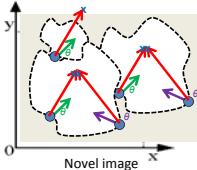


Generalized Hough Transform

Detection procedure:

For each edge point:

- Use its gradient orientation θ to index into stored table
- Use retrieved r vectors to vote for reference point

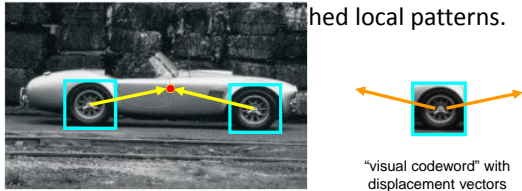


		...
		...
⋮		

Assuming translation is the only transformation here, i.e., orientation and scale are fixed.

Generalized Hough for object detection

- Instead of indexing displacements by gradient orientation, index by learned local patterns.




training image

B. Leibe, A. Leonardis, and B. Schiele, [Combined Object Categorization and Segmentation with an Implicit Shape Model](#), ECCV Workshop on Statistical Learning in Computer Vision 2004

Source: L. Lázebník

Generalized Hough for object detection

- Instead of indexing displacements by gradient orientation, index by "visual codeword"



test image

B. Leibe, A. Leonardis, and B. Schiele, [Combined Object Categorization and Segmentation with an Implicit Shape Model](#), ECCV Workshop on Statistical Learning in Computer Vision 2004

Source: L. Lázebník

Summary

- **Fitting** problems require finding any supporting evidence for a model, even within clutter and missing features
 - associate features with an explicit model
- **Voting** approaches, such as the **Hough transform**, make it possible to find likely model parameters without searching all combinations of features
 - Hough transform approach for lines, circles, ..., arbitrary shapes defined by a set of boundary points, recognition from patches


7

Today

- Feature-based alignment
 - 2D transformations
 - Affine fit
 - RANSAC

8


Motivation: Recognition



9

Figures from David Lowe

Motivation: medical image registration




Slide credit: Kristen Grauman

10

Motivation: mosaics

(In detail next week)




11

Image from http://graphics.cs.cmu.edu/courses/15-463/2010_fall/

Alignment problem

- We have previously considered how to **fit a model to image evidence**
 - e.g., a line to edge points
- In alignment, we will **fit the parameters of some transformation** according to a set of matching feature pairs (“correspondences”).



Slide credit: Kristen Grauman

12

Parametric (global) warping

Examples of parametric warps:

translation rotation aspect

affine perspective

13
Source: Alyosha Efros

Parametric (global) warping

$\mathbf{p} = (x, y)$ $\mathbf{p}' = (x', y')$

Transformation T is a coordinate-changing machine:
 $\mathbf{p}' = T(\mathbf{p})$

What does it mean that T is **global**?

- Is the same for any point \mathbf{p}
- can be described by just a few numbers (parameters)

Let's represent T as a matrix:

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \mathbf{M} \begin{bmatrix} x \\ y \end{bmatrix}$$

14
Source: Alyosha Efros

Scaling

Scaling a coordinate means multiplying each of its components by a scalar

Uniform scaling means this scalar is the same for all components:

$\times 2$

15
Source: Alyosha Efros

Scaling

Non-uniform scaling: different scalars per component:

$X \times 2,$
 $Y \times 0.5$

16
Source: Alyosha Efros

Scaling

Scaling operation: $x' = ax$
 $y' = by$

Or, in matrix form:

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \underbrace{\begin{bmatrix} a & 0 \\ 0 & b \end{bmatrix}}_{\text{scaling matrix } S} \begin{bmatrix} x \\ y \end{bmatrix}$$

17
Source: Alyosha Efros

What transformations can be represented with a 2x2 matrix?

2D Scaling?
 $x' = s_x * x$
 $y' = s_y * y$

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} s_x & 0 \\ 0 & s_y \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}$$

2D Rotate around (0,0)?
 $x' = \cos \Theta * x - \sin \Theta * y$
 $y' = \sin \Theta * x + \cos \Theta * y$

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} \cos \Theta & -\sin \Theta \\ \sin \Theta & \cos \Theta \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}$$

2D Shear?
 $x' = x + sh_x * y$
 $y' = sh_y * x + y$

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} 1 & sh_x \\ sh_y & 1 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}$$

18
Source: Alyosha Efros

What transformations can be represented with a 2x2 matrix?

