

# Robust Estimation

**Introduction.** A common computational problem in vision is to estimate the parameters of a model from image data.

Examples of parameterized models to be fit to image data include lines and ellipses, camera calibration models, image motion models, 3D planar regions, 3D models, and human face models.

## Key Difficulties:

- The models must be fit to noisy image data.
- Initial guesses for the models must be generated automatically.
- Multiple occurrences of the models are often represented in the data. But the number and types of models are typically unknown a priori. They must also be determined from the data.
- The data typically contains (structured) outliers, i.e., observations that do not belong to the model being fitted. These must somehow be ignored during the model fitting.

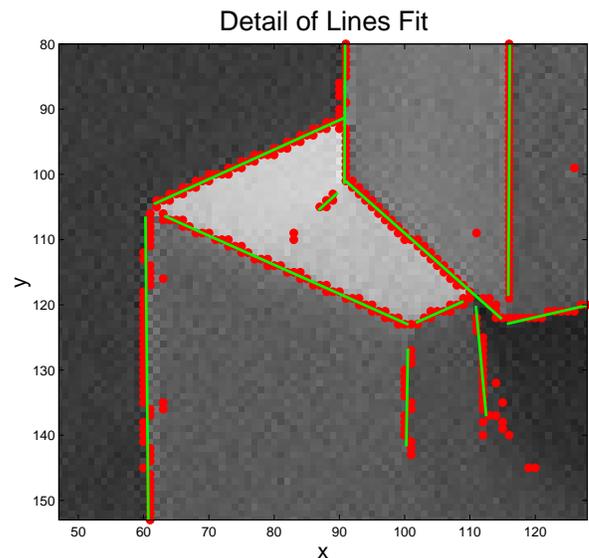
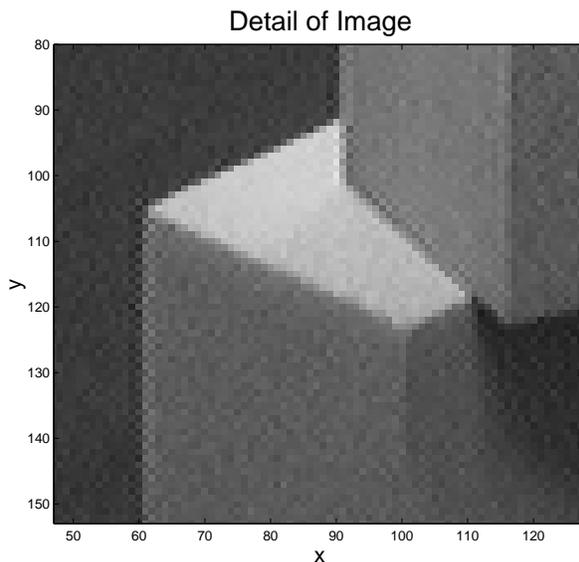
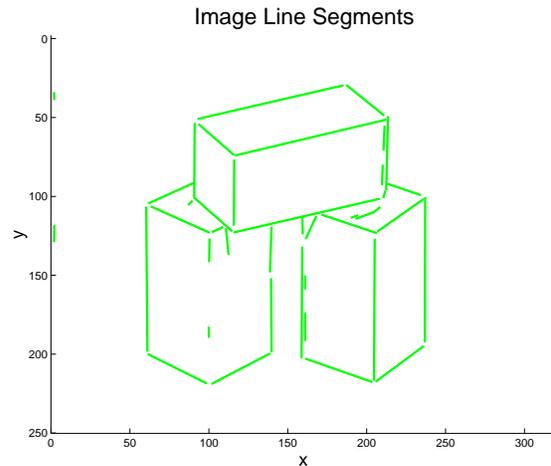
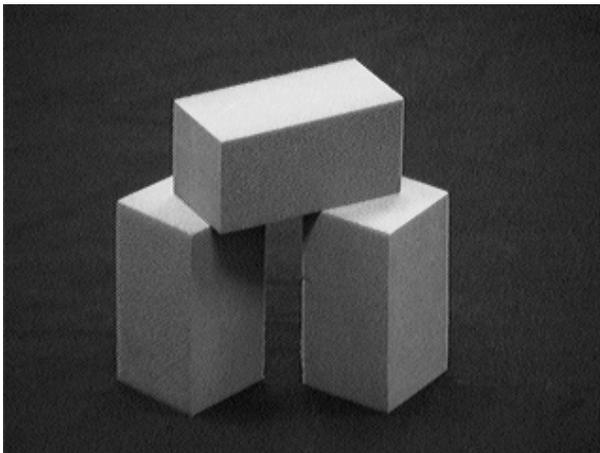
**Reading:** Chapter 15 of Forsyth and Ponce (skip Sec. 15.6 for now).

## Working Example Problem: Fitting Lines

**Example Problem:** Find the best fitting line(s) to a set of image edgel positions,

$$D = \{\vec{x}_k\}_{k=1}^K.$$

For simplicity, we ignore edgel strengths and orientations here.



Robust line estimator output. The Canny edgel positions  $\{\vec{x}_k\}$  are marked (red dots), along with the fitted line segments (green).

## Current Goals

We use the example of estimating image lines to:

- Introduce robust M-estimation and show how it deals with outliers.
- Discuss leverage points.
- Introduce methods for generating initial guesses.
- Introduce issues in model selection.
  - We have more to say about this difficult issue later in this course. Here we briefly consider selecting the number of models (i.e., lines), but not the model type (eg. lines versus curves).
- Discuss the general types of errors to be expected.

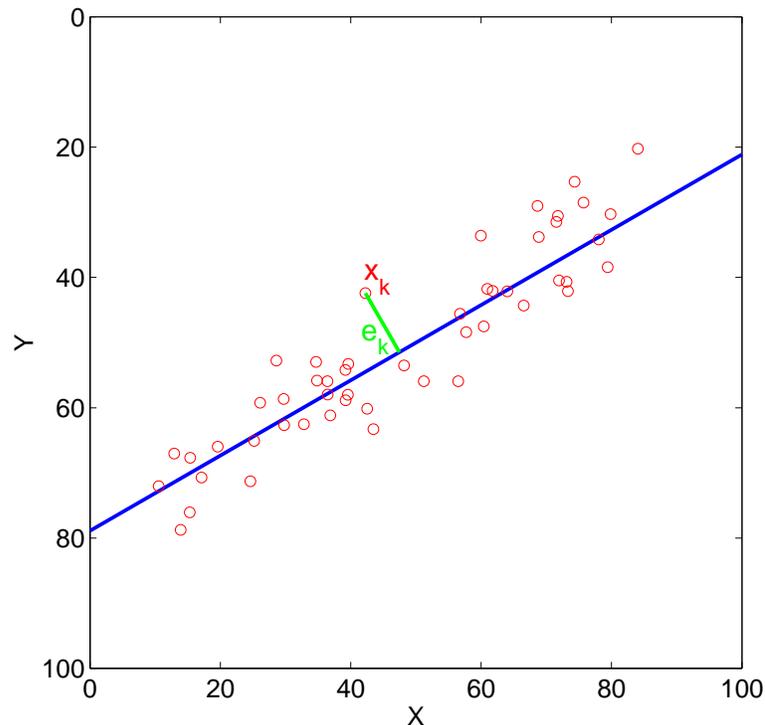
# Perpendicular Error

The equation for points  $\vec{x}$  on one infinite line  $L$  is

$$\vec{n} \cdot \vec{x} + c = 0 .$$

Here we normalize  $\|\vec{n}\| = 1$ . The error in an observed point  $\vec{x}_k$  (relative to the line  $L$ ) is defined to be the perpendicular distance

$$e(\vec{x}_k; \vec{n}, c) \equiv \vec{n} \cdot \vec{x}_k + c .$$



Note that the error perpendicular to the line is used here. This is different from the usual linear regression model for which only the error in the  $y$  component is minimized.

