




## Fitting a transformation: feature-based alignment

Tues Oct 13




## Motivation: Recognition



Figures from David Lowe

## Motivation: medical image registration



## Motivation: mosaics




Image from [http://graphics.cs.cmu.edu/courses/15-463/2010\\_fa](http://graphics.cs.cmu.edu/courses/15-463/2010_fa)

## Last week

- Interest point detection
  - Harris corner detector
  - Laplacian of Gaussian, automatic scale selection
- Invariant descriptors
  - Rotation according to dominant gradient direction
  - Histograms for robustness to small shifts and translations (SIFT descriptor)

## Review questions

- What is the purpose of the “ratio test” for local feature matching?
- What aspects of the SIFT descriptor design promote robustness to lighting changes?  
Robustness to rotation and translation?
- Does extracting multiple keypoints for multiple local maxima in scale space help recall or precision during feature matching?
- How far in the image plane can an object rotate before the SIFT descriptors will not match?

## Multi-view: what's next


---

Additional questions we need to address to achieve these applications:

- Fitting a parametric transformation given putative matches
- Dealing with outlier correspondences
- Exploiting geometry to restrict locations of possible matches
- Triangulation, reconstruction
- Efficiency when indexing so many keypoints


## Coming up: robust feature-based alignment

---



Source: L. Lazebnik


## Coming up: robust feature-based alignment



- Extract features

Source: L. Lazebnik


## Coming up: robust feature-based alignment



- Extract features
- Compute *putative matches*

Source: L. Lazebnik


## Coming up: robust feature-based alignment



- Extract features
- Compute *putative matches*
- Loop:
  - *Hypothesize* transformation  $T$  (small group of putative matches that are related by  $T$ )

Source: L. Lazebnik

## Coming up: robust feature-based alignment



- Extract features
- Compute *putative matches*
- Loop:
  - *Hypothesize* transformation  $T$  (small group of putative matches that are related by  $T$ )
  - *Verify* transformation (search for other matches consistent with  $T$ )

Source: L. Lazebnik

## Coming up: robust feature-based alignment



- Extract features
- Compute *putative matches*
- Loop:
  - *Hypothesize* transformation  $T$  (small group of putative matches that are related by  $T$ )
  - *Verify* transformation (search for other matches consistent with  $T$ )


Source: L. Lazebnik

## Today

- Feature-based alignment
  - 2D transformations
  - Affine fit
  - RANSAC for robust fitting

## Alignment as fitting


- Previous lectures: fitting a model to features in one image




Find model  $M$  that minimizes

$$\sum_i \text{residual}(x_i, M)$$

- Alignment: fitting a model to a transformation between pairs of features (*matches*) in two images



$T$



Find transformation  $T$  that minimizes

$$\sum_i \text{residual}(T(x_i), x'_i)$$

Slide credit: Lana Lazebnik

## Parametric (global) warping

Examples of parametric warps:



translation



rotation



aspect



affine



perspective


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:

$$\mathbf{p}' = \mathbf{M}\mathbf{p}$$


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

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:



Source: Alyosha Efros

## Scaling

*Non-uniform scaling*: different scalars per component:



Source: Alyosha Efros

## Scaling

Scaling operation:

$$x' = ax$$

$$y' = by$$

Or, in matrix form:

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

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?

$$x' = x + sh_x * y$$

$$y' = sh_y * x + y$$

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

Source: Alyosha Efros

What transformations can be represented with a 2x2 matrix?

2D Mirror about Y axis?

$$x' = -x$$

$$y' = y$$

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

2D Mirror over (0,0)?

$$x' = -x$$

$$y' = -y$$

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

2D Translation?

$$x' = x + t_x$$

$$y' = y + t_y$$

**NO!**

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

Source: Alyosha Efros

## Homogeneous coordinates

To convert to homogeneous coordinates:

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

homogeneous image  
coordinates

Converting *from* homogeneous coordinates

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

## 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:


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

Source: Alyosha Efros

## Translation

Homogeneous Coordinates

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



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

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




# Today


- Feature-based alignment
  - 2D transformations
  - Affine fit
  - RANSAC for robust fitting

## Alignment problem

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




# 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

## Let's start with affine transformations

- Simple fitting procedure (linear least squares)
- Approximates viewpoint changes for roughly planar objects and roughly orthographic cameras
- Can be used to initialize fitting for more complex models




