

Introduction to Computer Vision

Alvin Lin

August 2018 - December 2018

Line Fitting

Line fitting is an operation used for fitting curves and learning data manifolds. The goal is to find a line that best explains some observed data. Given a data point x_i , a target point y_i , and line parameters w, b , a line model would be in the form of $y_i = wx_i + b$. In general, we want to minimize the error between our line model and the actual data. We can calculate the sum of squared error to obtain this metric.

$$\sum_{i=1}^n (y_i - wx_i + b)^2$$

Samples vs Model Parameters

With m samples and n model parameters, we have three cases that can result:

- $m = n$: a unique solution to model this data.
- $m > n$: over-determined system of equations. No solution exists, so we minimize error through line fitting.
- $m < n$: under-determined system of equations. Infinite solutions exists.

Robustness

Squared error can be a source of bias in the presence of noise points in the data. One fix for this is expectation maximization (not discussed in this course). Alternatively, we can use M-estimators or the RANSAC algorithm.

RANSAC

Random Sample Consensus (RANSAC) is used for parametric matching and model fitting. Given a dataset of n fitting the best possible line through brute force would require searching through 2^n , which is often not feasible. RANSAC assumes the data contains both inliers and outliers and goes through the following process.

- Select a random subset of the original data to be hypothetical inliers.
- A model is fitted to the set of hypothetical inliers.
- All other data points are then tested against the fitted model using some model-specific loss function.
- The estimated model is reasonably good if sufficiently many points have been classified as part of the consensus set.
- Afterwards, the model may be improved by re-estimating it using all members of the consensus set.

In general this works because the probability that an inlier is selected p is defined as:

$$p = 1 - (1 - w^n)^k$$

where w is the ratio of inliers to total points, n is the minimum number of points required, and k is the number of iterations. RANSAC is robust even when there are significant number of outliers and is useful in many advanced computer vision applications, but it is not always able to find the optimal set. It usually performs badly when the number inliers is less than 50%. There is no upper bound on the time it takes to compute these parameters and RANSAC requires setting many problem-specific thresholds. Additionally, it can only estimate one model for a particular dataset.

You can find all my notes at <http://omgimanerd.tech/notes>. If you have any questions, comments, or concerns, please contact me at alvin@omgimanerd.tech