## Least Squares Problem

Given a set of edgel positions  $D = \{\vec{x}_k\}_{k=1}^K$ , consider estimating the line parameters  $(\vec{n}, c)$  by minimizing the squared error

$$\mathcal{O}(\vec{n}, c) \equiv \sum_{k=1}^K (\vec{n}^T \vec{x}_k + c)^2, \quad \text{for } \|\vec{n}\| = 1. \quad (1)$$

By the theory of Lagrange multipliers, the solution must satisfy

$$\frac{\partial \mathcal{L}}{\partial (\vec{n}, c, \lambda)}(\vec{n}, c, \lambda) = \vec{0}, \quad \text{for } \mathcal{L}(\vec{n}, c, \lambda) \equiv \mathcal{O}(\vec{n}, c) - \lambda(\|\vec{n}\|^2 - 1).$$

The derivatives of  $\mathcal{L}$  with respect to  $\vec{n}$ ,  $c$ , and  $\lambda$  give (respectively):

$$\begin{aligned} \sum_{k=1}^K 2(\vec{n}^T \vec{x}_k + c)\vec{x}_k - 2\lambda\vec{n} &= \vec{0}, \\ \sum_{k=1}^K 2(\vec{n}^T \vec{x}_k + c) &= 0, \\ \|\vec{n}\|^2 - 1 &= 0. \end{aligned} \quad (2)$$

## Least Squares Solution

The solution to the LS problem above satisfies three necessary conditions:

$$\sum_{k=1}^K (\vec{n}^T \vec{x}_k + c) \vec{x}_k - \lambda \vec{n} = \vec{0}, \quad (3)$$

$$\sum_{k=1}^K \vec{n}^T \vec{x}_k + c = 0, \quad (4)$$

$$\|\vec{n}\|^2 - 1 = 0. \quad (5)$$

First, it is straightforward to show from Eqn. (4) that  $c$  must satisfy

$$c = -\vec{n}^T \vec{m}, \quad (6)$$

where  $\vec{m} = \frac{1}{K} \sum_{k=1}^K \vec{x}_k$  is the mean edgel location. Substituting  $-\vec{m}^T \vec{n}^T$  for  $c$  into the objective function (on the previous page) yields

$$\mathcal{O} = \sum_{k=1}^K (\vec{n}^T (\vec{x}_k - \vec{m}))^2. \quad (7)$$

In this form it is clear that the optimal line passes through the mean of the edgel positions.

Using (6) and minor algebraic manipulation, we can also rewrite the first constraint above in (3) as

$$\sum_{k=1}^K \vec{x}_k (\vec{x}_k - \vec{m})^T \vec{n} = \lambda \vec{n} \quad (8)$$

This is further simplified using the covariance matrix of the edgel locations, i.e.,

$$C = \frac{1}{K} \sum_{k=1}^K (\vec{x}_k - \vec{m})(\vec{x}_k - \vec{m})^T = \frac{1}{K} \left( \sum_{k=1}^K \vec{x}_k \vec{x}_k^T \right) - \vec{m} \vec{m}^T \quad (9)$$

That is, Eqn. (8) is equivalent to

$$C \vec{n} = \mu \vec{n}, \quad (10)$$

which shows that  $\vec{n}$  is an eigenvector of the covariance matrix  $C$ .

Finally, using (5), (6) and (10), we can simplify the objective function to obtain:

$$\mathcal{O} = \sum_{k=1}^K (\vec{n}^T (\vec{x}_k - \vec{m}))^2 = K \vec{n}^T C \vec{n} = K \mu. \quad (11)$$

Therefore, to minimize  $\mathcal{O}$ , we select the eigenvector corresponding to the smallest eigenvalue of  $C$ .

## Least Squares Solution

The gradient equations (2) can be simplified using

$$\vec{m} = \frac{1}{K} \sum_{k=1}^K \vec{x}_k, \quad C = \frac{1}{K} \sum_{k=1}^K (\vec{x}_k - \vec{m})(\vec{x}_k - \vec{m})^T,$$

namely the mean and covariance of the data.

With  $\vec{m}$  and  $C$  one can show that the least squares solution,  $(\vec{n}, c)$ , must satisfy

$$\begin{aligned} C\vec{n} &= \mu\vec{n}, \quad \|\vec{n}\| = 1, \\ c &= -\vec{n}^T \vec{m}. \end{aligned} \tag{12}$$

That is,

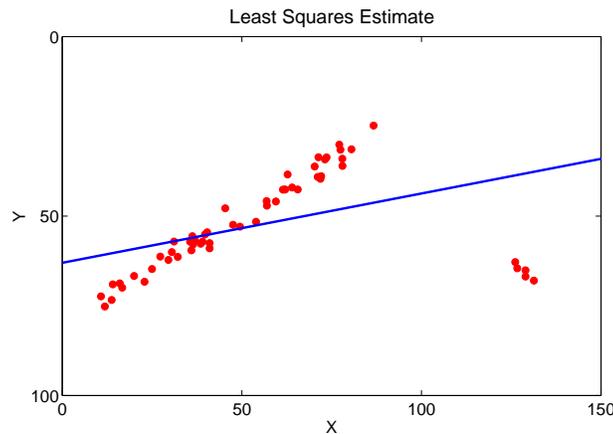
- $\vec{n}$  must be an eigenvector of  $C$ ;
- $c$  must be chosen such that the estimated line passes through the mean of the points;
- and for the solution to be a local minimum, it also can be shown that  $\vec{n}$  must be the eigenvector for the *minimum* eigenvalue of  $C$ .

The solution is therefore unique and easy to compute. No initial guess is required.

## Least Squares and Outliers

Unfortunately, least squares solutions are sensitive to outliers in the data.

For example, consider the set of edgel point data  $\{\vec{x}_k\}_{k=1}^K$  depicted by red dots:



The least squares estimate (blue line, above) is strongly influenced by the small cluster of outliers.

In other words, this LS approach is not directly suitable for fitting data with outliers or when multiple solutions are expected.

## Hough Transform

The Hough transform (HT) is an early method for dealing with outliers and multiple solutions in parameter estimation problems.

For our working example, image lines  $\vec{n}^T(\theta)\vec{x} + c = 0$  are parameterized by  $\theta$  and  $c$ , with  $\vec{n}(\theta) = (\cos(\theta), \sin(\theta))^T$ .

A discrete voting space is formed by quantizing  $\theta \in [0, \pi)$  and  $c$ . For each edgel  $\vec{x}_k$ , and each discrete  $\theta_i$ , we add one vote to the bin  $(\theta_i, c_j)$  which has minimal absolute error, that is

$$\vec{n}^T(\theta_i) \vec{x}_k + c_j \in [-\delta c/2, \delta c/2). \quad (13)$$

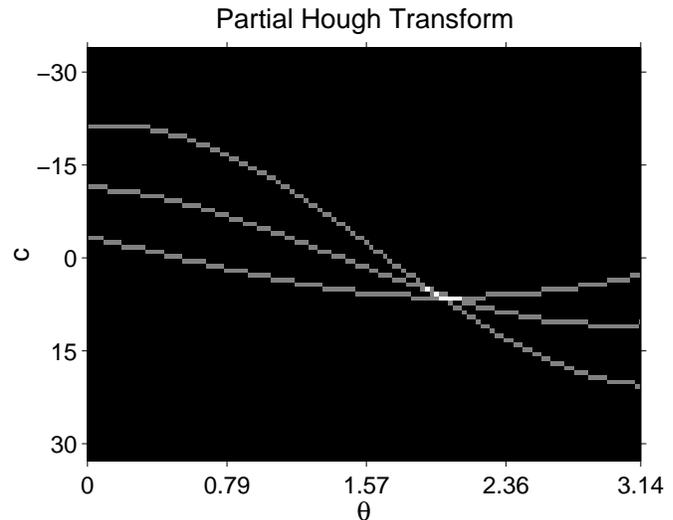
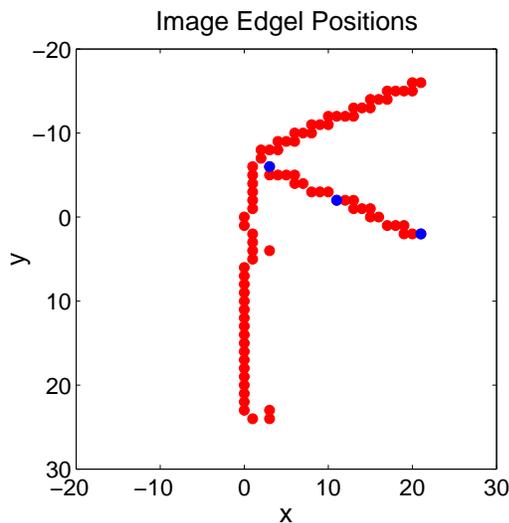
Here  $\delta c$  is the bin spacing for  $c$ .

Eqn. (13) dictates that each edgel  $\vec{x}_k$  votes for a sinusoidal curve in the Hough space (see example below).

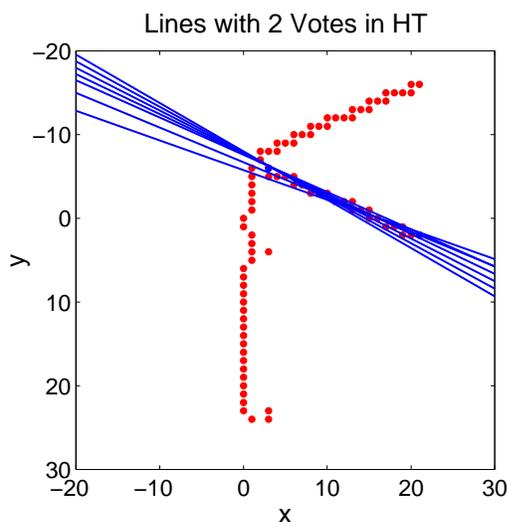
The accumulation of votes for all the edgels  $\vec{x}_k$  provides the Hough transform  $H$ .

