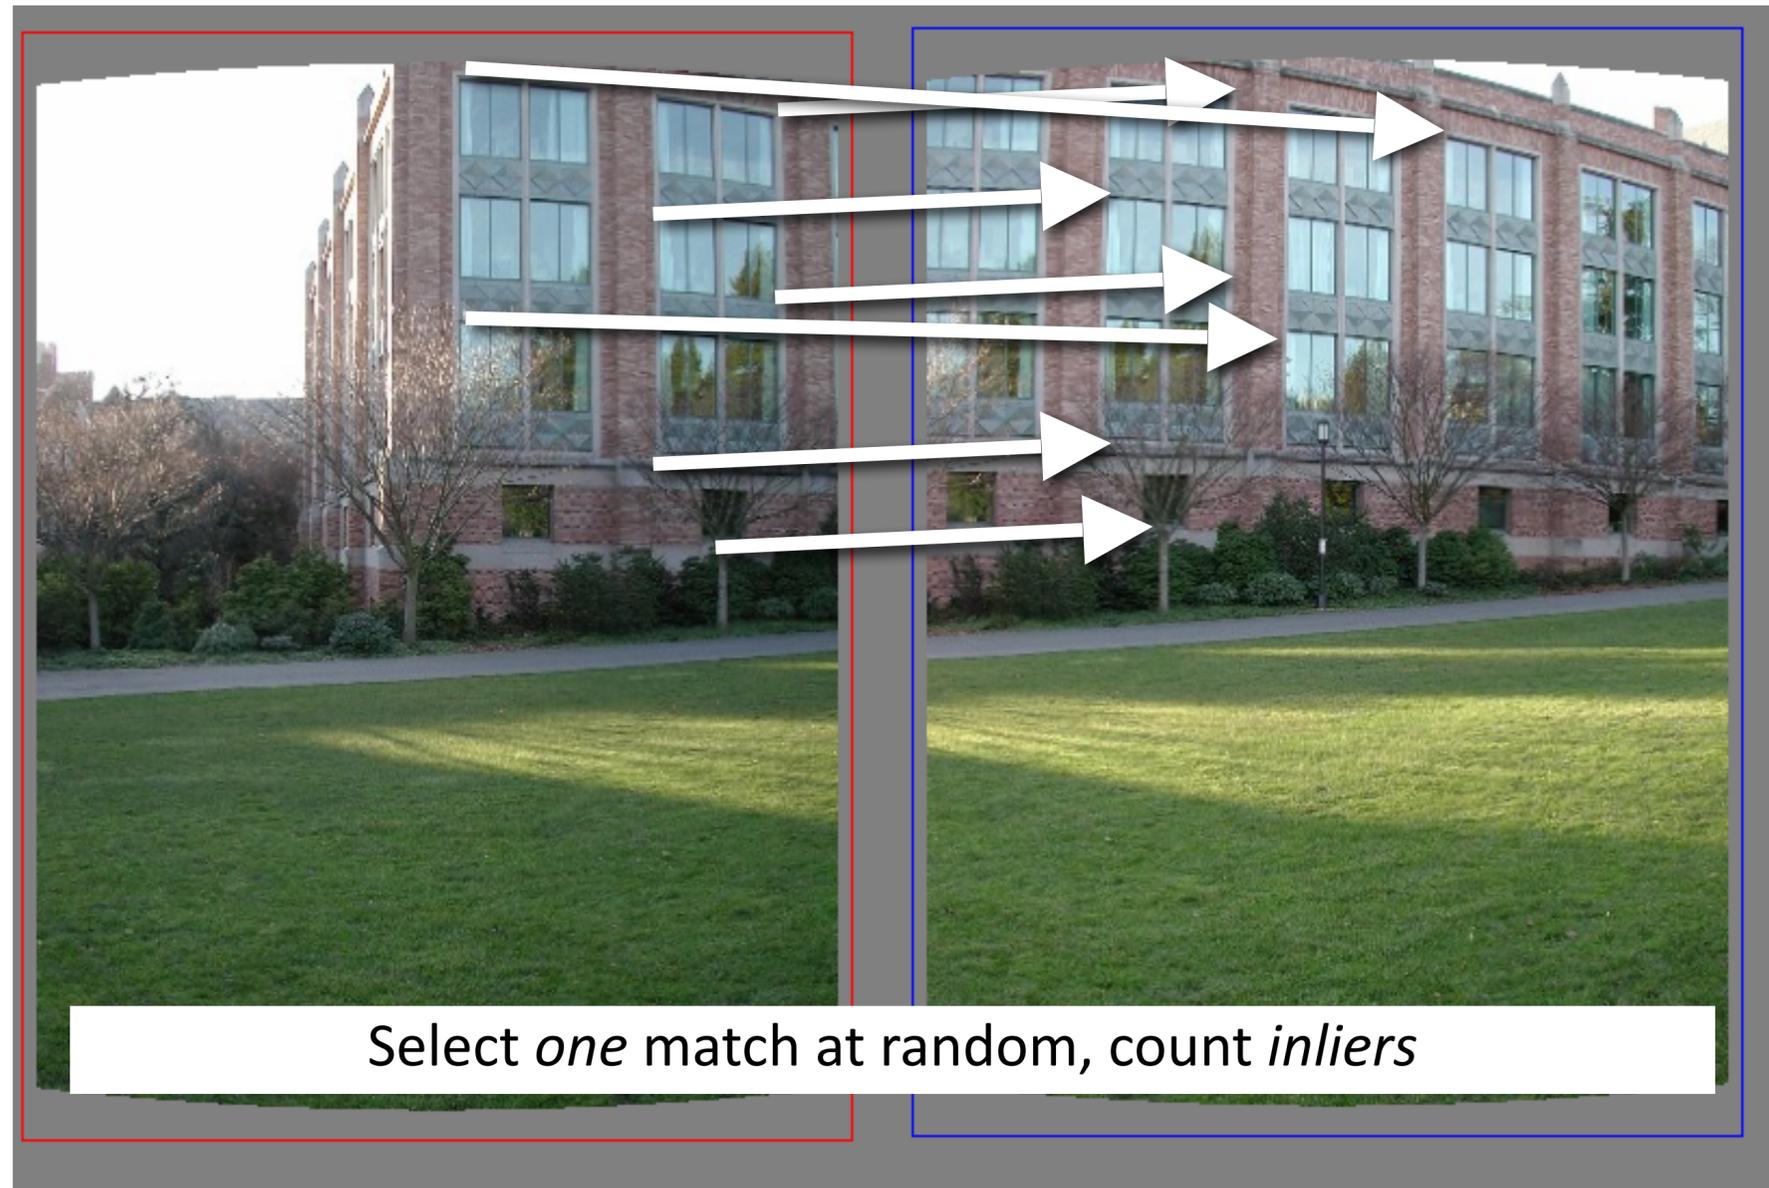
- Use biggest set of inliers
- Do least-square fit



