



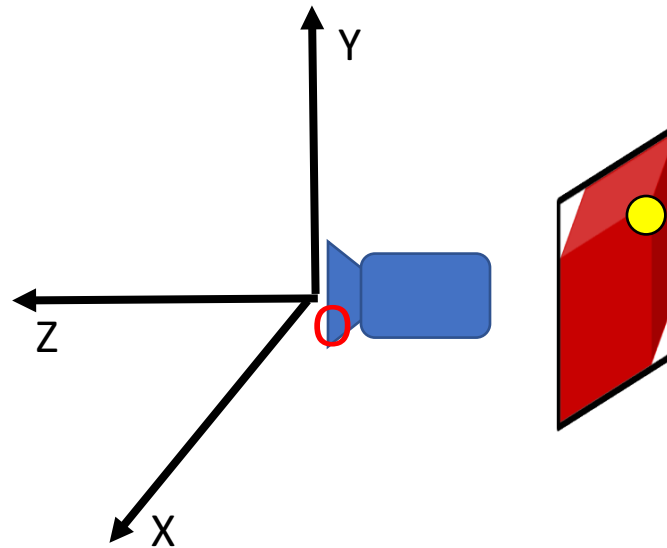

A special case of  
calibration

# Camera calibration



# Camera calibration = pose estimation

- Estimating where camera is relative to object in world
- = Estimating where object is relative to camera





# What if object of interest is plane?

- Not that uncommon....



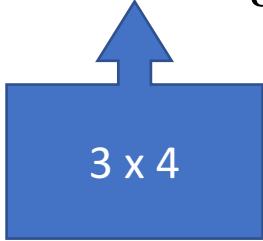
What if object of interest is plane?



# What if object of interest is a plane?

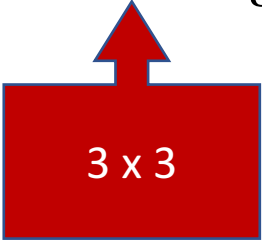
- Imagine that plane is equipped with two axes.
- Points on the plane are represented by *two* euclidean coordinates
- ...Or 3 homogenous coordinates

3D object

$$\vec{\mathbf{x}}_{img} \equiv P \vec{\mathbf{x}}_w$$


3 x 4

2D object (plane)

$$\vec{\mathbf{x}}_{img} \equiv H \vec{\mathbf{x}}_w$$


3 x 3

# What if object of interest is a plane?

Homography

$$\vec{X}_{img} \equiv H \vec{X}_w$$

3 x 3

- Homography maps points on the plane to pixels in the image



# Fitting homographies

- How many parameters does a homography have?
- Given a single point on the plane and corresponding image location, what does that tell us?

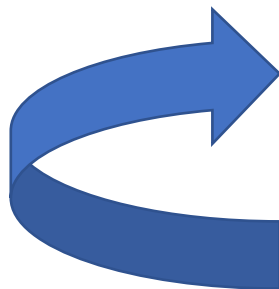
$$\vec{\mathbf{x}}_{img} \equiv H \vec{\mathbf{x}}_w$$

$$\begin{bmatrix} \lambda x \\ \lambda y \\ \lambda \end{bmatrix} = \begin{bmatrix} H_{11} & H_{12} & H_{13} \\ H_{21} & H_{22} & H_{23} \\ H_{31} & H_{32} & H_{33} \end{bmatrix} \begin{bmatrix} x_w \\ y_w \\ 1 \end{bmatrix}$$



# Fitting homographies

- How many parameters does a homography have?
- Given a single point on the plane and corresponding image location, what does that tell us?


$$\begin{bmatrix} \lambda x \\ \lambda y \\ \lambda \end{bmatrix} = \begin{bmatrix} H_{11} & H_{12} & H_{13} \\ H_{21} & H_{22} & H_{23} \\ H_{31} & H_{32} & H_{33} \end{bmatrix} \begin{bmatrix} x_w \\ y_w \\ 1 \end{bmatrix}$$

The parameter  $\lambda$  in the third row of the left matrix is circled in red.