# Hough Example

For example, the Hough transform for only the three blue edgels below is shown on the right.



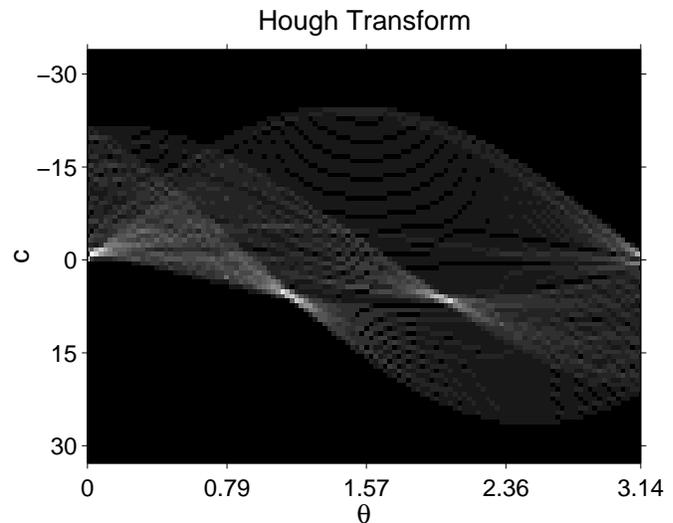
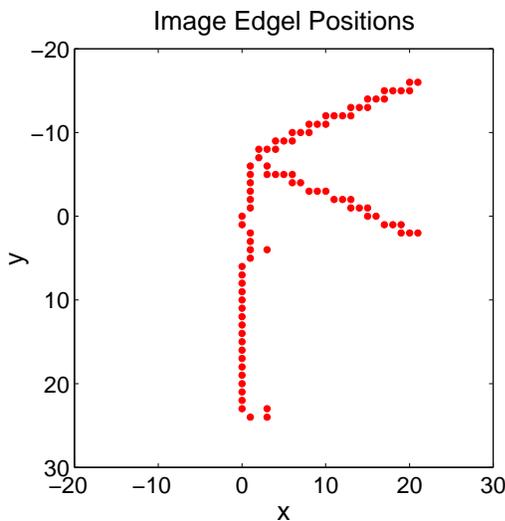
The maximum number of votes here is 2, not 3, since the three edgels are not precisely colinear. The centers of the bins with 2 votes provide the parameters for blue lines below.



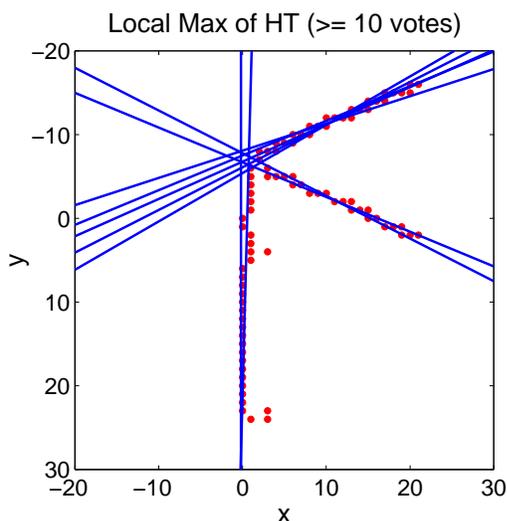
Increasing the bin size  $\delta c$  would allow one (or more) bins to have votes from all three edgels. In practice, selecting appropriate bin sizes for the HT can be tricky.

## Hough Example (Cont.)

The Hough transform for all the edgels in the previous example is shown on the right. Note the symmetry  $H(\theta_i, c_j) = H(\theta_i + \pi, -c_j)$ , which is due to the symmetry of the error  $e = \vec{n}^T(\theta)\vec{x} + c$ .



Three strong peaks are visible in the HT, corresponding roughly to the three line segments apparent in the data. (Enlarge the neighbourhood of these peaks when viewing an electronic copy.)



The blue lines result from local non-max suppression of the HT above, then thresholding the results at 10 votes.

Results are sensitive to the particular values used for this threshold and the bin sizes. Good values for these parameters often depend on the image data.

## Hough Transform Summary

The Hough transform is a reasonable approach for simple data, for which clear peaks can be expected. However,

- **Goldilocks' problem:** It is difficult to set appropriate bin sizes:
  - **Too large:** Poor resolution of parameters.
  - **Too small:** Peaks in HT broken into pieces, each with fewer votes.
  - **Just right:** Isolated peaks at appropriate parameter values.
- Poor detector performance (i.e. the trade-off between false positives and false negatives). In noisy and/or cluttered datasets many extraneous bins can have vote counts comparable to, or larger than, the counts for the desired models.
- In practice, fractional votes for neighbouring Hough bins are also included to help smooth the HT and reduce discretization artifacts.
- The number of bins grows exponentially with the dimension of the unknown parameters (eg.  $n^d$  for  $n$  bins in each of  $d$  dimensions).

For more complex data sets, a much better alternative is to use robust M-estimation, which we discuss next.

## Robust M-estimation

We seek line parameters  $(\vec{n}, c)$  which (locally) solve the minimization problem

$$\min \mathcal{O}(\vec{n}, c), \quad \text{for } \|\vec{n}\| = 1, \quad (14)$$

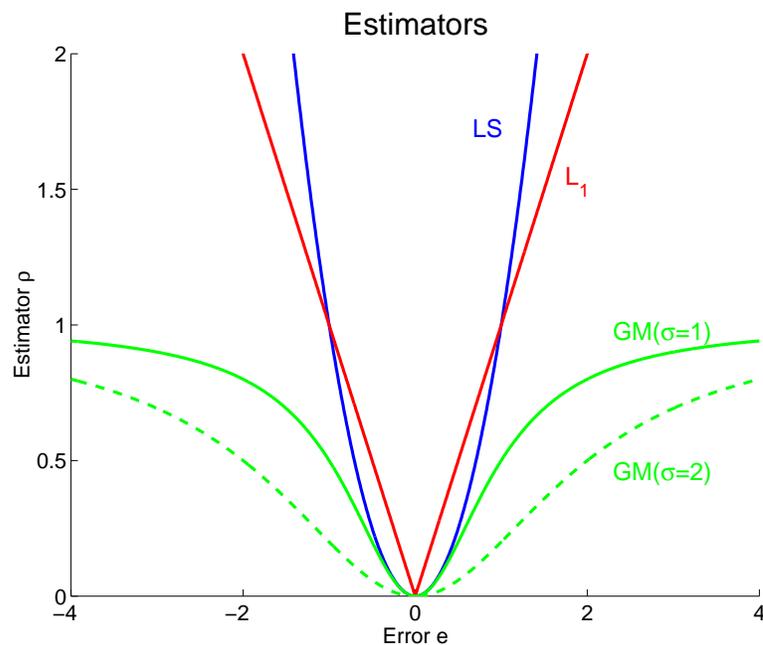
$$\mathcal{O}(\vec{n}, c) \equiv \sum_{k=1}^K \rho(e(\vec{x}_k; \vec{n}, c)). \quad (15)$$

Here  $\rho(e)$  is the “estimator”, which provides a cost for any given error  $e$ . Common choices include:

$$\rho(e) = e^2, \quad \text{Least squares (LS) estimator,}$$

$$\rho(e) = |e|, \quad L_1 \text{ estimator,}$$

$$\rho(e) = \frac{e^2}{\sigma^2 + e^2}, \quad \text{Geman-McLure (GM) estimator.}$$