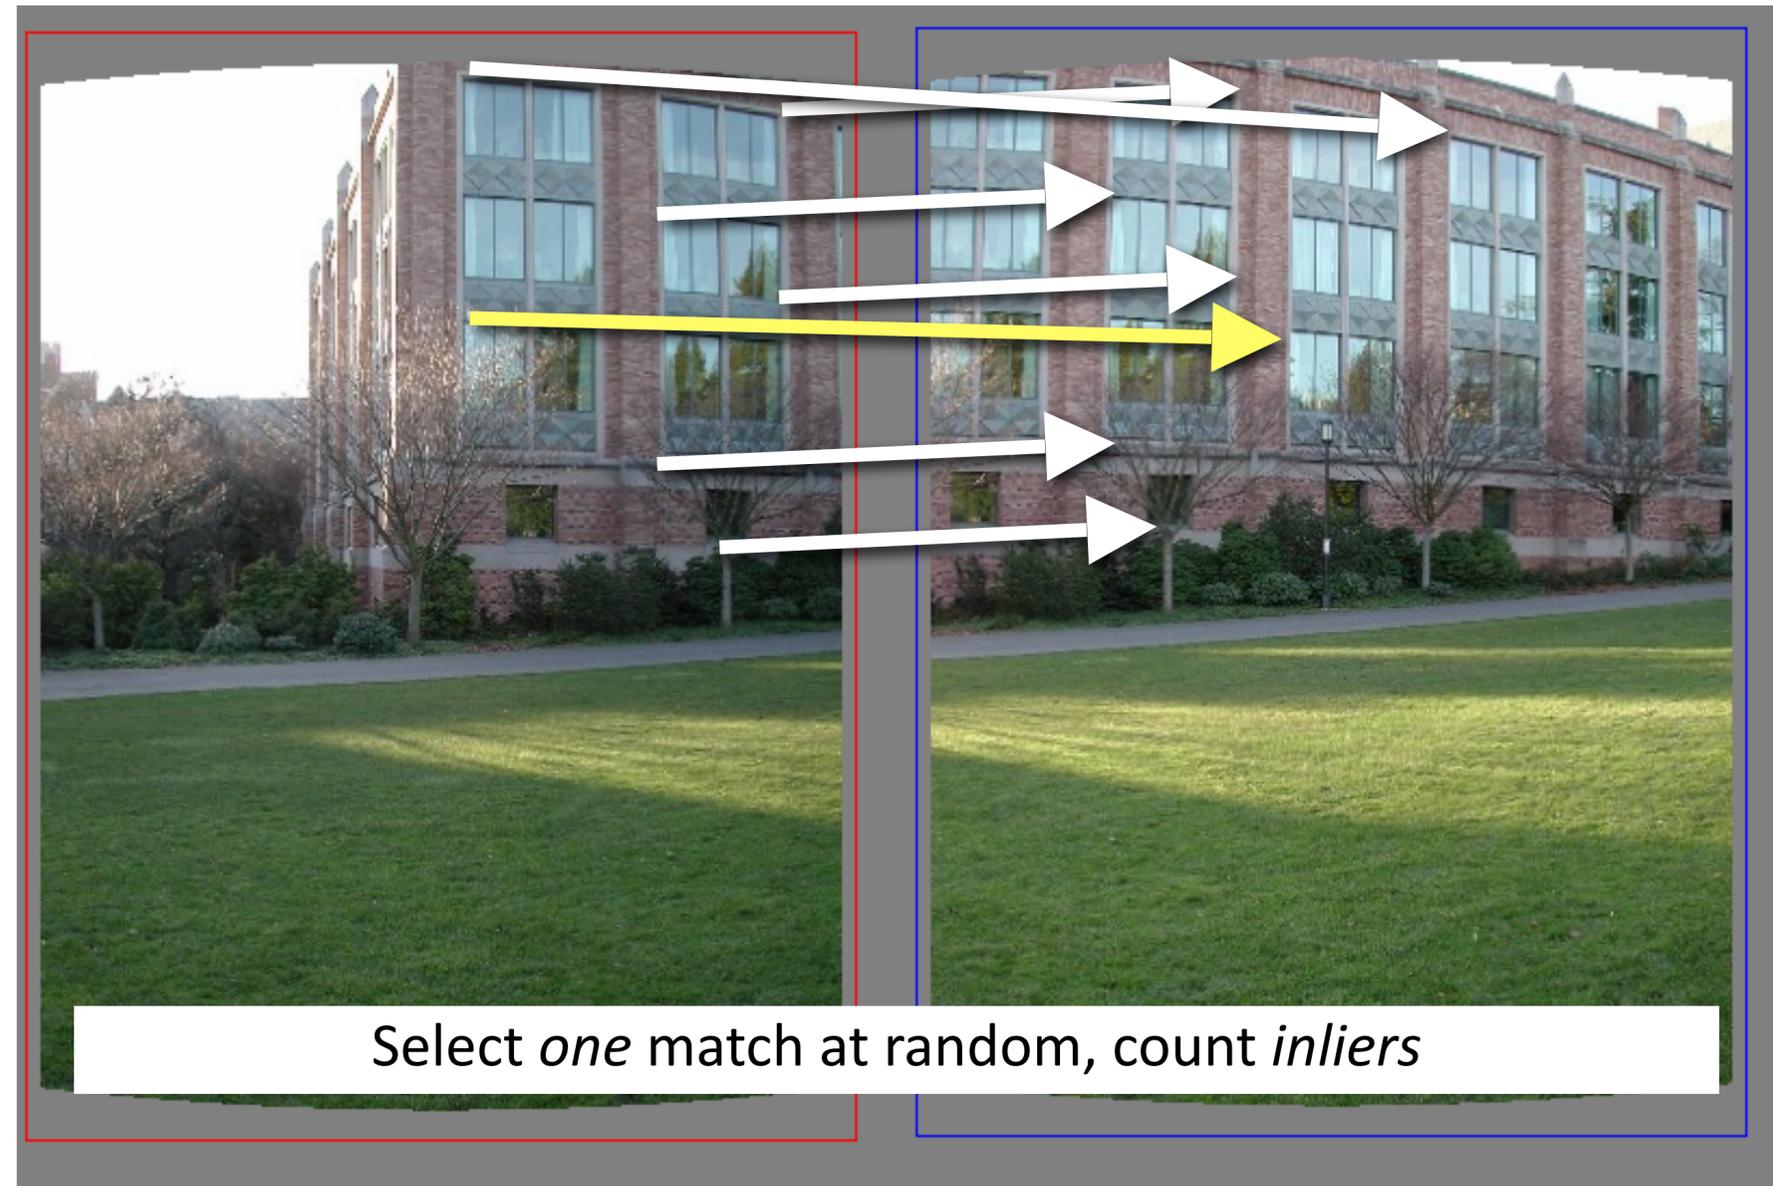
# Example: fitting a translation



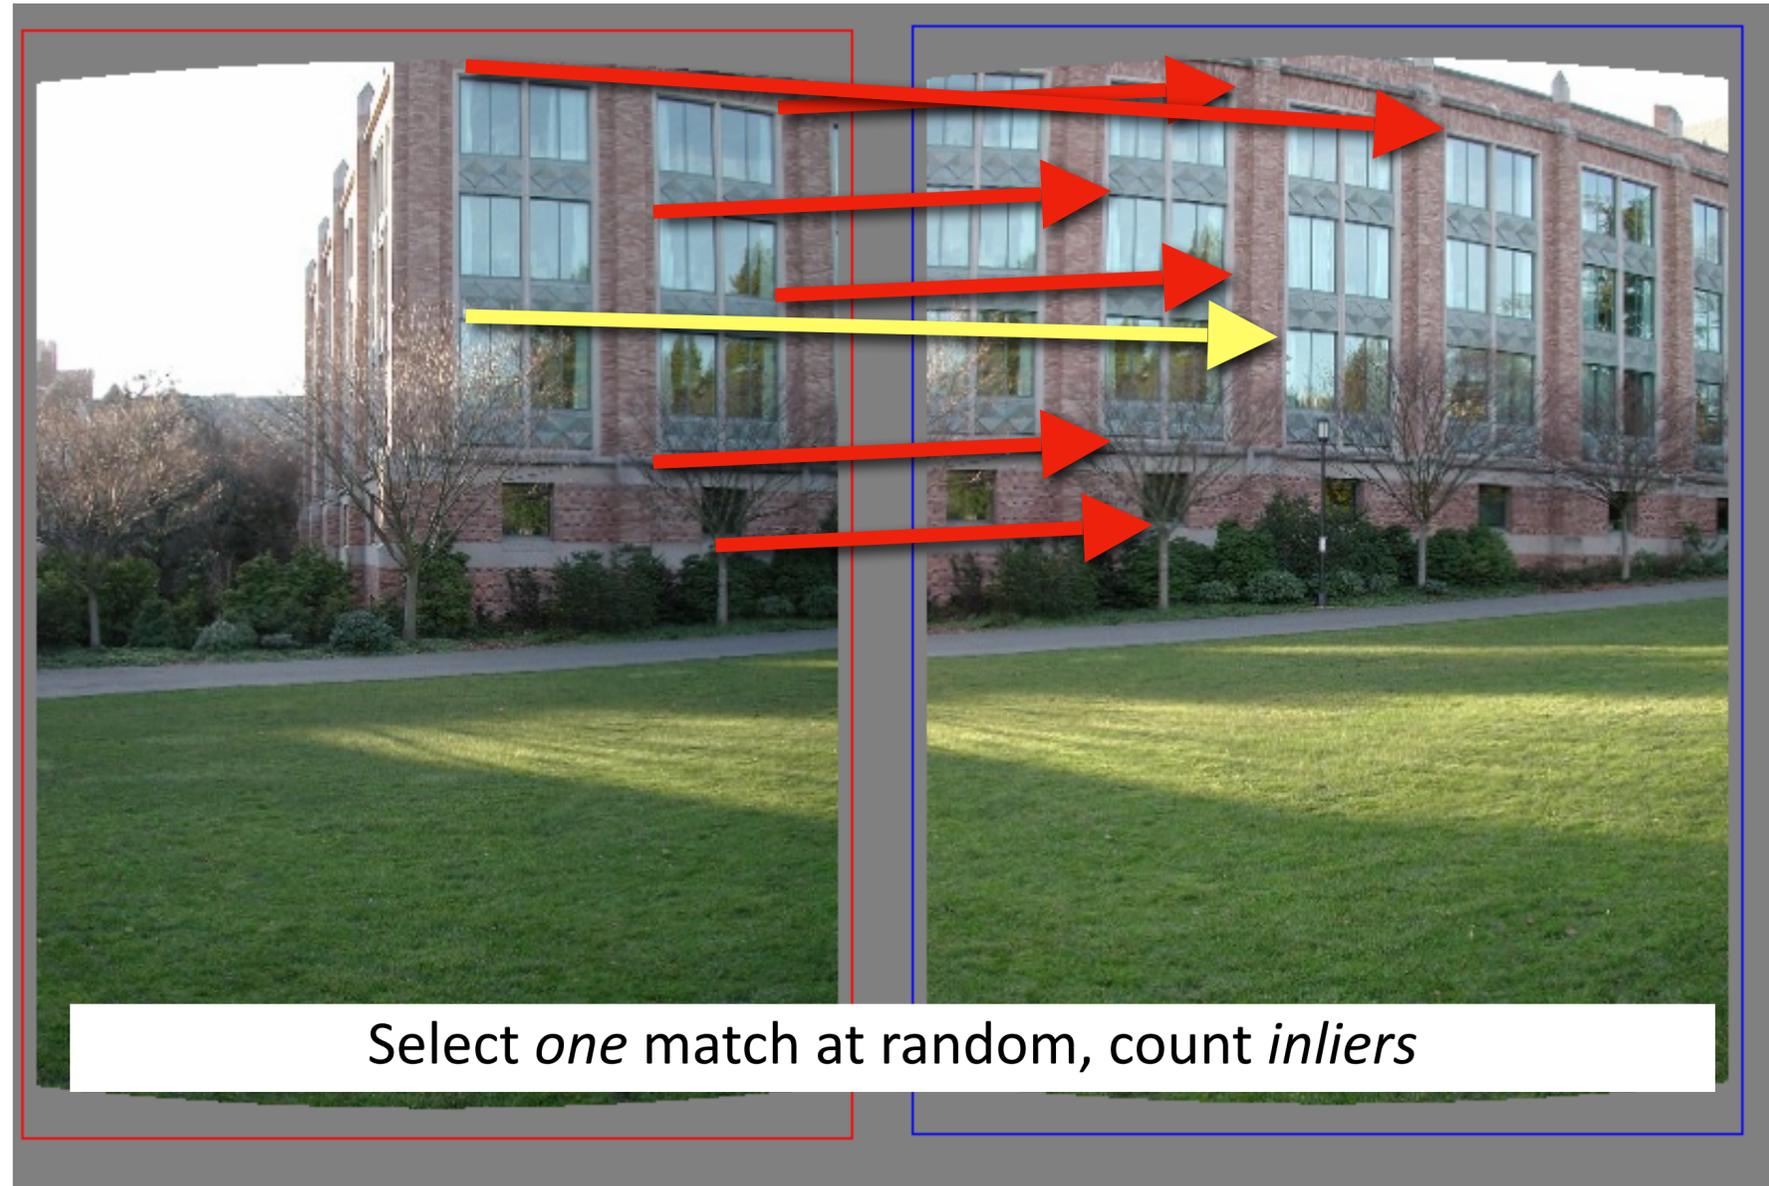
