


Fitting



Book: Szeliski A.2, 4.3.2, 6.1.4, Forsyth 22.1, 10.1-10.4


Fitting

- We've learned how to detect edges, corners, blobs.
- We would like to form a higher-level, representation of the features in the image by grouping together multiple features.



Fitting

- Choose a parametric model to represent a set of features



simple model: lines



simple model: circles



complicated model: car

Fitting: Issues

Case study: Line detection




- **Noise** in the measured feature locations
- **Extraneous data:** clutter (outliers), multiple lines
- **Missing data:** occlusions

Fitting: Overview

- If we know which points belong to the line, how do we find the “optimal” line parameters?
 - Least squares
- What if there are outliers?
 - Robust fitting, RANSAC
- What if there are many lines?
 - Voting methods: RANSAC, Hough transform
- What if we're not even sure it's a line?
 - Model selection

Total least squares



Total least squares


Point: (x_i, y_i)

Line: $ax+by=d$ ($a^2+b^2=1$)

Distance: $|ax_i + by_i - d|$

Find (a, b, d) to minimize:

$$E = \sum_{i=1}^n (ax_i + by_i - d)^2$$



$$\frac{\partial E}{\partial d} = \sum_{i=1}^n -2(ax_i + by_i - d) = 0$$


$$d = \frac{a}{n} \sum_{i=1}^n x_i + \frac{b}{n} \sum_{i=1}^n y_i = a\bar{x} + b\bar{y}$$

$$E = \sum_{i=1}^n (a(x_i - \bar{x}) + b(y_i - \bar{y}))^2$$

$$\frac{dE}{dN} = 0$$


Least squares: Robustness to noise

Least squares fit to the red points:



Least squares: Robustness to noise

Least squares fit with an outlier:




Problem: squared error heavily penalizes outliers

RANSAC


- Robust fitting can deal with a few outliers – what if we have very many?
- Random sample consensus (RANSAC):
Very general framework for model fitting in the presence of outliers
- Outline
 - Choose a small subset of points uniformly at random
 - Fit a model to that subset
 - Find all remaining points that are “close” to the model and reject the rest as outliers
 - Do this many times and choose the best model

M. A. Fischler, R. C. Bolles. [Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography](#). Comm. of the ACM, Vol 24, pp 381-395, 1981.

RANSAC for line fitting example




RANSAC for line fitting example




Least-squares fit

RANSAC for line fitting example




1. Randomly select minimal subset of points

RANSAC for line fitting example




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

RANSAC for line fitting example




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

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**

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


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


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

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

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

Choosing the parameters

- Initial number of points **s**
 - Typically minimum number needed to fit the model
- Distance threshold **t**
 - Choose **t** so probability for inlier is p (e.g. 0.95)
 - Zero-mean Gaussian noise with std. dev. σ : $t^2 = 3.84\sigma^2$
- Number of iterations **N**
 - Choose **N** so that, with probability p , at least one random sample is free from outliers (e.g. $p=0.99$) (outlier ratio: e)
- Consensus set size **d**
 - Should match expected inlier ratio

Fitting: The Hough transform



Voting schemes

- Let each feature vote for all the models that are compatible with it
- Hopefully the noise features will not vote consistently for any single model
- Missing data doesn't matter as long as there are enough features remaining to agree on a good model

Hough transform

- An early type of voting scheme
- General outline:
 - Discretize *parameter space* into bins
 - For each feature point in the image, put a vote in every bin in the parameter space that could have generated this point
 - Find bins that have the most votes






Image space




Hough parameter space

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


Parameter space representation

- A line in the image corresponds to a point in Hough space




Parameter space representation

- What does a point (x_0, y_0) in the image space map to in the Hough space?




Parameter space representation

- What does a point (x_0, y_0) in the image space map to in the Hough space?
 - Answer: the solutions of $b = -x_0m + y_0$
 - This is a line in Hough space



Parameter space representation


- Where is the line that contains both (x_0, y_0) and (x_1, y_1) ?