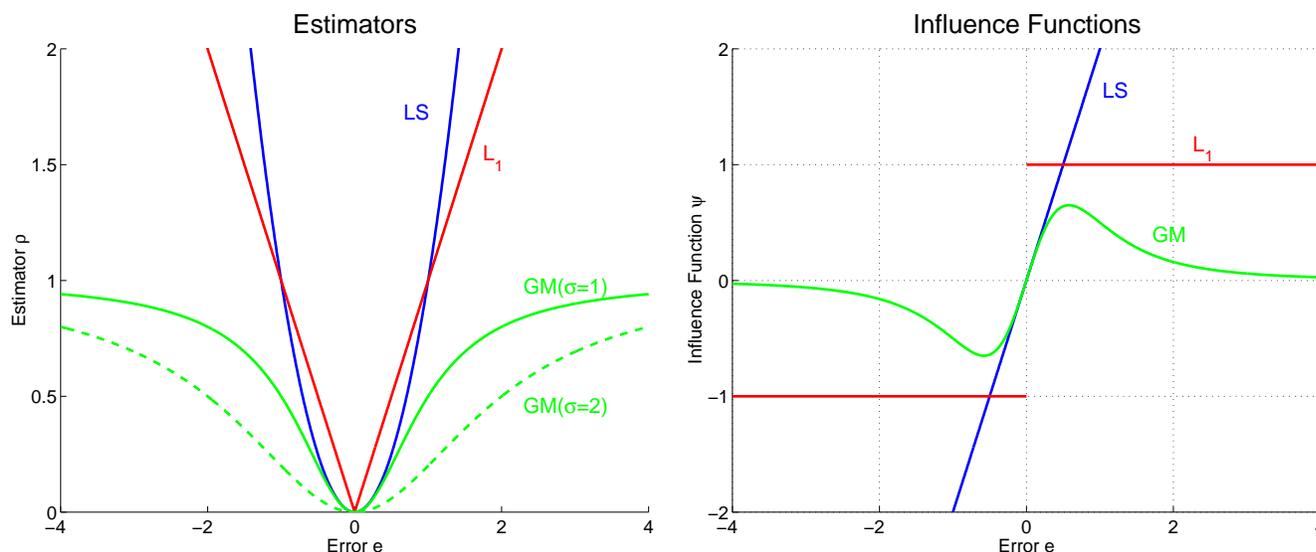


# Influence Function

The influence function describes the sensitivity of the overall estimate,  $(\vec{n}, c)$ , on data with error  $e$ .

**Definition:** Given the estimator,  $\rho(e)$ , the influence function,  $\psi(e)$ , is defined to be:

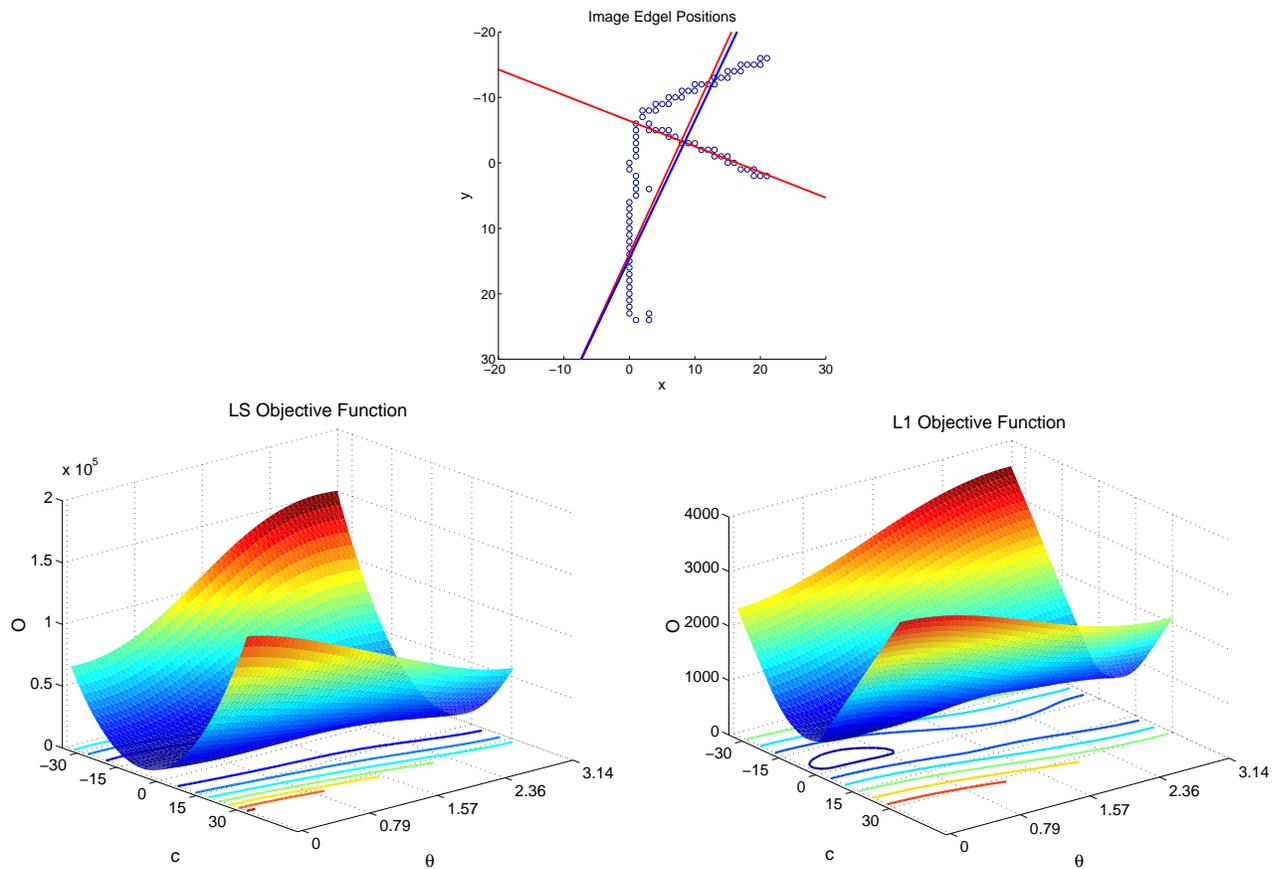
$$\psi(e) = \frac{d\rho}{de}(e). \quad (16)$$



*Remark:* Unlike LS, the influence functions for  $L_1$  and GM are bounded. Moreover, GM is a ‘redescending estimator’, that is,  $\psi(e) \rightarrow 0$  as  $|e| \rightarrow \infty$ .

## Objective Function Examples: LS and L<sub>1</sub>

Given a data set  $\{\vec{x}_k\}_{k=1}^K$ , we can plot the objective function  $\mathcal{O}(\vec{n}(\theta), c)$  for  $\vec{n}(\theta) = (\cos(\theta), \sin(\theta))^T$ . Note that, since  $\vec{n}(\theta + \pi) = -\vec{n}(\theta)$ , the objective function is periodic in  $\theta$  with  $\mathcal{O}(\vec{n}(\theta + \pi), c) = \mathcal{O}(\vec{n}(\theta), -c)$ .

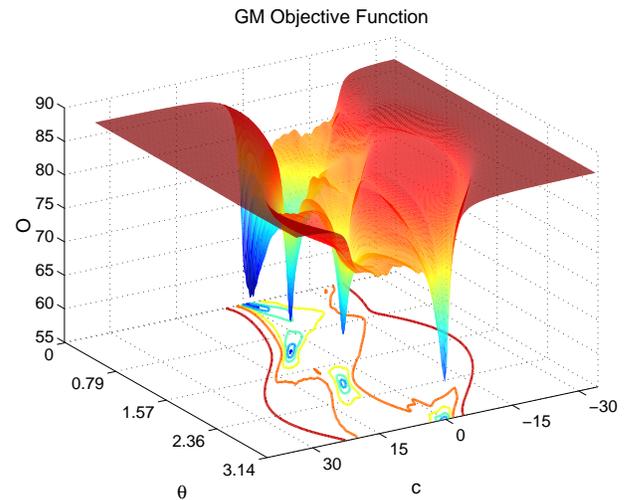
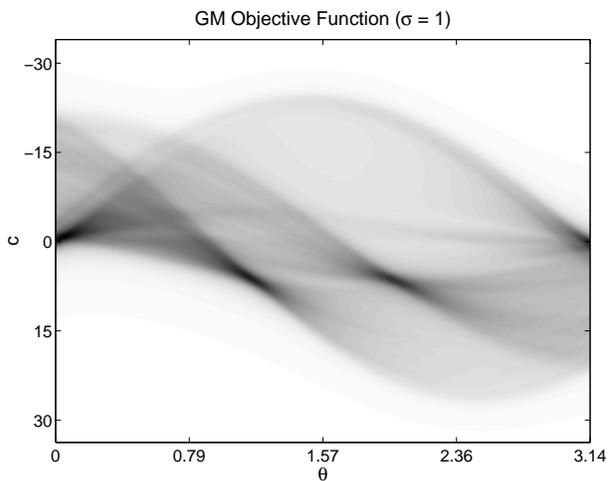


The LS estimator always provides a *unique* local minimum (blue line in top figure). The L1 estimator can have multiple local minima. In this case, there are two local minima (red lines in top figure).

These surfaces (and the fits above) indicate the significant influence of data with large errors.