- Convince yourself that this gives 2 linear equations!

# Fitting homographies

- Homography has 9 parameters
- But can't determine scale factor, so only 8: 4 points!

$$A\mathbf{h} = 0 \text{ s.t. } \|\mathbf{h}\| = 1$$

- Or because we will have noise:

$$\min_{\mathbf{h}} \|A\mathbf{h}\|^2 \text{ s.t. } \|\mathbf{h}\| = 1$$

# Fitting homographies



a



b

# Homographies for image alignment

- A general mapping from one plane to another!
- Can also be used to align one photo of a plane to another photo of the same plane



Image 1




Image 2



Original plane

# Homographies for image alignment

- Can also be used to align one photo of a plane to another photo of the same plane



# Image Alignment Algorithm

Given images A and B


1. Compute image features for A and B
2. Match features between A and B
3. Compute homography between A and B

What could go wrong?

# Fitting in general


- Fitting: find the parameters of a model that best fit the data
- Other examples:
  - least squares linear regression

# Least squares: linear regression





# Linear regression



$$\text{Cost}(m, b) = \sum_{i=1}^n |y_i - (mx_i + b)|^2$$

# Linear regression

$$\begin{bmatrix} x_1 & 1 \\ x_2 & 1 \\ \vdots & \\ x_n & 1 \end{bmatrix} \begin{bmatrix} m \\ b \end{bmatrix} = \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{bmatrix}$$

# Outliers

outliers





inliers

# Robustness



Problem: Fit a line to these datapoints




Least squares fit


# Idea

- Given a hypothesized line
- Count the number of points that “agree” with the line
  - “Agree” = within a small distance of the line
  - I.e., the **inliers** to that line
- For all possible lines, select the one with the largest number of inliers

# Counting inliers




# Counting inliers



**Inliers: 3**

# Counting inliers






# How do we find the best line?

- Unlike least-squares, no simple closed-form solution
- Hypothesize-and-test
  - Try out many lines, keep the best one
  - Which lines?

# RANSAC (Random Sample Consensus)

Line fitting example




Algorithm:

1. **Sample** (randomly) the number of points required to fit the model ( $\#=2$ )
2. **Solve** for model parameters using samples
3. **Score** by the fraction of inliers within a preset threshold of the model

**Repeat** 1-3 until the best model is found with high confidence

# RANSAC

Line fitting example




Algorithm:

1. **Sample** (randomly) the number of points required to fit the model ( $n=2$ )
2. **Solve** for model parameters using samples
3. **Score** by the fraction of inliers within a preset threshold of the model

**Repeat** 1-3 until the best model is found with high confidence

# RANSAC

Line fitting example




$$N_I = 6$$

Algorithm:

1. **Sample** (randomly) the number of points required to fit the model (#=2)
2. **Solve** for model parameters using samples
3. **Score** by the fraction of inliers within a preset threshold of the model

**Repeat** 1-3 until the best model is found with high confidence

# RANSAC



Algorithm:

1. **Sample** (randomly) the number of points required to fit the model ( $\#=2$ )
2. **Solve** for model parameters using samples
3. **Score** by the fraction of inliers within a preset threshold of the model


**Repeat** 1-3 until the best model is found with high confidence

# RANSAC


- Idea:
  - All the inliers will agree with each other on the translation vector; 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