2D Mirror about Y axis?

$$\begin{matrix} x' = -x \\ y' = y \end{matrix} \quad \begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} -1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}$$

2D Mirror over (0,0)?

$$\begin{matrix} x' = -x \\ y' = -y \end{matrix} \quad \begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} -1 & 0 \\ 0 & -1 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}$$

2D Translation?

$$\begin{matrix} x' = x + t_x \\ y' = y + t_y \end{matrix} \quad \text{NO!}$$

19
Source: Alyosha Efros

2D Linear Transformations

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix}$$

Only linear 2D transformations can be represented with a 2x2 matrix.

Linear transformations are combinations of ...

- Scale,
- Rotation,
- Shear, and
- Mirror

20
Source: Alyosha Efros

Homogeneous coordinates

Convenient coordinate system to represent many useful transformations

To convert to homogeneous coordinates:

$$(x, y) \Rightarrow \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

homogeneous image coordinates

Converting *from* homogeneous coordinates:

$$\begin{bmatrix} x \\ y \\ w \end{bmatrix} \Rightarrow (x/w, y/w)$$

21

Slide credit: Kristen Grauman

Homogeneous Coordinates

Q: How can we represent 2d translation as a 3x3 matrix using homogeneous coordinates?

$$x' = x + t_x$$

$$y' = y + t_y$$

A: Using the rightmost column:


$$\text{Translation} = \begin{bmatrix} 1 & 0 & t_x \\ 0 & 1 & t_y \\ 0 & 0 & 1 \end{bmatrix}$$

22
Source: Alyosha Efros

Translation

Homogeneous Coordinates

$$\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & t_x \\ 0 & 1 & t_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = \begin{bmatrix} x + t_x \\ y + t_y \\ 1 \end{bmatrix}$$



23
Source: Alyosha Efros

Basic 2D Transformations

Basic 2D transformations as 3x3 matrices

$$\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & t_x \\ 0 & 1 & t_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

Translate

$$\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \begin{bmatrix} s_x & 0 & 0 \\ 0 & s_y & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

Scale

$$\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \begin{bmatrix} \cos\Theta & -\sin\Theta & 0 \\ \sin\Theta & \cos\Theta & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

Rotate

$$\begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \begin{bmatrix} 1 & sh_x & 0 \\ sh_y & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

Shear

24
Source: Alyosha Efros


2D Affine Transformations

$$\begin{bmatrix} x' \\ y' \\ w' \end{bmatrix} = \begin{bmatrix} a & b & c \\ d & e & f \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ w \end{bmatrix}$$

Affine transformations are combinations of ...

- Linear transformations, and
- Translations

Parallel lines remain parallel



Slide credit: Kristen Grauman

25


Today

- Feature-based alignment
 - 2D transformations
 - Affine fit
 - RANSAC

26

Alignment problem

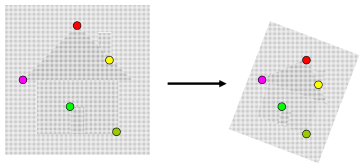
- We have previously considered how to **fit a model to image evidence**
 - e.g., a line to edge points
- In alignment, we will fit the parameters of some **transformation** according to a set of matching feature pairs ("correspondences").



27

Kristen Grauman

Image alignment

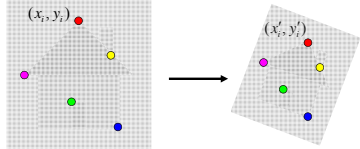


- Two broad approaches:
 - Direct (pixel-based) alignment
 - Search for alignment where most pixels agree
 - Feature-based alignment
 - Search for alignment where *extracted features* agree
 - Can be verified using pixel-based alignment

Slide credit: Kristen Grauman 28

Fitting an affine transformation

- Assuming we know the correspondences, how do we get the transformation?