# Random Sample Consensus



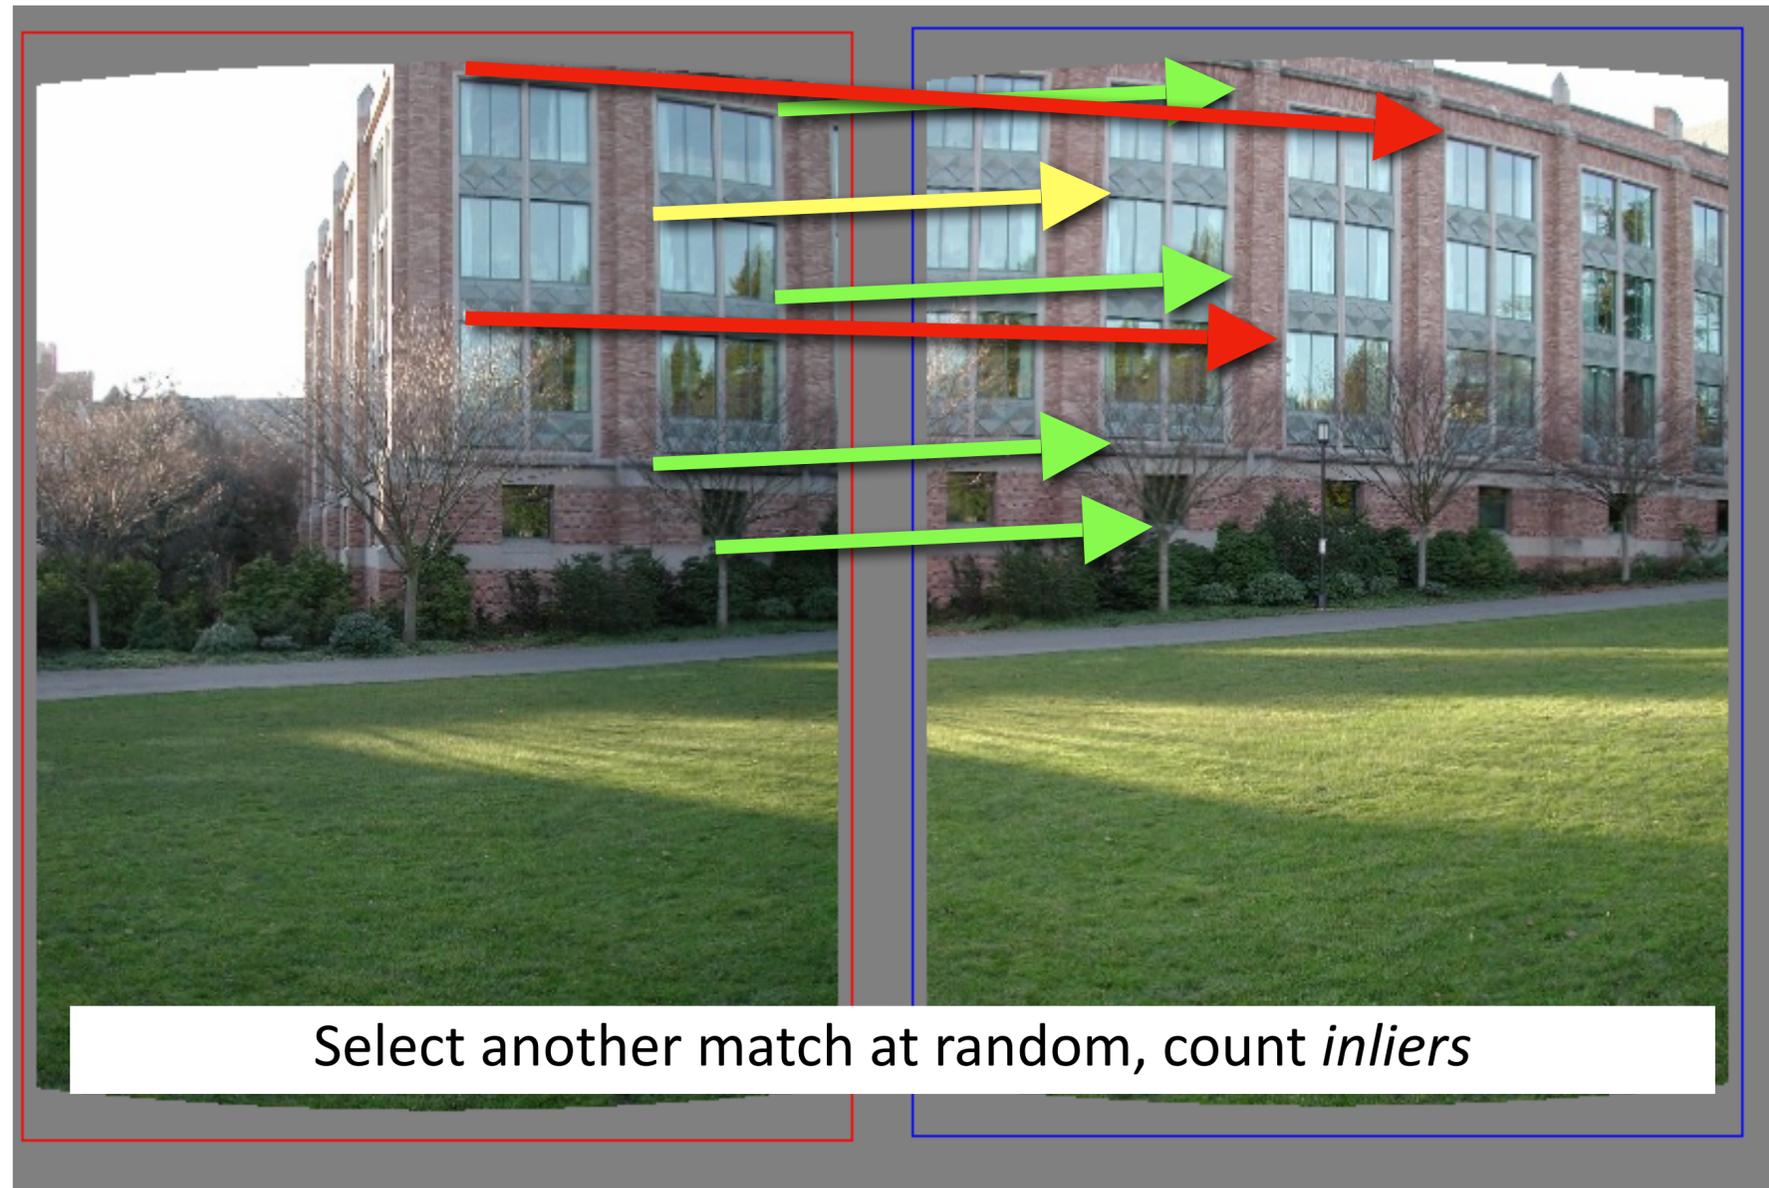
# Random Sample Consensus