# Translations




# Random Sample Consensus

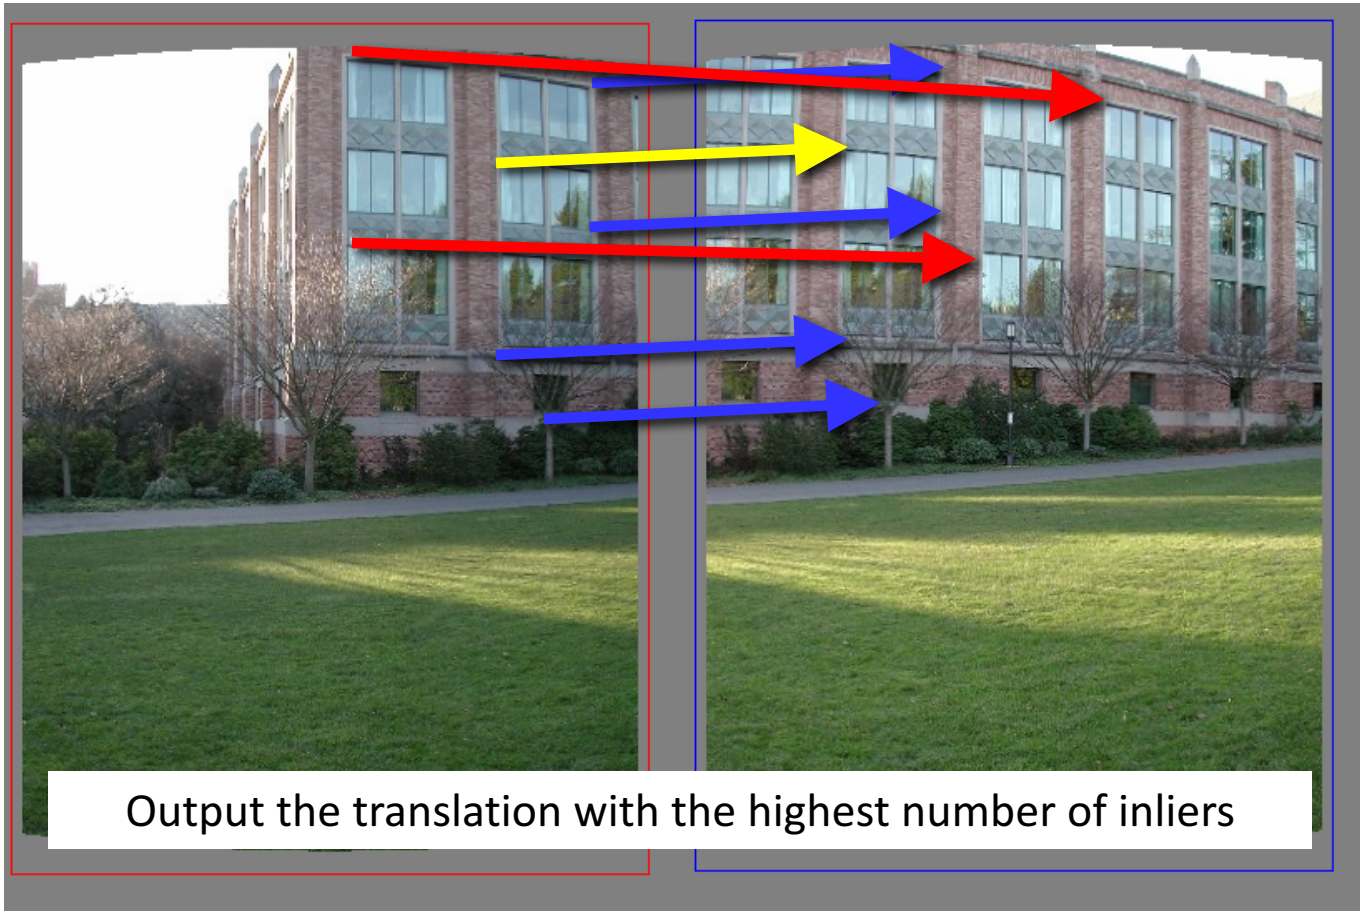




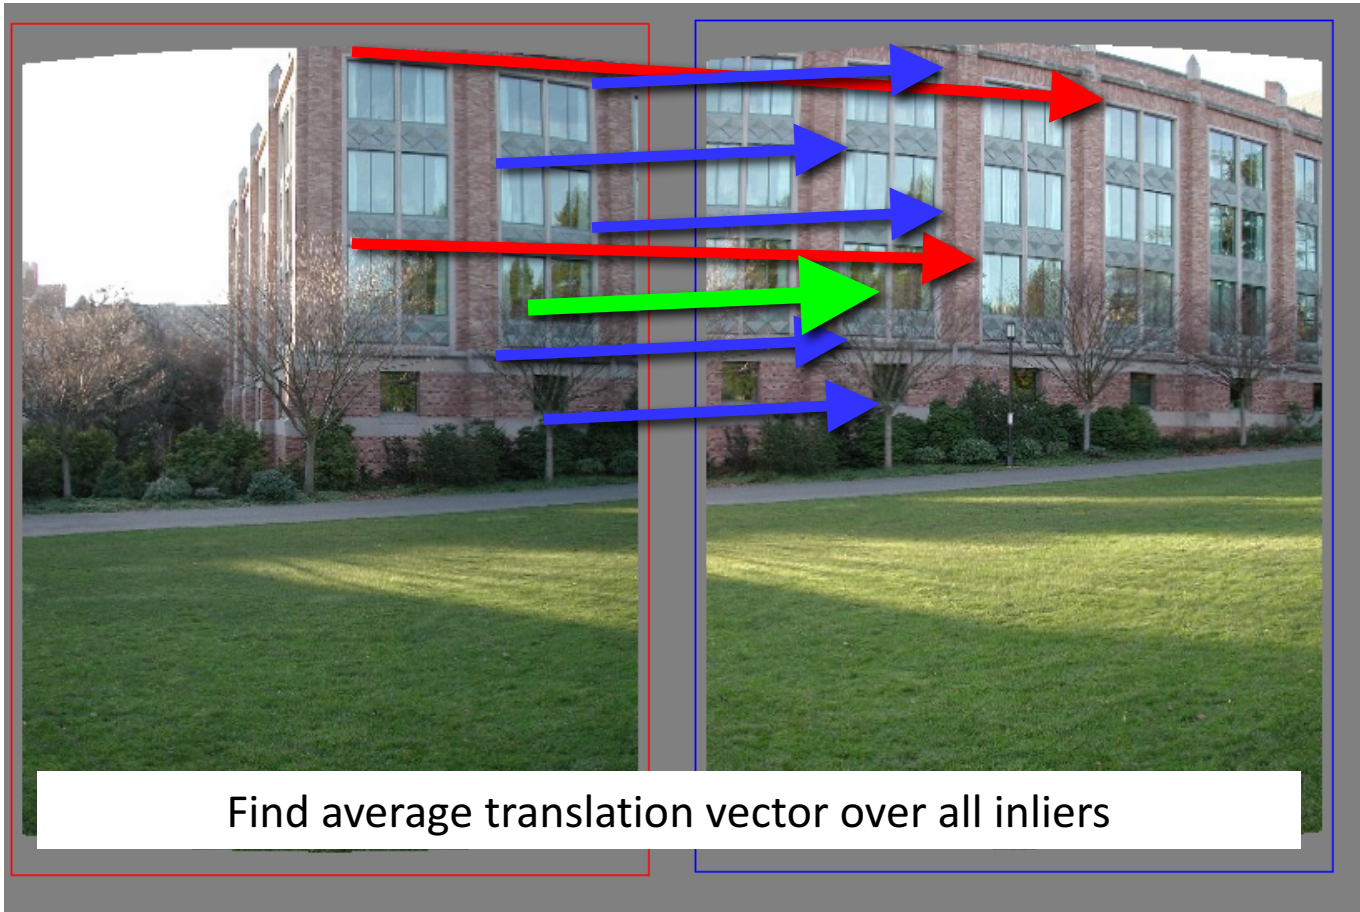
# Random Sample Consensus



# Random Sample Consensus



# Final step: least squares fit



# RANSAC

- **Inlier threshold** related to the amount of noise we expect in inliers
  - Often model noise as Gaussian with some standard deviation (e.g., 3 pixels)
- **Number of rounds** related to the percentage of outliers we expect, and the probability of success we'd like to guarantee
  - Suppose there are 20% outliers, and we want to find the correct answer with 99% probability
  - How many rounds do we need?