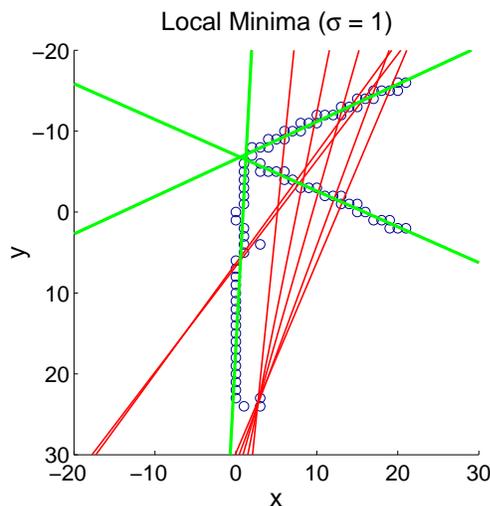
## Objective Function Examples: GM (with $\sigma = 1$ )

In contrast, for the GM estimator the influence of a data point with large error (essentially) vanishes. For the same data as before:



For the GM estimator, the objective function exhibits 3 strong minima (the green lines below). The outer pair of the deep valleys above are actually parts of the same valley, due to periodicity in  $\theta$ .

There are many weaker local minima (red lines below).

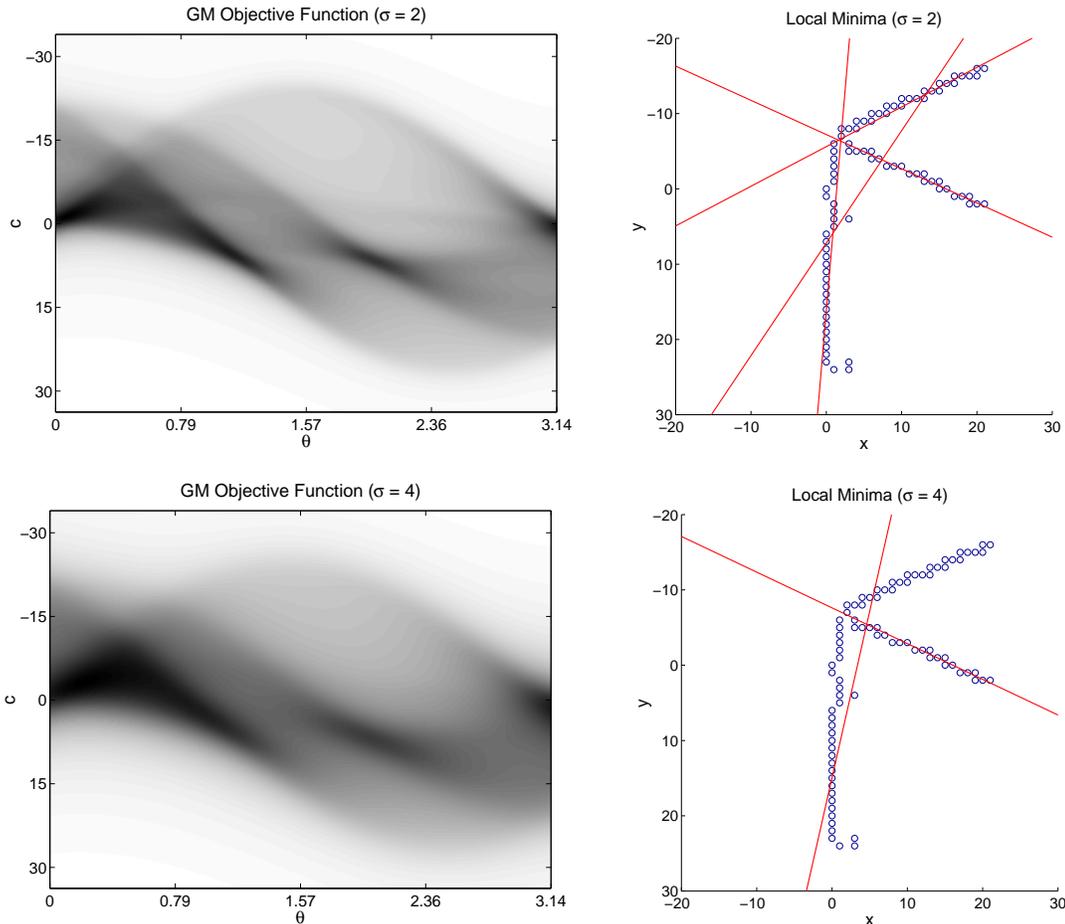


We show only those local minima for which  $\mathcal{O} < K - 10$ , where  $K$  is the number of data points.

Note that  $\rho(e) \leq 1$ , so  $K$  is the maximum possible value for  $\mathcal{O}$ .

# Objective Function Examples: GM, Varying $\sigma$

Increasing  $\sigma$  in the GM-estimator smooths the objective function:



- Note there are fewer local minima for larger  $\sigma$ . (We show *all* the local minimum beyond minor fluctuations.)
- As  $\sigma$  is increased, outliers receive greater influence, and therefore the bias in the estimates often increases (e.g., bottom right plot).
- As  $\sigma \rightarrow \infty$  the objective function can be shown to approximate the one for least squares (with a unique solution).

## Robust Objective Function and the Hough Transform

The robust objective function  $\mathcal{O}(\theta, c)$  is closely related to a smoothed Hough transform (e.g., see the figures on p. 10 and 15).

**Theorem.** Suppose  $\rho(e)$  is a smooth redescending estimator such as GM, with  $\rho(e) \rightarrow 1$  as  $|e| \rightarrow \infty$ . Let  $H(\theta_i, c_j)$  be the Hough transform of  $K$  edgels using small bin sizes  $\delta\theta$  and  $\delta c$ . Consider the discrete 1D filter kernel  $f(c_j)$  obtained by sampling  $1 - \rho(c)$  at  $c_j = j\delta c$ ,  $j = 0, \pm 1, \pm 2, \dots$ . Define

$$S(\theta_i, c_j) = (f *_{2} H)(\theta_i, c_j).$$

Here  $*_{2}$  denotes convolution with respect to the second argument. So  $S(\theta_i, c_j)$  is simply  $H(\theta_i, c_j)$  blurred in the  $c$ -direction by  $\rho(c_j)$ . Then

$$\mathcal{O}(\theta_i, c_j) = K - S(\theta_i, c_j) + O(\delta\theta + \delta c),$$

as  $\delta\theta, \delta c \rightarrow 0$ . That is, the blurred Hough transform  $S$  approximates  $\mathcal{O}$ , up to a sign change and additive constant.

The smoothness of  $\mathcal{O}(\theta, c)$  is useful for locating local minima precisely. Indeed, as mentioned above, the HT is often smoothed in practice. Moreover, on the next few pages we show how to find local minima **without** densely sampling the objective function, which can provide a considerable computational savings over HT.

## Weight Function

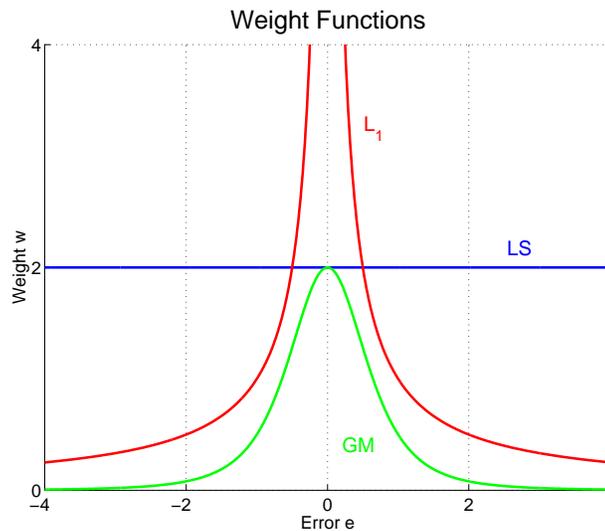
It is sometimes convenient to re-express a robust estimator in a form resembling weighted least squares.

The gradient of the objective  $\mathcal{O} = \sum_i \rho(e_i(\vec{\theta}))$  with respect to model parameters  $\vec{\theta}$  is

$$\frac{\partial \mathcal{O}}{\partial \vec{\theta}} = \sum_i \frac{\partial \rho}{\partial e}(e_i) \frac{\partial e_i}{\partial \vec{\theta}}(\vec{\theta}) = \sum_i w(e_i) e_i(\vec{\theta}) \frac{\partial e_i}{\partial \vec{\theta}}(\vec{\theta}) \quad (17)$$

with weights

$$w(e_i) = \frac{1}{e_i} \frac{\partial \rho}{\partial e}(e_i) = \frac{1}{e_i} \psi(e_i). \quad (18)$$



*Remarks:*

- For  $\rho(e) = e^2$ , the weights are constant,  $w = 2$ , and (17) is just LS.
- If the weights were independent of  $\vec{\theta}$ , then (17) would be the gradient of a *weighted* LS estimator.
- For robust estimators, for bounded weights, Eqn (17) has the form of *weighted* LS, but it is nonlinear as the weights depend on  $\vec{\theta}$ .