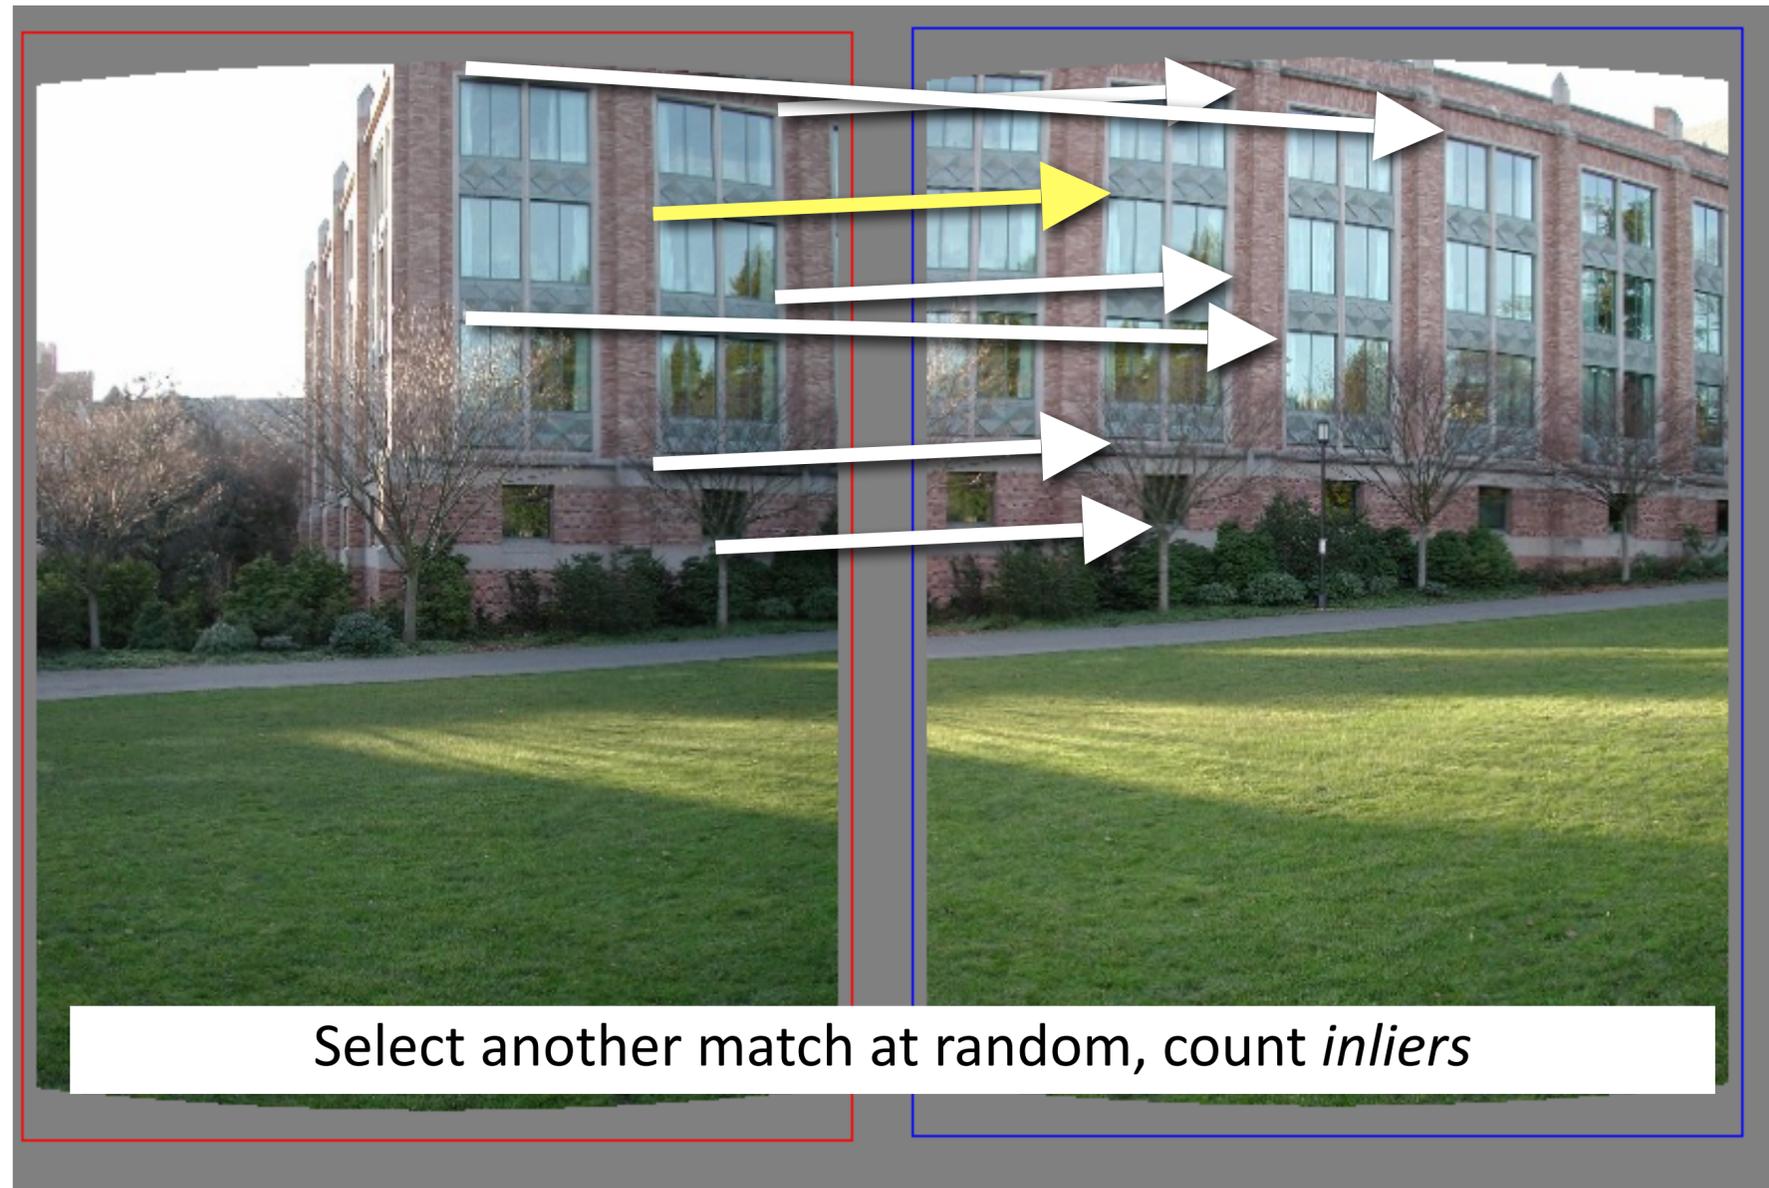
# Random Sample Consensus



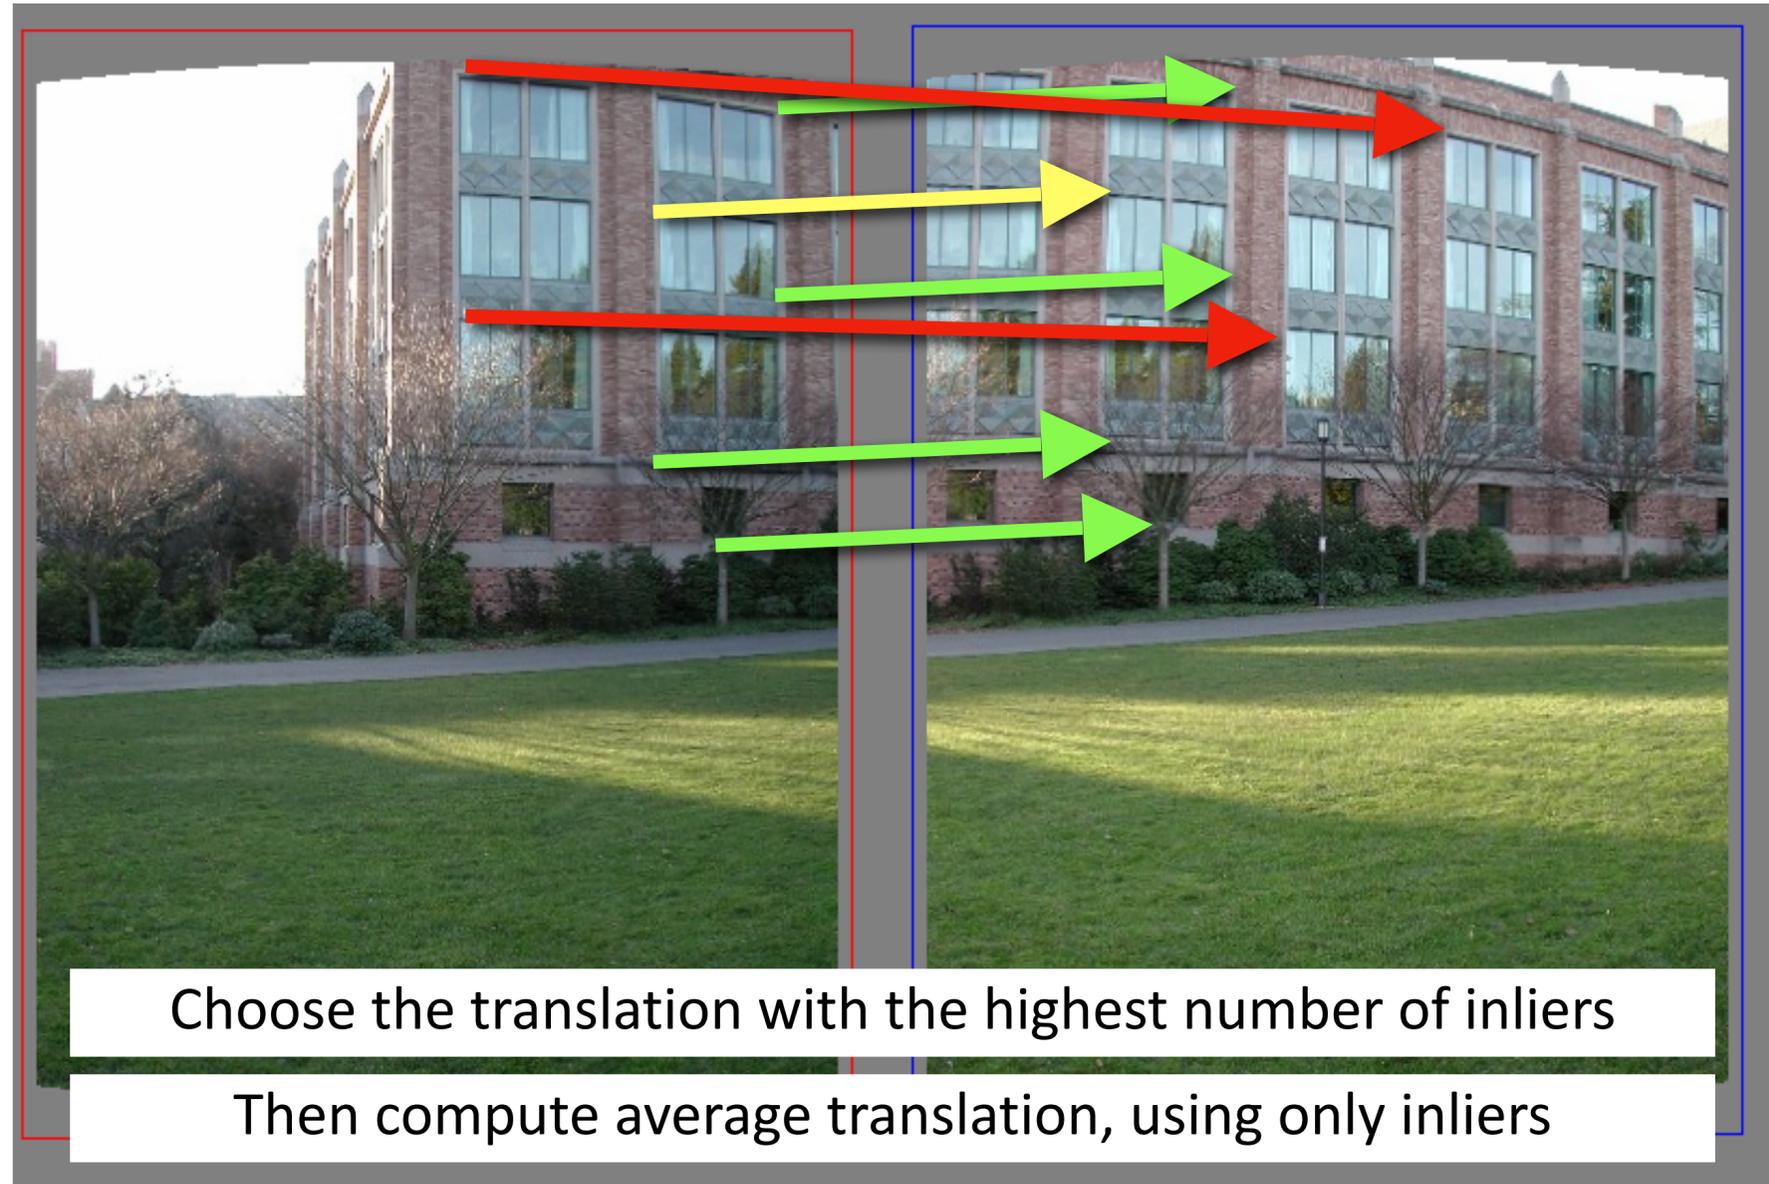
# Random Sample Consensus