Parameter space representation

- Where is the line that contains both (x_0, y_0) and (x_1, y_1) ?
 - It is the intersection of the lines $b = -x_0m + y_0$ and $b = -x_1m + y_1$

Image space



Hough parameter space




Parameter space representation

- Problems with the (m,b) space:
 - Unbounded parameter domains
 - Vertical lines require infinite m

Parameter space representation

- Problems with the (m,b) space:
 - Unbounded parameter domains
 - Vertical lines require infinite m
- Alternative: *polar representation*



Each point (x,y) will add a sinusoid in the (θ, ρ) parameter space

Algorithm outline

- Initialize accumulator H to all zeros
- For each feature point (x,y) in the image

For $\theta = 0$ to 180

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


$$H(\theta, \rho) = H(\theta, \rho) + 1$$

end


end

- Find the value(s) of (θ, ρ) where $H(\theta, \rho)$ is a local maximum
- The detected line in the image is given by
$$\rho = x \cos \theta + y \sin \theta$$


H: accumulator array (votes)



Basic illustration



features




votes

Hough transform demo:


<http://liquify.eu/swf/HoughTransform.swf>

Other shapes



Square




Circle




Several lines




A more complicated image




Effect of noise



Effect of noise




features




votes

Peak gets fuzzy and hard to locate

Random points



features



votes


Uniform noise can lead to spurious peaks in the array

Dealing with noise

- Choose a good grid / discretization
 - **Too coarse:** large votes obtained when too many different lines correspond to a single bucket
 - **Too fine:** miss lines because some points that are not exactly collinear cast votes for different buckets
- Increment neighboring bins (smoothing in accumulator array)
- Try to get rid of irrelevant features
 - E.g., take only edge points with significant gradient magnitude

Incorporating image gradients

- Recall: when we detect an edge point, we also know its gradient direction
- But this means that the line is uniquely determined!
- Modified Hough transform:



$$\theta = \tan^{-1} \left(\frac{\partial f}{\partial y} / \frac{\partial f}{\partial x} \right)$$

For each edge point (x,y)

θ = gradient orientation at (x,y)

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


$H(\theta, \rho) = H(\theta, \rho) + 1$

end

Generalized Hough transform

- We want to find a template defined by its reference point (center) and several distinct types of landmark points in stable spatial configuration


Template




Generalized Hough transform

- Template representation:
for each type of landmark
point, store all possible
displacement vectors
towards the center

Template




Model




Generalized Hough transform

- Detecting the template:
 - For each feature in a new image, look up that feature type in the model and vote for the possible center locations associated with that type in the model

Test image




Model




Application in recognition

- Index displacements by “visual codeword”



training image



visual codeword with
displacement vectors

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

Application in recognition


- Index displacements 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

Image alignment



Book: Forsyth 12.1, Kriegman 2007 paper.

Image alignment: Challenges



Small degree of overlap
Intensity changes




Occlusion,
clutter

Feature-based alignment: Overview

- Alignment as fitting
 - Affine transformations
 - Homographies
- Robust alignment
 - Descriptor-based feature matching
 - RANSAC
- Application: searching the night sky

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




Find transformation T that minimizes


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

2D transformation models


- Similarity
(translation, scale, rotation)



- Affine




- Projective
(homography)




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 objects



Fitting an affine transformation

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



$$\mathbf{x}'_i = \mathbf{M}\mathbf{x}_i + \mathbf{t}$$


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

Want to find $\mathbf{M}$, $\mathbf{t}$ to minimize

$$\sum_{i=1}^n \| \mathbf{x}'_i - \mathbf{M}\mathbf{x}_i - \mathbf{t} \|^2$$

Fitting an affine transformation

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



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

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


Fitting an affine transformation

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

- Linear system with six unknowns
- Each match gives us two linearly independent equations: need at least three to solve for the transformation parameters

Fitting a plane projective transformation

- **Homography:** plane projective transformation (transformation taking a quad to another arbitrary quad)



Homography

- The transformation between two views of a planar surface



The transformation between images from two cameras that share the same center



Fitting a homography

- Recall: homogeneous coordinates

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

Converting *to* homogeneous
image coordinates

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

Converting *from* homogeneous
image coordinates

Equation for homography:

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

Fitting a homography

- Equation for homography:

$$\lambda \begin{bmatrix} x'_i \\ y'_i \\ 1 \end{bmatrix} = \begin{bmatrix} h_{11} & h_{12} & h_{13} \\ h_{21} & h_{22} & h_{23} \\ h_{31} & h_{32} & h_{33} \end{bmatrix} \begin{bmatrix} x_i \\ y_i \\ 1 \end{bmatrix} \quad \begin{aligned} \lambda \mathbf{x}'_i &= \mathbf{H} \mathbf{x}_i \\ \mathbf{x}'_i \times \mathbf{H} \mathbf{x}_i &= 0 \end{aligned}$$

$$\begin{bmatrix} x'_i \\ y'_i \\ 1 \end{bmatrix} \times \begin{bmatrix} \mathbf{h}_1^T \mathbf{x}_i \\ \mathbf{h}_2^T \mathbf{x}_i \\ \mathbf{h}_3^T \mathbf{x}_i \end{bmatrix} = \begin{bmatrix} y'_i \mathbf{h}_3^T \mathbf{x}_i - \mathbf{h}_2^T \mathbf{x}_i \\ \mathbf{h}_1^T \mathbf{x}_i - x'_i \mathbf{h}_3^T \mathbf{x}_i \\ x'_i \mathbf{h}_2^T \mathbf{x}_i - y'_i \mathbf{h}_1^T \mathbf{x}_i \end{bmatrix}$$

$$\begin{bmatrix} 0^T & -\mathbf{x}_i^T & y'_i \mathbf{x}_i^T \\ \mathbf{x}_i^T & 0^T & -x'_i \mathbf{x}_i^T \\ -y'_i \mathbf{x}_i^T & x'_i \mathbf{x}_i^T & 0^T \end{bmatrix} \begin{pmatrix} \mathbf{h}_1 \\ \mathbf{h}_2 \\ \mathbf{h}_3 \end{pmatrix} = 0 \quad \begin{aligned} &3 \text{ equations,} \\ &\text{only 2 linearly} \\ &\text{independent} \end{aligned}$$


Direct linear transform

$$\begin{bmatrix} 0^T & \mathbf{x}_1^T & -y'_1 \mathbf{x}_1^T \\ \mathbf{x}_1^T & 0^T & -x'_1 \mathbf{x}_1^T \\ \dots & \dots & \dots \\ 0^T & \mathbf{x}_n^T & -y'_n \mathbf{x}_n^T \\ \mathbf{x}_n^T & 0^T & -x'_n \mathbf{x}_n^T \end{bmatrix} \begin{pmatrix} \mathbf{h}_1 \\ \mathbf{h}_2 \\ \mathbf{h}_3 \end{pmatrix} = 0 \quad \mathbf{A} \mathbf{h} = 0$$

- H has 8 degrees of freedom (9 parameters, but scale is arbitrary)
- One match gives us two linearly independent equations
- Homogeneous least squares: find $\mathbf{h}$ minimizing $\|\mathbf{A}\mathbf{h}\|^2$
 - Eigenvector of $\mathbf{A}^T\mathbf{A}$ corresponding to smallest eigenvalue
 - Four matches needed for a minimal solution


Robust feature-based alignment

- So far, we've assumed that we are given a set of “ground-truth” correspondences between the two images we want to align
- What if we don't know the correspondences?



Robust feature-based alignment



- So far, we've assumed that we are given a set of “ground-truth” correspondences between the two images we want to align
- What if we don't know the correspondences?



Robust feature-based alignment




Robust feature-based alignment




- Extract features

Robust feature-based alignment




- Extract features
- Compute *putative matches*

Robust feature-based alignment



- Extract features
- Compute *putative matches*
- Loop:
 - *Hypothesize* transformation T

Robust feature-based alignment



- Extract features
- Compute *putative matches*
- Loop:
 - *Hypothesize* transformation T
 - *Verify* transformation (search for other matches consistent with T)

Robust feature-based alignment



- Extract features
- Compute *putative matches*
- Loop:
 - *Hypothesize* transformation T
 - *Verify* transformation (search for other matches consistent with T)