$$\begin{bmatrix} x'_i \\ y'_i \end{bmatrix} = \begin{bmatrix} m_1 & m_2 \\ m_3 & m_4 \end{bmatrix} \begin{bmatrix} x_i \\ y_i \end{bmatrix} + \begin{bmatrix} t_1 \\ t_2 \end{bmatrix}$$

Slide credit: Kristen Grauman 29

An aside: Least Squares Example

Say we have a set of data points (X_1, X'_1) , (X_2, X'_2) , (X_3, X'_3) , etc. (e.g. person's height vs. weight)

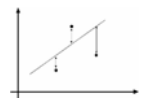
We want a nice compact formula (a line) to predict X' s from X s: $Xa + b = X'$

We want to find a and b

How many (X, X') pairs do we need?

$$\begin{matrix} X_1 a + b = X'_1 \\ X_2 a + b = X'_2 \end{matrix} \quad \begin{bmatrix} X_1 & 1 \\ X_2 & 1 \end{bmatrix} \begin{bmatrix} a \\ b \end{bmatrix} = \begin{bmatrix} X'_1 \\ X'_2 \end{bmatrix} \quad Ax=B$$


What if the data is noisy?

$$\begin{bmatrix} X_1 & 1 \\ X_2 & 1 \\ X_3 & 1 \\ \dots & \dots \end{bmatrix} \begin{bmatrix} a \\ b \end{bmatrix} = \begin{bmatrix} X'_1 \\ X'_2 \\ X'_3 \\ \dots \end{bmatrix} \quad \min \|Ax - B\|^2$$


overconstrained Source: Ayosha Efros

Fitting an affine transformation

- Assuming we know the correspondences, how do we get the transformation?



$$\begin{bmatrix} x'_i \\ y'_i \end{bmatrix} = \begin{bmatrix} m_1 & m_2 \\ m_3 & m_4 \end{bmatrix} \begin{bmatrix} x_i \\ y_i \end{bmatrix} + \begin{bmatrix} t_1 \\ t_2 \end{bmatrix}$$

Slide credit: Kristen Grauman


Fitting an affine transformation

$$\begin{bmatrix} x_i & y_i & 0 & 0 & 1 & 0 \\ 0 & 0 & x_i & y_i & 0 & 1 \\ \dots & \dots & \dots & \dots & \dots & \dots \end{bmatrix} \begin{bmatrix} m_1 \\ m_2 \\ m_3 \\ m_4 \\ t_1 \\ t_2 \end{bmatrix} = \begin{bmatrix} \dots \\ x'_i \\ y'_i \\ \dots \end{bmatrix}$$

- How many matches (correspondence pairs) do we need to solve for the transformation parameters?

32
Kristen Grauman


Affine: # correspondences?



How many correspondences needed for affine?

Alyosha Efros


Fitting an affine transformation



Example from UBC SIFT Demo

37

Fitting an affine transformation



Example from UBC SIFT Demo

38

Fitting an affine transformation



39
Figures from David Lowe, ICCV 1999


Today

- Feature-based alignment
 - 2D transformations
 - Affine fit
 - RANSAC

40

Outliers


- **Outliers** can hurt the quality of our parameter estimates, e.g.,
 - an erroneous pair of matching points from two images
 - an edge point that is noise, or doesn't belong to the line we are fitting.



41

Kristen Grauman


Outliers affect least squares fit



Slide credit: Kristen Grauman

42

Outliers affect least squares fit



Slide credit: Kristen Grauman

43

RANSAC

- RANdom Sample Consensus
- **Approach:** we want to avoid the impact of outliers, so let's look for "inliers", and use those only.
- **Intuition:** if an outlier is chosen to compute the current fit, then the resulting line won't have much support from rest of the points.