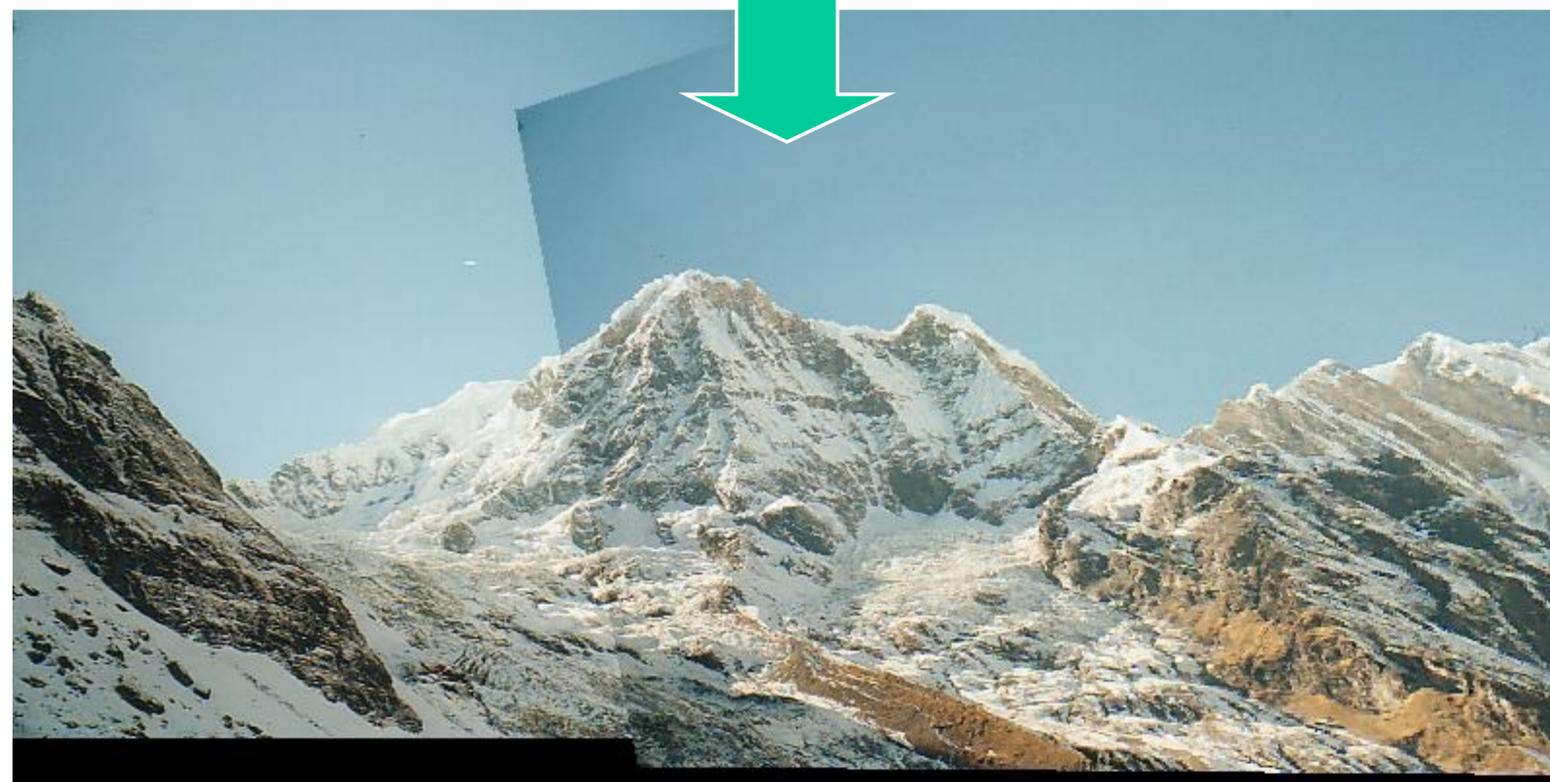
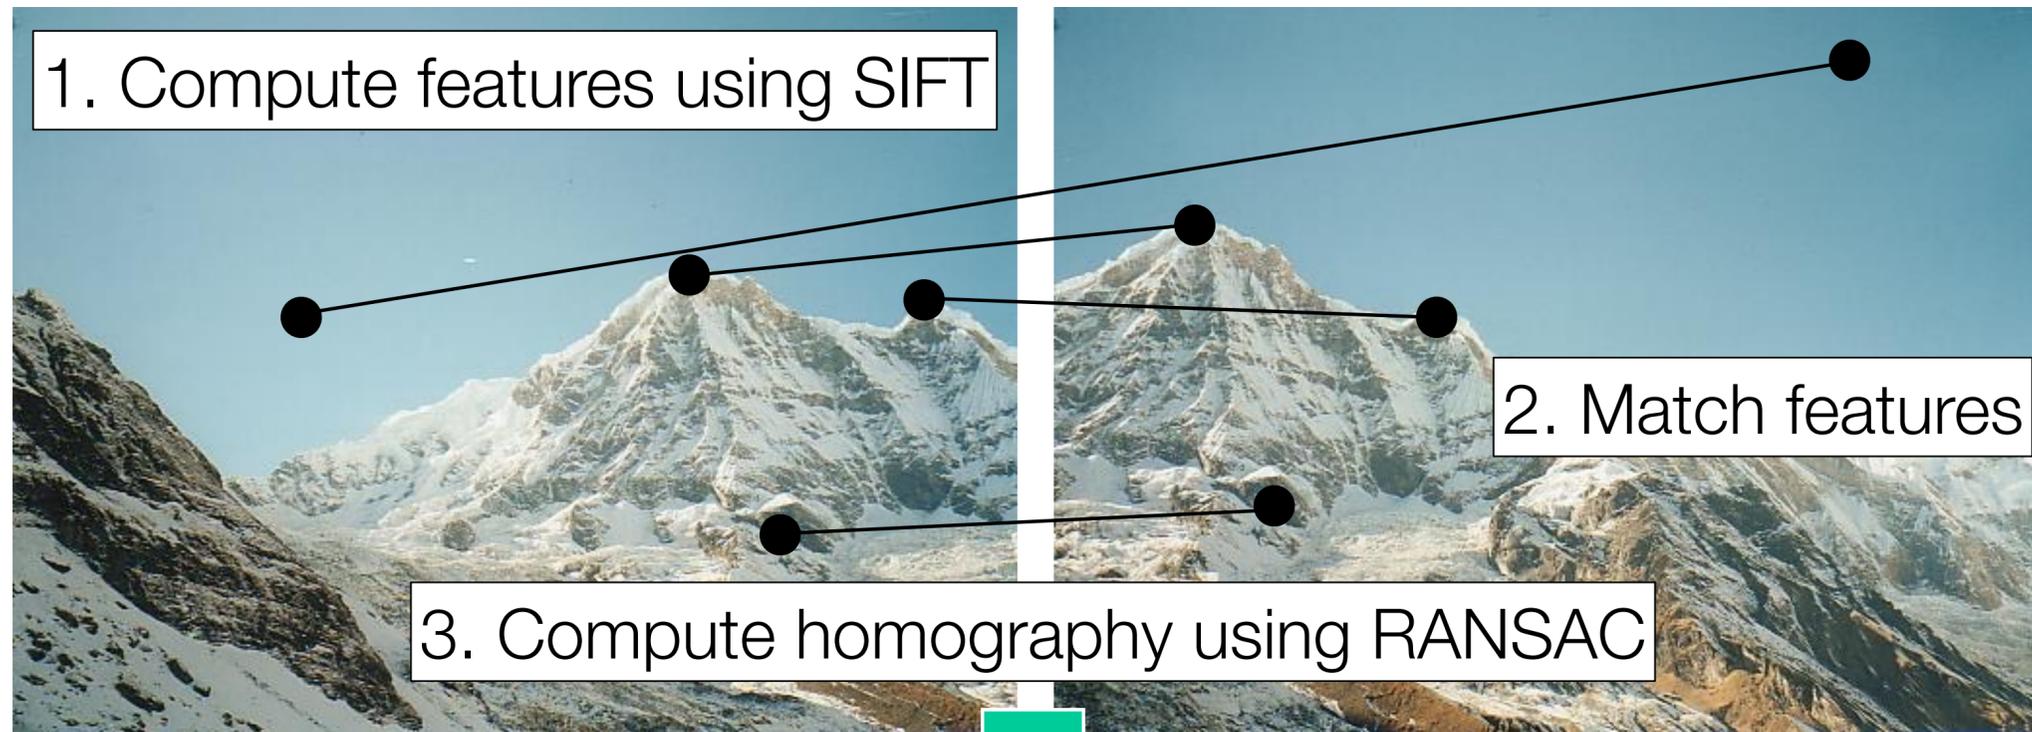
# Random Sample Consensus