## Weighted Mean and Covariance

Back to our robust line fitting problem: Minimize  $\mathcal{O}(\vec{n}, c)$ , but without densely sampling the parameter space.

Extending our previous solution, we now use the weighted mean and covariance:

$$\vec{m} = \frac{1}{s} \sum_{k=1}^K w(e_k) \vec{x}_k, \quad (19)$$

$$C = \frac{1}{s} \sum_{k=1}^K w(e_k) (\vec{x}_k - \vec{m})(\vec{x}_k - \vec{m})^T, \quad (20)$$

where  $s \equiv \sum_{k=1}^K w(e_k)$  and  $e_k \equiv \vec{n}^T \vec{x}_k + c$ .

We show in the subsequent notes that a necessary condition for  $(\vec{n}, c)$  to be a solution of (14) is that

$$C\vec{n} = \mu\vec{n}, \quad (21)$$

$$c = -\vec{n}^T \vec{m}, \quad \|\vec{n}\| = 1. \quad (22)$$

where  $\mu$  is the minimum eigenvalue of the  $2 \times 2$  covariance matrix  $C$ .

Properties of the solution:

- $C$  typically depends on  $(\vec{n}, c)$  through the weights  $w(e_k)$  (as does  $\vec{m}$ ), and therefore (21) is a *nonlinear* eigenvalue problem for  $\vec{n}$ .
- For LS, (21) becomes the *linear* eigenvalue problem in (12).
- Eqn. (22) implies the line must pass through the weighted mean  $\vec{m}$ .

## Notes: Conditions for Objective Function Minima

Suppose  $\rho(e)$  is an estimator with a bounded weight function  $w(e)$ . We wish to minimize  $\mathcal{O}(\vec{n}, c)$ , subject to the constraint  $\|\vec{n}\| = 1$ . A standard approach for such a constrained optimization problem is to introduce a Lagrange multiplier  $\lambda$ , and consider

$$\mathcal{L}(\vec{n}, c, \lambda) = \mathcal{O}(\vec{n}, c) - \lambda(\|\vec{n}\|^2 - 1).$$

A necessary condition for  $(\vec{n}, c)$  is that this Lagrangian  $\mathcal{L}$  has a stationary point at  $(\vec{n}, c, \lambda)$ , that is

$$\frac{\partial \mathcal{L}}{\partial(\vec{n}, c, \lambda)}(\vec{n}, c, \lambda) = \vec{0}.$$

For  $\mathcal{O}(\vec{n}, c)$  as in (14) the derivatives of  $\mathcal{L}$  with respect to  $\vec{n}$ ,  $c$ , and  $\lambda$  give (respectively):

$$\begin{aligned} sA\vec{n} + s\vec{m}c - 2\lambda\vec{n} &= \vec{0}, \\ s\vec{m}^T\vec{n} + sc &= 0, \\ \|\vec{n}\|^2 - 1 &= 0. \end{aligned} \tag{23}$$

Here  $s \equiv \sum_{k=1}^K w(e_k)$  is the sum of the weights for errors  $e_k \equiv \vec{n}^T \vec{x}_k + c$ ,  $\vec{m} \equiv \frac{1}{s} \sum_{k=1}^K w(e_k) \vec{x}_k$  is the weighted mean of the data  $\{\vec{x}_k\}_{k=1}^K$ , and  $A \equiv \frac{1}{s} \sum_{k=1}^K w(e_k) \vec{x}_k \vec{x}_k^T$  is the weighted correlation matrix. Solving the second equation for  $c$  gives  $c = -\vec{m}^T \vec{n}$ . Substituting this into the first equation gives the (*nonlinear*) eigenvalue problem for  $\vec{n}$ ,

$$\frac{s}{2}(A - \vec{m}\vec{m}^T)\vec{n} = \lambda\vec{n}. \tag{24}$$

A short calculation shows that the weighted covariance matrix in (20) satisfies  $C = A - \vec{m}\vec{m}^T$ . Therefore (24) takes the desired form,  $C\vec{n} = \mu\vec{n}$ , for  $\mu = 2\lambda/s$ . Finally, it can be shown that the objective function is locally minimized only when  $\vec{n}$  is chosen to be the eigenvector for the *minimum* eigenvalue of  $C$ .

## Iteratively Reweighted Least Squares Algorithm

There are many ways to compute minima of  $\mathcal{O}$ . Here we consider one approach, motivated by Eqn (17), that is intuitive and often effective. (That is, there may be better ways to optimize  $\mathcal{O}$ .)

For the robust line fitting problem, given an initial guess for the line parameters,  $(\vec{n}_0, c_0)$ , we can update the parameters as follows:

1. Compute the weights  $w_k = w(\vec{n}_0^T \vec{x}_k + c_0)$ .
2. Compute  $\vec{m}$  and  $C$  as in (19) and (20).
3. Set  $\vec{n}$  to be the eigenvector for the minimum eigenvalue of  $C$ .
4. Set  $c = -\vec{n}^T \vec{m}$ .

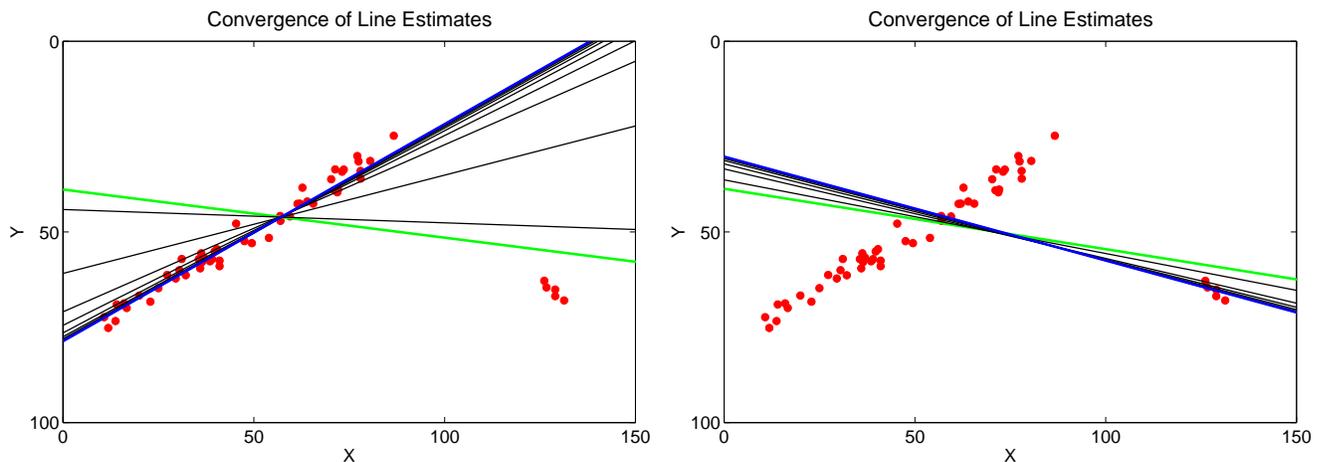
Note that the updated estimate for  $(\vec{n}, c)$  is not necessarily a local minima of (14), since  $\vec{m}$  and  $C$  were computed using  $(\vec{n}_0, c_0)$  instead of  $(\vec{n}, c)$ .

Resetting  $(\vec{n}_0, c_0)$  to be the newly computed  $(\vec{n}, c)$ , we iterate the above steps until the update of both  $\vec{n}$  and  $c$  are small.

This is known as the iteratively reweighted least squares (IRLS) algorithm. See the Matlab tutorial `robustDemo.m` (available from the course homepage) for a demonstration.

## Example Results

For the GM estimator (with  $\sigma$  equal to the standard deviation of the noise for the edgels near the dominant line) the iterations of the IRLS algorithm are shown below (black lines). The initial guesses (green) and the converged states (blue) are also shown.



The solution on the left indicates the outliers have been rejected, while the iterations on the right show the algorithm converging to one of several undesirable local minimum. How can we avoid these?

## Analysis of Influence

To understand these extraneous local minima it is useful to first consider the influence of outliers in detail.

Given a solution  $(\vec{n}_0, c_0)$  of (14), we consider the effect of one additional outlier  $\vec{x}_{K+1}$ . We wish to estimate the perturbation to the solution caused by this outlier. The size of this perturbation will determine the *influence* of that outlier.