Slide credit: Kristen Grauman

44


RANSAC: General form

- RANSAC loop:
 1. Randomly select a *seed group* of points on which to base transformation estimate
 2. Compute transformation from seed group
 3. Find *inliers* to this transformation
 4. If the number of inliers is sufficiently large, re-compute estimate of transformation on all of the inliers
- Keep the transformation with the largest number of inliers

Slide credit: Kristen Grauman

45


RANSAC for line fitting example



Source: R. Raguram

46
Lana Lazebnik

RANSAC for line fitting example




Least-squares fit

Source: R. Raguram

47
Lana Lazebnik

RANSAC for line fitting example

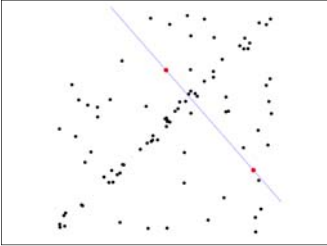


1. Randomly select minimal subset of points

Source: R. Raguram

48
Lana Lazebnik

RANSAC for line fitting example

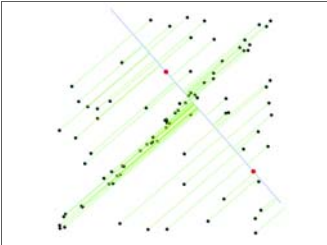


1. Randomly select minimal subset of points
2. Hypothesize a model

Source: R. Raguram

49
Lana Lazebnik

RANSAC for line fitting example

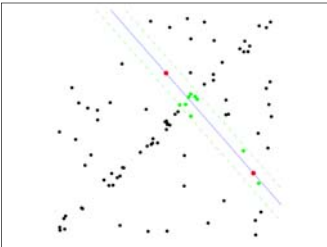


1. Randomly select minimal subset of points
2. Hypothesize a model
3. Compute error function

Source: R. Raguram

50
Lana Lazebnik

RANSAC for line fitting example

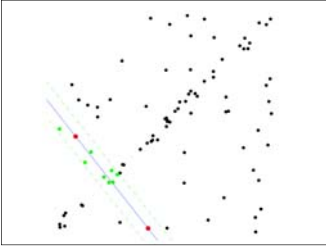


1. Randomly select minimal subset of points
2. Hypothesize a model
3. Compute error function
4. Select points consistent with model

Source: R. Raguram

51
Lana Lazebnik

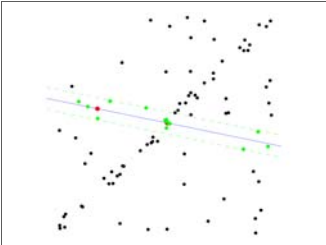
RANSAC for line fitting example



1. Randomly select minimal subset of points
2. Hypothesize a model
3. Compute error function
4. Select points consistent with model
5. Repeat *hypothesize-and-verify loop*

Source: R. Raguram
Lana Lazebnik

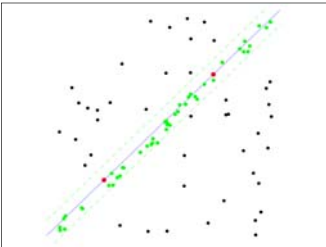
RANSAC for line fitting example



1. Randomly select minimal subset of points
2. Hypothesize a model
3. Compute error function
4. Select points consistent with model
5. Repeat *hypothesize-and-verify loop*

Source: R. Raguram
Lana Lazebnik


RANSAC for line fitting example



1. Randomly select minimal subset of points
2. Hypothesize a model
3. Compute error function
4. Select points consistent with model
5. Repeat *hypothesize-and-verify loop*

Source: R. Raguram
Lana Lazebnik

RANSAC for line fitting example



1. Randomly select minimal subset of points
2. Hypothesize a model
3. Compute error function
4. Select points consistent with model
5. Repeat hypothesize-and-verify loop

Source: R. Raguram

Lana Lazebnik

RANSAC for line fitting

Repeat N times:

- Draw s points uniformly at random
- Fit line to these s points
- Find inliers to this line among the remaining points (i.e., points whose distance from the line is less than t)
- If there are d or more inliers, accept the line and refit using all inliers

56

Lana Lazebnik

RANSAC pros and cons

- Pros
 - Simple and general
 - Applicable to many different problems
 - Often works well in practice
- Cons
 - Lots of parameters to tune
 - Doesn't work well for low inlier ratios (too many iterations, or can fail completely)

57


Lana Lazebnik

Today

- Feature-based alignment
 - 2D transformations
 - Affine fit
 - RANSAC

58

Coming up: alignment and image stitching



59

Questions?

See you Thursday!

60