# Random Sample Consensus

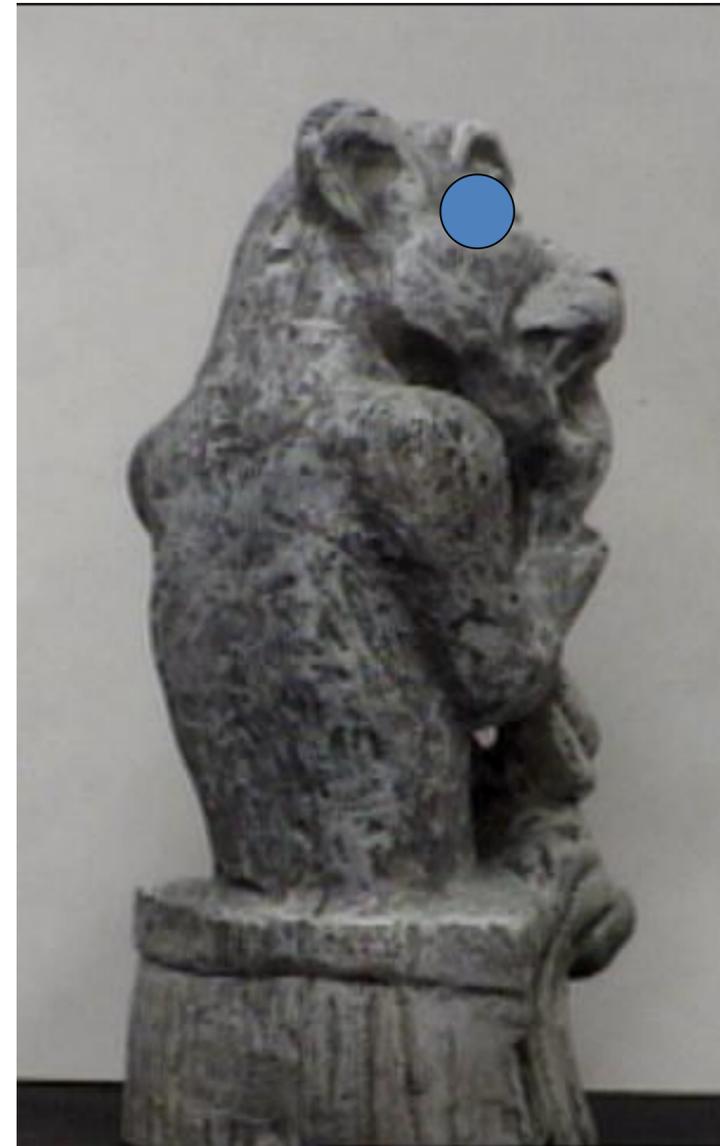
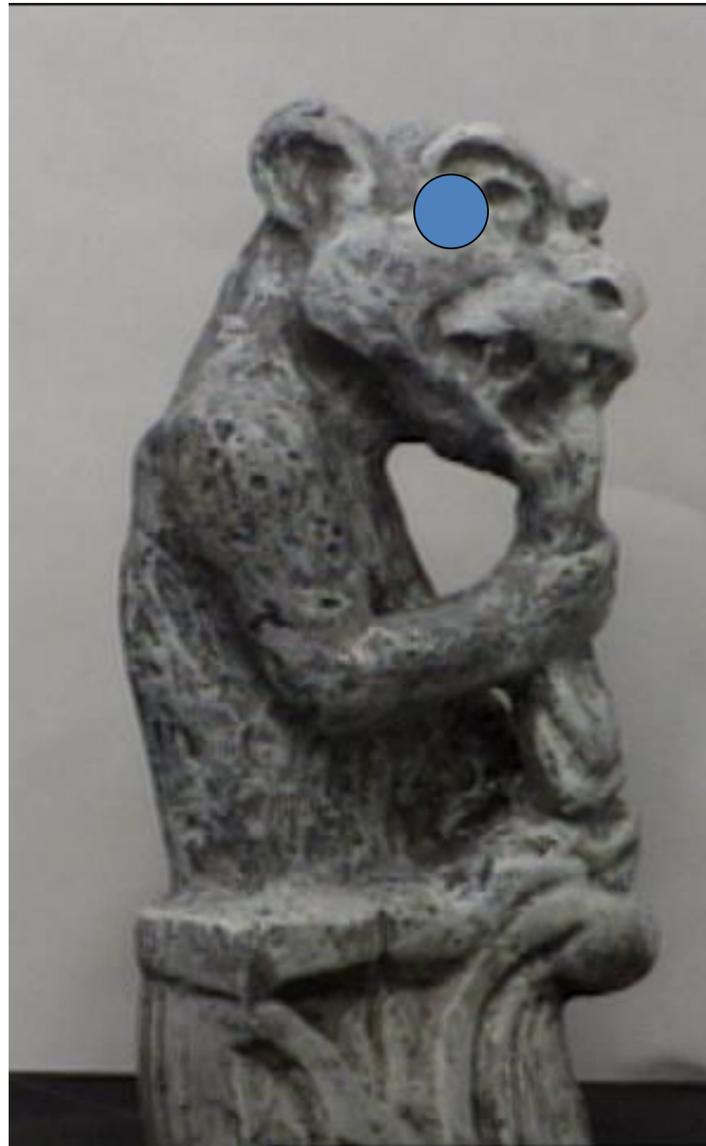


# Warping with a homography (PS8)



# A similar geometric problem: triangulation

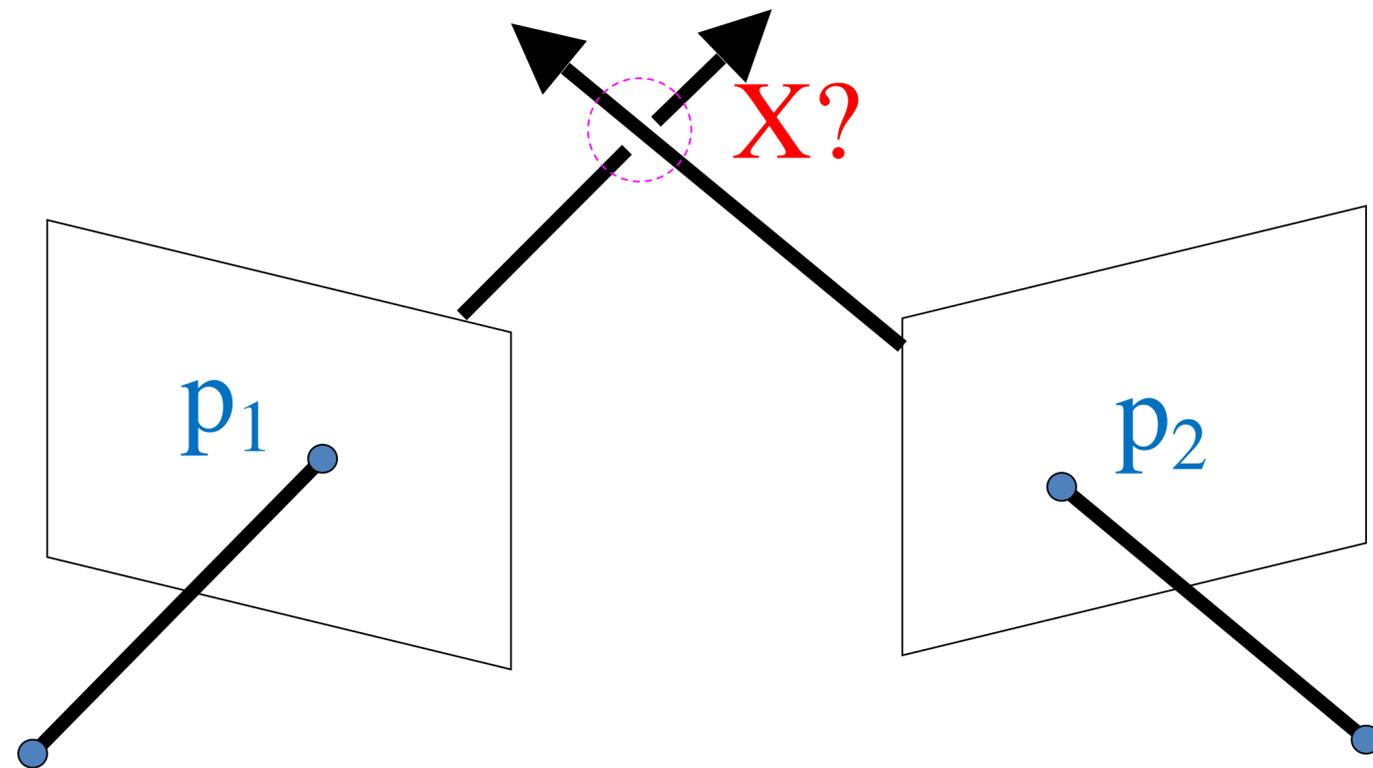
Given projection  $\mathbf{p}_i$  of unknown 3D point  $\mathbf{X}$  in two or more images (with known cameras  $\mathbf{P}_i$ ), find  $\mathbf{X}$