Lana Lazebnik



## Fitting an affine transformation

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



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

## 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?


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

What if the data is noisy?

$$\begin{bmatrix} X_1 & 1 \\ X_2 & 1 \\ X_3 & 1 \\ \dots & \dots \end{bmatrix} \begin{bmatrix} a \\ b \end{bmatrix} = \begin{bmatrix} X'_1 \\ X'_2 \\ X'_3 \\ \dots \end{bmatrix}$$

overconstrained


$$\min \|Ax - B\|^2$$



Source: Alyosha Efros

## Fitting an affine transformation

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



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



$$\begin{bmatrix} m_1 \\ m_2 \\ m_3 \\ m_4 \\ t_1 \\ t_2 \end{bmatrix} = \begin{bmatrix} \vdots \\ \vdots \end{bmatrix}$$

## Fitting an affine transformation

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

- How many matches (correspondence pairs) do we need to solve for the transformation parameters?
- Once we have solved for the parameters, how do we compute the coordinates of the corresponding point for  $(x_{new}, y_{new})$ ?
- Where do the matches come from?

## Recall: Scale Invariant Feature Transform (SIFT) descriptor [Lowe 2004]




Interest points and their scales and orientations (random subset of 50)

SIFT descriptors

<http://www.vlfeat.org/overview/sift.html>

## Recall: SIFT (preliminary) matches



<http://www.vlfeat.org/overview/sift.html>

## Fitting an affine transformation



Affine model approximates perspective projection of planar objects.


Figures from David Lowe, ICCV 1999

## Today


- Feature-based alignment
  - 2D transformations
  - Affine fit
  - RANSAC for robust fitting

# 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.



## Outliers affect least squares fit



## Outliers affect least squares fit



## 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** (**transformation**) won't have much support from rest of the **points** (**matches**).

## 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

Lana Lazebnik

## RANSAC for line fitting example

---




Source: R. Raguram

Lana Lazebnik

## RANSAC for line fitting example

---




Least-squares fit

Source: R. Raguram

Lana Lazebnik

## RANSAC for line fitting example

---



1. Randomly select minimal subset of points


Source: R. Raguram

Lana Lazebnik



## RANSAC for line fitting example

---




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

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

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**

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


56

Lana Lazebnik

Source: R. Raguram

## RANSAC for line fitting example

**Uncontaminated sample**



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


57

Lana Lazebnik

Source: R. Raguram

## 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

---


That is an example fitting a model  
(line)...

What about fitting a transformation  
(translation)?

## RANSAC: General form


- RANSAC loop:
  1. Randomly select a *seed group* of points on which to base transformation estimate (e.g., a group of matches)
  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

### RANSAC example: Translation




Source: Rick Szeliski

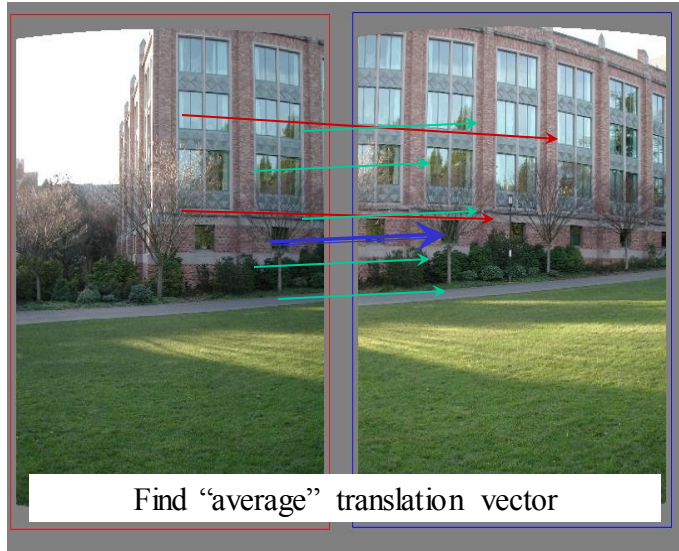
## RANSAC example: Translation



## RANSAC example: Translation




## RANSAC example: Translation



## 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)
  - Can't always get a good initialization of the model based on the minimum number of samples



Lana Lazebnik