# How many rounds?

- If we have to choose  $k$  samples each time
  - with an inlier ratio  $p$
  - and we want the right answer with probability  $P$

proportion of inliers $p$							
$k$	95%	90%	80%	75%	70%	60%	50%
2	2	3	5	6	7	11	17
3	3	4	7	9	11	19	35
4	3	5	9	13	17	34	72
5	4	6	12	17	26	57	146
6	4	7	16	24	37	97	293
7	4	8	20	33	54	163	588
8	5	9	26	44	78	272	1177

$P = 0.99$

To ensure that the random sampling has a good chance of finding a true set of inliers, a sufficient number of trials  $S$  must be tried. Let  $p$  be the probability that any given correspondence is valid and  $P$  be the total probability of success after  $S$  trials. The likelihood in one trial that all  $k$  random samples are inliers is  $p^k$ . Therefore, the likelihood that  $S$  such trials will all fail is

$$1 - P = (1 - p^k)^S \quad (6.29)$$

and the required minimum number of trials is


$$S = \frac{\log(1 - P)}{\log(1 - p^k)}. \quad (6.30)$$






proportion of inliers $p$							
k	95%	90%	80%	75%	70%	60%	50%
2	2	3	5	6	7	11	17
3	3	4	7	9	11	19	35
4	3	5	9	13	17	34	72
5	4	6	12	17	26	57	146
6	4	7	16	24	37	97	293
7	4	8	20	33	54	163	588
8	5	9	26	44	78	272	1177

$$P = 0.99$$

# How big is $k$ ?

- For alignment, depends on the motion model
  - Here, each sample is a correspondence (pair of matching points)



Name	Matrix	# D.O.F.	Preserves:	Icon
translation	$\begin{bmatrix} \mathbf{I} & \mathbf{t} \end{bmatrix}_{2 \times 3}$	2	orientation + ...	
rigid (Euclidean)	$\begin{bmatrix} \mathbf{R} & \mathbf{t} \end{bmatrix}_{2 \times 3}$	3	lengths + ...	
similarity	$\begin{bmatrix} s\mathbf{R} & \mathbf{t} \end{bmatrix}_{2 \times 3}$	4	angles + ...	
affine	$\begin{bmatrix} \mathbf{A} \end{bmatrix}_{2 \times 3}$	6	parallelism + ...	
projective	$\begin{bmatrix} \tilde{\mathbf{H}} \end{bmatrix}_{3 \times 3}$	8	straight lines	

# RANSAC pros and cons

- Pros
  - Simple and general
  - Applicable to many different problems
  - Often works well in practice
- Cons
  - Parameters to tune
  - Sometimes too many iterations are required
  - Can fail for extremely low inlier ratios




# RANSAC

- An example of a “voting”-based fitting scheme
- Each hypothesis gets voted on by each data point, best hypothesis wins
- There are many other types of voting schemes
  - E.g., Hough transforms...

# Hough transform

P.V.C. Hough, *Machine Analysis of Bubble Chamber Pictures*, Proc. Int. Conf. High Energy Accelerators and Instrumentation, 1959

Given a set of points, find the curve or line that explains the data points best








$$y = m x + b$$

# Hough Transform: Outline

1. Create a grid of parameter values
2. Each point votes for a set of parameters, incrementing those values in grid
3. Find maximum or local maxima in grid

# Hough transform





$$d = x \cos \theta + y \sin \theta$$

# Hough transform