# Triangulation

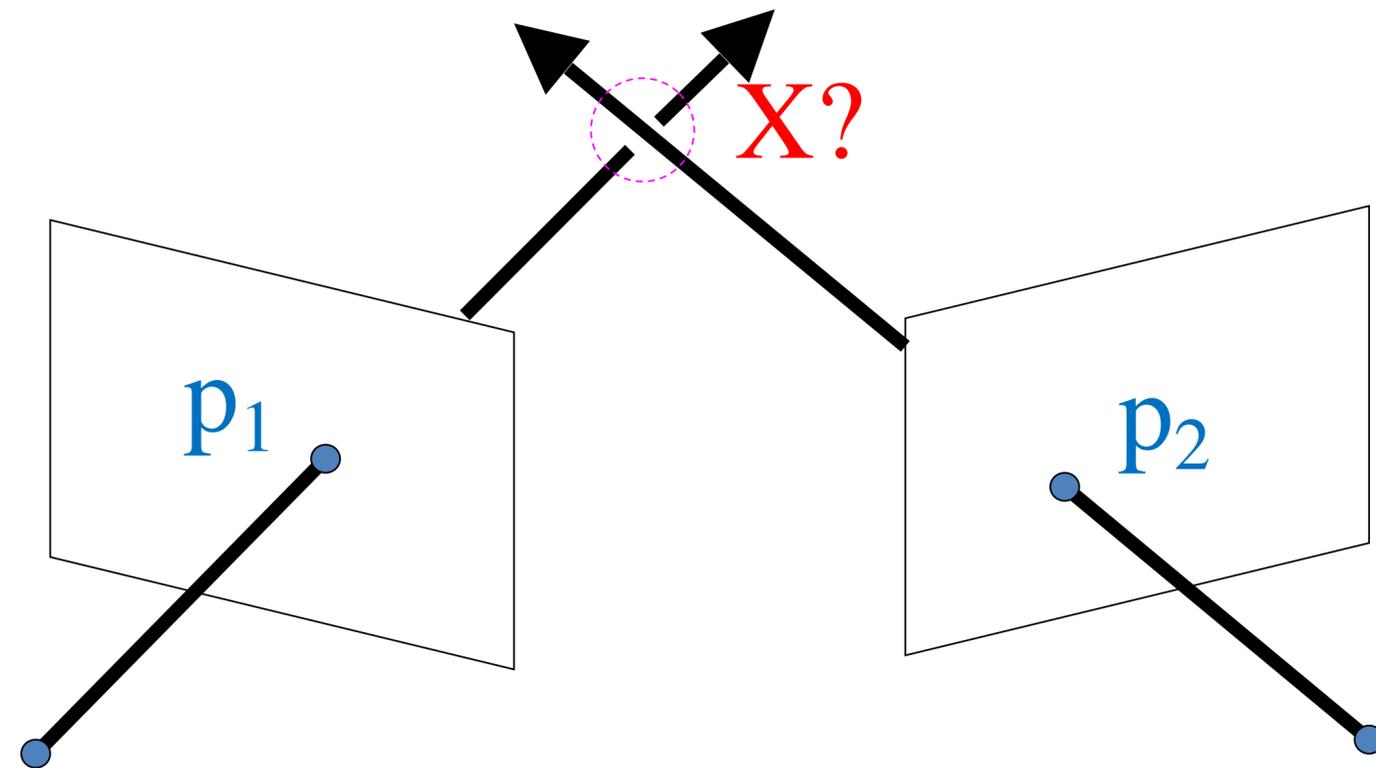
Given projection  $p_i$  of unknown 3D point  $X$  in two or more images (with known cameras  $P_i$ ), find  $X$

**Why is the calibration here important?**



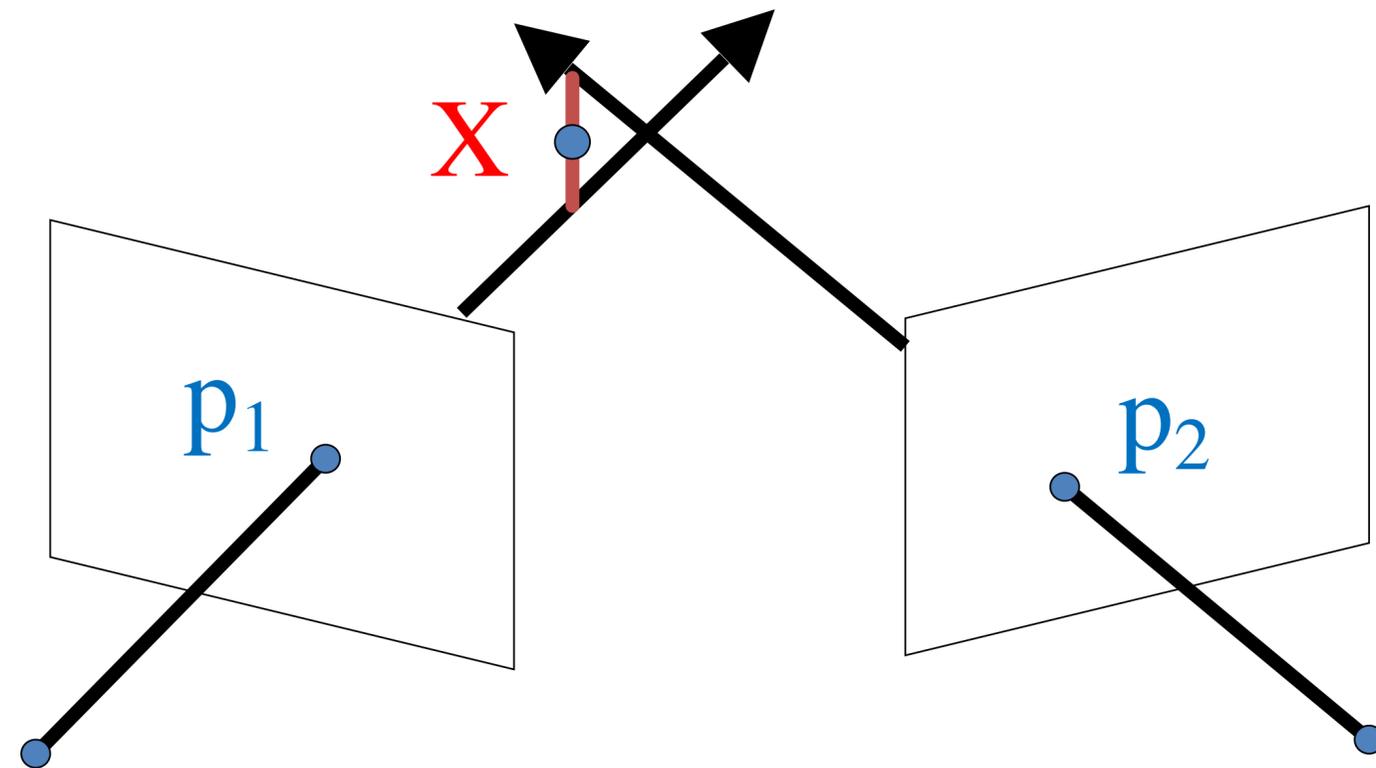
# Triangulation

Rays in principle should intersect, but in practice usually don't exactly due to noise, numerical errors.



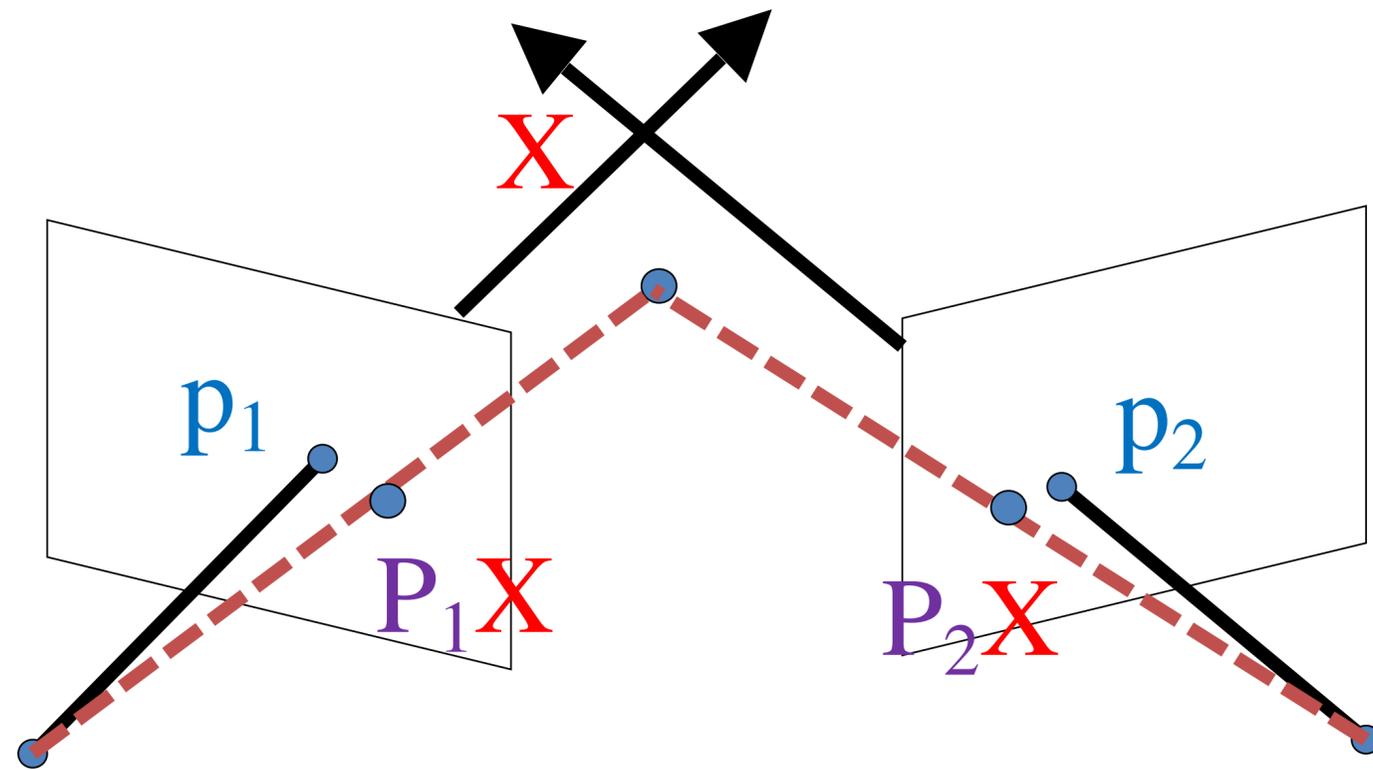
# Triangulation – Geometry

Find shortest segment between viewing rays, set **X** to be the midpoint of the segment.