To simplify the analysis, we freeze the weights for the original data,  $w_k = w(\vec{n}_0^T \vec{x}_k + c_0)$  for  $k = 1, \dots, K$ , and set  $w_{K+1} = w(\vec{n}_0^T \vec{x}_{K+1} + c_0)$ . With all these weights frozen, we seek the solution of the weighted least squares problem

$$\min \mathcal{O}_{WLS}(\vec{n}, c, \epsilon), \quad \text{for } \|\vec{n}\| = 1, \quad (25)$$

$$\mathcal{O}_{WLS}(\vec{n}, c, \epsilon) \equiv \left[ \sum_{k=1}^K \frac{w_k}{2} (\vec{n}^T \vec{x}_k + c)^2 \right] + \epsilon \frac{w_{K+1}}{2} (\vec{n}^T \vec{x}_{K+1} + c)^2. \quad (26)$$

For  $\epsilon = 0$  we recover Eqns (21) and (22) for the original solution  $(\vec{n}_0, c_0)$ . For  $\epsilon = 1$  the solution is the first update of the IRLS algorithm given the initial guess  $(\vec{n}_0, c_0)$  and the additional point  $\vec{x}_{K+1}$ .

## Influence of Outliers on Line Parameters

**Theorem.** Let  $\psi_{K+1} = \psi(\vec{n}_0^T \vec{x}_{K+1} + c_0)$  be the influence function evaluated at  $\vec{x}_{K+1}$ , and  $\vec{t}_0$  be a unit vector tangent to the initial line (i.e.,  $\vec{t}_0 \perp \vec{n}_0$ ). Then the optimal solution  $(\vec{n}(\epsilon), c(\epsilon))$  of (25) satisfies

$$\left. \frac{dc}{d\epsilon} \right|_{\epsilon=0} = -\frac{\psi_{K+1}}{s}, \quad (27)$$

$$\left. \frac{d\vec{n}}{d\epsilon} \right|_{\epsilon=0} = -\frac{\psi_{K+1}}{(\mu_1 - \mu_2)s} \vec{t}_0 \vec{t}_0^T (\vec{x}_{K+1} - \vec{m}), \quad (28)$$

where  $s \equiv \sum_{k=1}^K w_k$ ,  $\vec{m} = \frac{1}{s} \sum_{k=1}^K w_k \vec{x}_k$ , and  $\mu_1 > \mu_2$  are the two eigenvalues of  $C = \sum_{k=1}^K w_k (\vec{x}_k - \vec{m})(\vec{x}_k - \vec{m})^T$ .

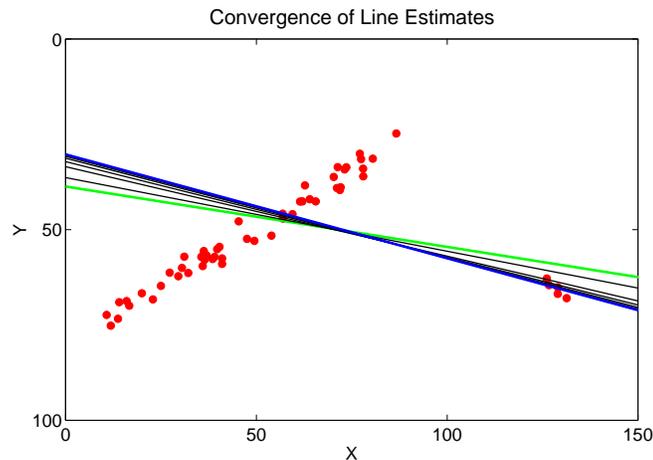
The proof of this theorem is left to the reader.

### Observations:

- Eqn. (27) shows  $dc/d\epsilon$  is proportional to  $\psi(e)$ , motivating the name ‘influence function’ for  $\psi(e)$ .
- Eqn. (27) shows the sense in which  $c$  is stable when  $\psi(e)$  is bounded.
- Eqn. (28) shows  $\vec{n}$  may not be stable, even for a bounded influence function  $\psi(e)$  (since  $|\vec{t}_0^T (\vec{x}_{K+1} - \vec{m})|$  may be large). In addition, notice the effect of  $\mu_1 - \mu_2$ , which measures the eccentricity of the original data points.
- Points  $\vec{x}_{K+1}$  for which  $|\psi_{K+1} \vec{t}_0^T (\vec{x}_{K+1} - \vec{m})|$  is large are called **leverage points**.

## Controlling Leverage

The previous Theorem shows that even a redescending estimator, such as GM, is not robust to leverage points.



In this example (from p. 23) the estimator fails to reject a cluster of outliers in favour of points on the dominant line. For some initial guesses, this cluster has enough leverage to unduly influence the solution.

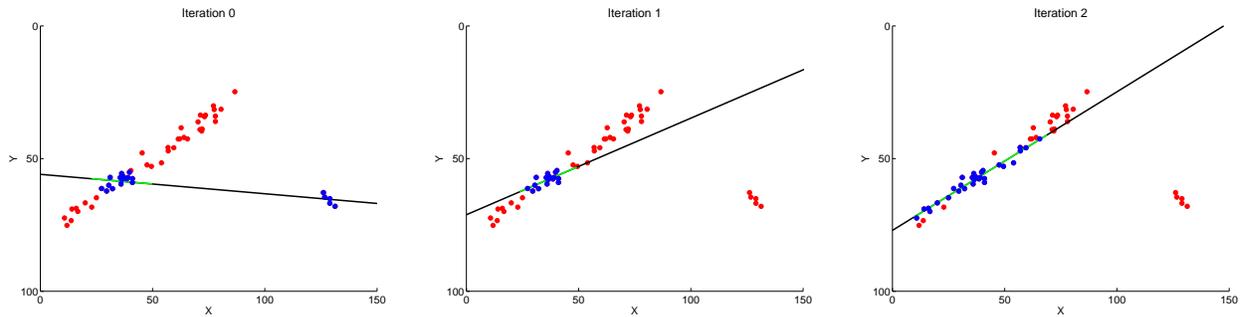
A simple strategy for dealing with leverage points is:

1. Estimate the distribution of data support along the fitted line (eg. project the weights  $w_k$  from  $\vec{x}_k$  onto the line and blur).
2. Determine a contiguous region of support along the fitted line, that is, an interval of support without any large gaps.
3. Reduce the weights  $w_k$  in the IRLS algorithm for points  $\vec{x}_k$  significantly outside the region of contiguous support (eg. set such weights to 0).

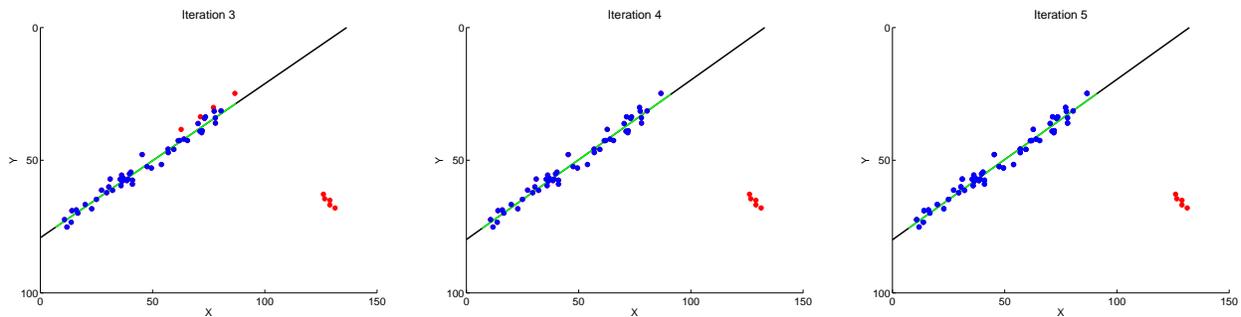
We show an example of this next.

# Leverage Example

We can reduce leverage problems by controlling the support interval within the robust line estimation algorithm.



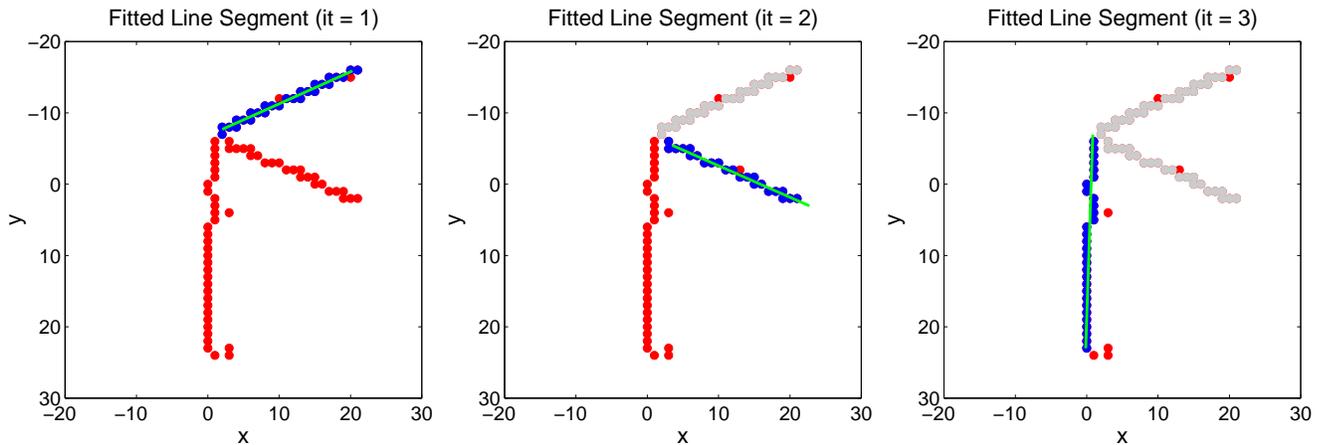
Iterations 0 to 2 above. Current line estimate (black). Edgels with significant weights (blue). The estimate for a contiguous region of support is also shown (green segment).