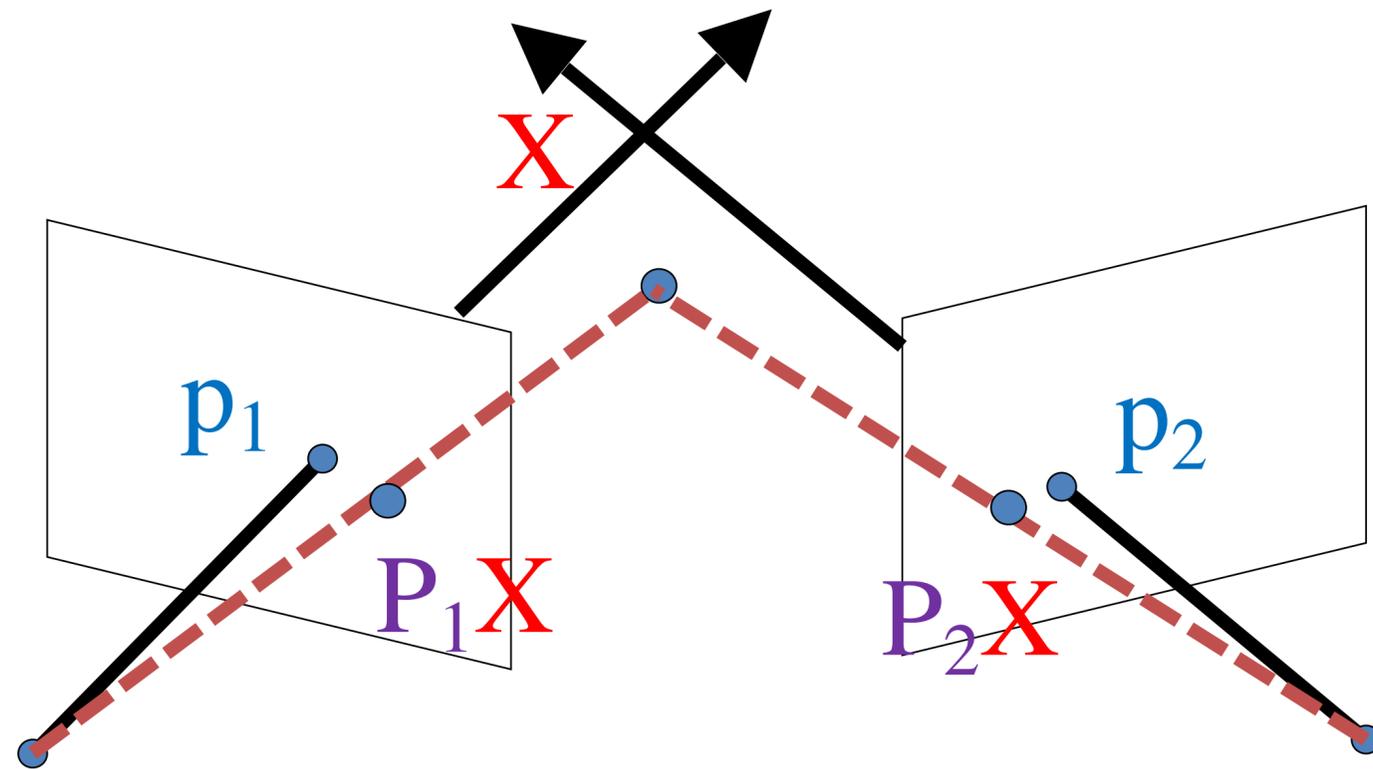


# Triangulation – Non-linear Optim.

Find  $X$  minimizing  $d(p_1, P_1 X)^2 + d(p_2, P_2 X)^2$   
where  $d$  is distance in image space



# Triangulation – Linear Optimization



# First: A better way to handle homogeneous coordinates in linear optimization

Recall: projection in homogeneous coordinates.

$$p_i \equiv PX_i$$

Remember: this implies  $PX_i$  &  $p_i$  are proportional/scaled copies of each other

$$p_i = \lambda PX_i, \lambda \neq 0$$

This implies their cross product is  $\mathbf{0}$ , since  $a \times b = \|a\| \|b\| \sin(\theta)$ .

$$p_i \times PX_i = \mathbf{0}$$

Handles the “divide by 0” issue we saw before.

# Triangulation – Linear Optimization

$$\begin{array}{l}
 p_1 \equiv P_1 X \\
 p_2 \equiv P_2 X
 \end{array}
 \begin{array}{l}
 \rightarrow \\
 \rightarrow
 \end{array}
 \begin{array}{l}
 p_1 \times P_1 X = \mathbf{0} \\
 p_2 \times P_2 X = \mathbf{0}
 \end{array}
 \begin{array}{l}
 \rightarrow \\
 \rightarrow
 \end{array}
 \begin{array}{l}
 [p_{1x}] P_1 X = \mathbf{0} \\
 [p_{2x}] P_2 X = \mathbf{0}
 \end{array}$$

Cross Prod.  
as matrix

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

$$\begin{array}{l}
 [p_{1x}] P_1 X = \mathbf{0} \\
 [p_{2x}] P_2 X = \mathbf{0}
 \end{array}
 \begin{array}{l}
 \rightarrow \\
 \rightarrow
 \end{array}
 \begin{array}{l}
 ([p_{1x}] P_1) X = \mathbf{0} \\
 ([p_{2x}] P_2) X = \mathbf{0}
 \end{array}
 \begin{array}{l}
 \rightarrow \\
 \rightarrow
 \end{array}
 \begin{array}{l}
 \text{Two eqns per} \\
 \text{camera for 3} \\
 \text{unknown in } X
 \end{array}$$

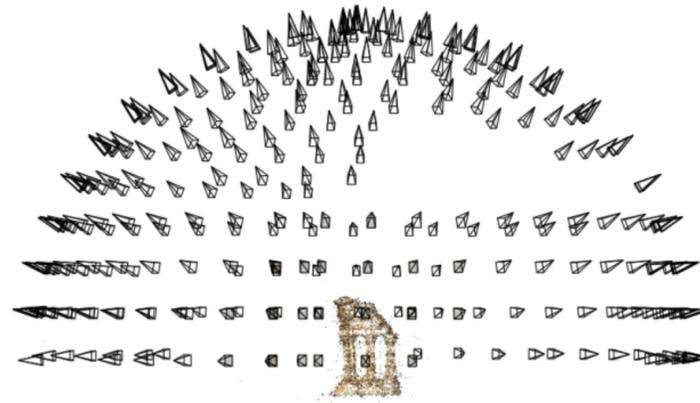
# Next time: Estimating 3D structure

- Given many images, how can we...
  1. Figure out where they were all taken from?
  2. Build a 3D model of the scene?

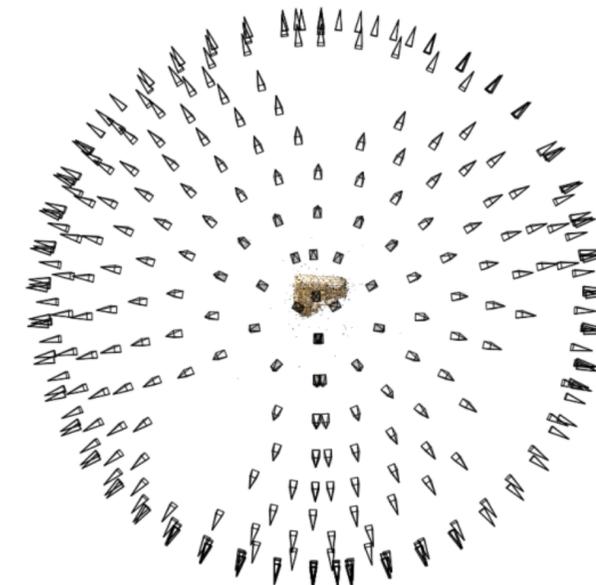


This is the **structure from motion** problem

# Structure from motion



Reconstruction (side)



(top)

- Input: images with pixels in correspondence
- Output
  - **Structure:** 3D location  $\mathbf{x}_i$  for each point  $p_i$
  - **Motion:** camera parameters  $\mathbf{R}_j$ ,  $\mathbf{t}_j$  possibly  $\mathbf{K}_j$
- Objective function: minimize reprojection error

$$p_{i,j} = (u_{i,j}, v_{i,j})$$

# Camera calibration & triangulation

- **Suppose we know 3D points**
  - And have matches between these points and an image
  - Computing camera parameters similar to homography estimation
- **Suppose we have know camera parameters, each of which observes a point**
  - We can solve for the 3D location
- Seems like a chicken-and-egg problem, but in SfM we can solve both at once

Next class: more 3D