The remaining iterations 3 to 5, at which point the algorithm has converged.

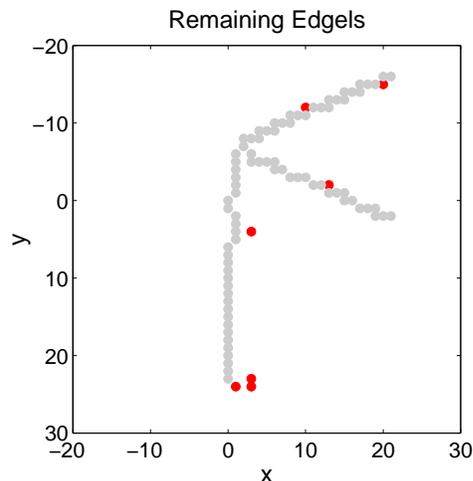
# Successive Estimation

This approach can be applied to find lines one at a time:



The first fitted segment is shown (left). Edgels having significant robust weights (blue), can then be removed from the active data set (red).

A second line segment is fit (middle) to the remaining edgels, and the new supporting edgels (blue) are removed. Then a third segment is fit.



This fitting process continues until no further lines with contiguous support regions can be fit. The remaining active edgels are outliers.

# Outline of a Robust Line Estimation Algorithm

## 1. Initial Guess.

- Randomly sample from the data. Eg. Set  $(\vec{n}_0, c_0)$  according to the position and orientation of a sampled edgel.

## 2. Iterative Fitting.

- Use the iteratively reweighted least squares algorithm to fit the line parameters. Maintain information about the support (from the data) for the line along its length, and use it to downweight data likely to cause leverage problems.

## 3. Verification.

- Given a converged solution, decide whether that line has sufficient data support (e.g., consider the sum of the weights). Discard solutions with insufficient support.

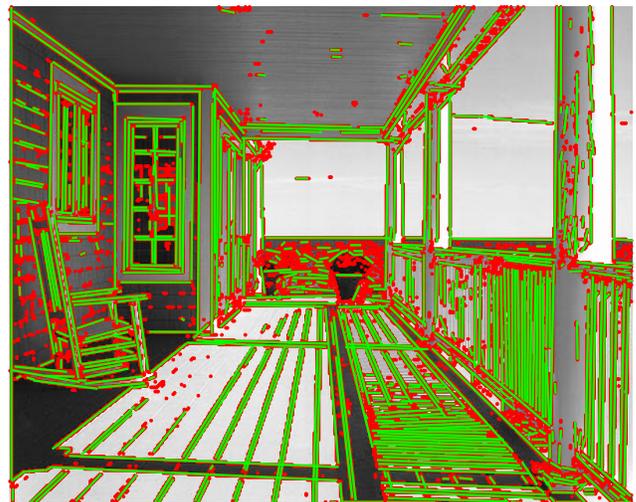
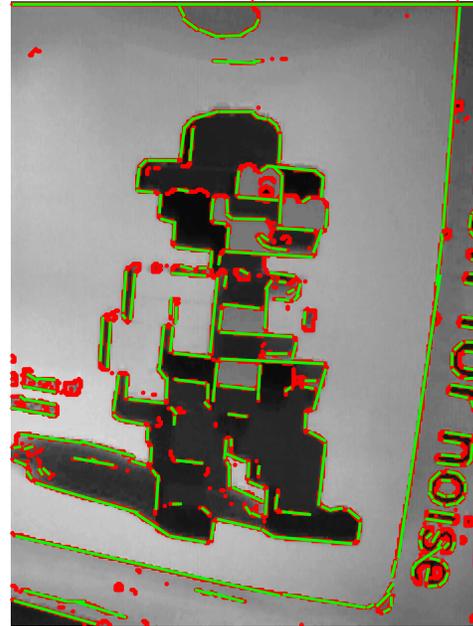
## 4. Model Selection.

- For each verified line segment, remove the edgels that provide significant support for it from the data set.

Repeat steps 1 through 4.

See page 2 and the next two pages for sample results. Similar strategies can be applied to many other estimation problems.

## Linefinder Results



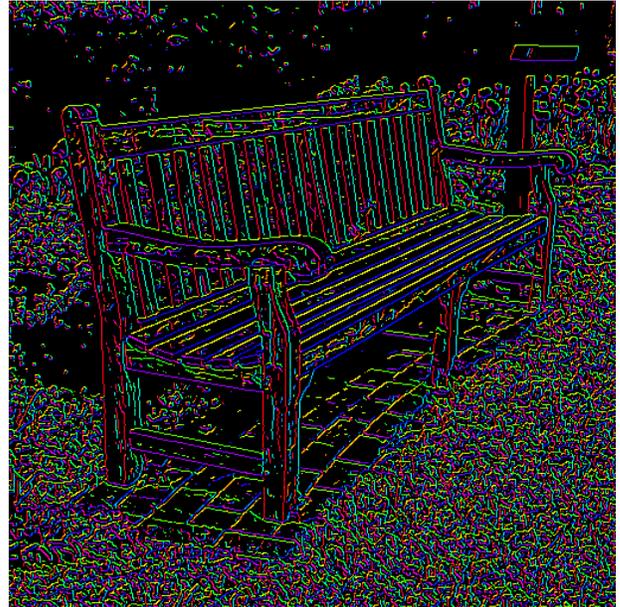
Edgel positions are shown in red, with fitted lines in green. (Try zooming into small regions, such as the rocking chair, to see the fit in detail)

# Linefinder Results

Parkbench Image



Colour-coded Edgel Orientation



Orientation Tensor



Estimated Lines



Random seeds for lines were obtained by sampling the edgels *only* within oriented regions (i.e., red regions in bottom-left image).

## Types of Errors

There are instances of several general types of errors visible in the previous line finder results:

- **Drop-outs:** Parts or all of a true line that are missing in the estimation results.
- **False positives:** An estimated line which does not correspond to any combination of true lines.
- **Over-segmentation:** Breaking one true line into several estimated segments (possibly colinear, parallel or otherwise).
- **Under-segmentation:** Joining two or more true lines into one estimated segment.
- **Parameter noise:** Small errors in the position and orientation estimates.
- **Model-type errors:** The somewhat inappropriate use of our current model (i.e., line segments fitted to edgels) for curves, texture, or thin bars in the image. More generally, model-type errors occur when the system can fit several types of models (eg. curves versus lines), but selects the wrong choice.

## Evaluation of Results

We have not provided any quantitative evaluation of the results which, almost certainly, would provide invaluable data. Why not?

What should we use for ground truth? Human data? Synthetic data?

How should the results from different line-finder algorithms be compared? What evaluation metric should be used? Different algorithms which use the line finder results can be expected to vary in their sensitivity to different types of errors.

In the end what matters is how cost-effective various algorithms are when included in a working robot. (For an example of this style of assessment, see Dougherty and Bowyer, 1998, Objective evaluation of edge detectors using a formally defined framework.)

Here we simply punt (i.e., accept our current gains and move on).

## Robust Estimation: Summary

For the estimation of model parameters from image data:

- A redescending M-estimator, such as the Geman-McLure estimator, downweights the influence of data with large errors.
- The use of such an estimator leads to a non-linear optimization problem for the model parameters.
- The optimization problem typically has multiple local minima, some of which provide useful estimates while others are extraneous.
- The iteratively reweighted least squares algorithm provides one way to find local minima of the resulting objective function.
- Leverage points can significantly skew the parameter estimates, and can be controlled by limiting the spatial extent of the data being fit. This can also reduce the number of extraneous minima.
- Initial guesses can be generated by randomly sampling the data (see RANSAC). Another approach uses continuation with decreasing  $\sigma$  (see deterministic annealing and graduated non-convexity).
- One way to choose the number of models is to iteratively fit individual models, removing the data providing support for each model before fitting the next (see also Bayesian model selection).