Introduction to Computer Vision

Michael J. Black

Nov 2009

Robust statistics and robust motion estimation

November 6, 2009 (Friday) 4:00-5:30 pm [refreshments beginning at 3:45 pm]
Marcuvitz Auditiorium (Room 220), Sidney Frank Hall, 185 Meeting St

Understanding Visual Scenes
Antonio Torralba, MIT

Human visual scene understanding is remarkable: with only a brief glanceat an image, an abundance of information is available - spatialstructure, scene category and the identity of main objects in the scene.In traditional computer vision, scene and object recognition are twovisual tasks generally studied separately. However, it is unclearwhether it is possible to build robust systems for scene and objectrecognition, matching human performance, based only on localrepresentations. Another key component of machine vision algorithms isthe access to data that describe the content of images. As the fieldmoves into integrated systems that try to recognize many object classesand learn about contextual relationships between objects, the lack oflarge annotated datasets hinders the fast development of robustsolutions. In the early days, the first challenge a computer visionresearcher would encounter would be the difficult task of digitizing aphotograph. Even once a picture was in digital form, storing a largenumber of pictures (say six) consumed most of the availablecomputational resources. In addition to the algorithmic advancesrequired to solve object recognition, a key component to progress isaccess to data in order to train computational models for the differentobject classes. This situation has dramatically changed in the lastdecade, especially via the internet, which has given computer visionresearchers access to billions of images and videos. In this talk I willdescribe recent work on visual scene understanding that try to buildintegrated models for scene and object recognition, emphasizing thepower of large database of annotated images in computer vision.

Goals

- Today
  - Project – high level
  - Robust statistics and robust motion estimation
- Fri
  - Project ideas and regularization
- Monday
  - Regularization and robust statistics

Project

Greater depth on something that interests you.

Worth 2 homeworks (30% of grade).

Individual (not group) projects.

Keep it simple.

Implement something from a paper.

Extend an assignment from class.

Implement your own code. If you want to use source from web, you need to get my approval first.

Assignment for Friday:

Think about what might interest you. Come up with a concrete proposal you can describe in words.

How to get a bad grade

- Treat this as a CS32 programming project

How to get a good grade

- Think about how we formulate problems throughout this class.

1) Precise problem definition
2) Mathematical formulation
3) Assumptions and approximations
4) Optimization
5) Experiments on synthetic and real data

How to get a good grade

- Clear writeup
  - I will grade you primarily based on your writeup.
  - Clear problem definition
  - Clear mathematical formulation
  - Clear statement of assumptions/approximations
  - Good illustration of ideas with experiments
  - Analysis of failures

Multiple Motions

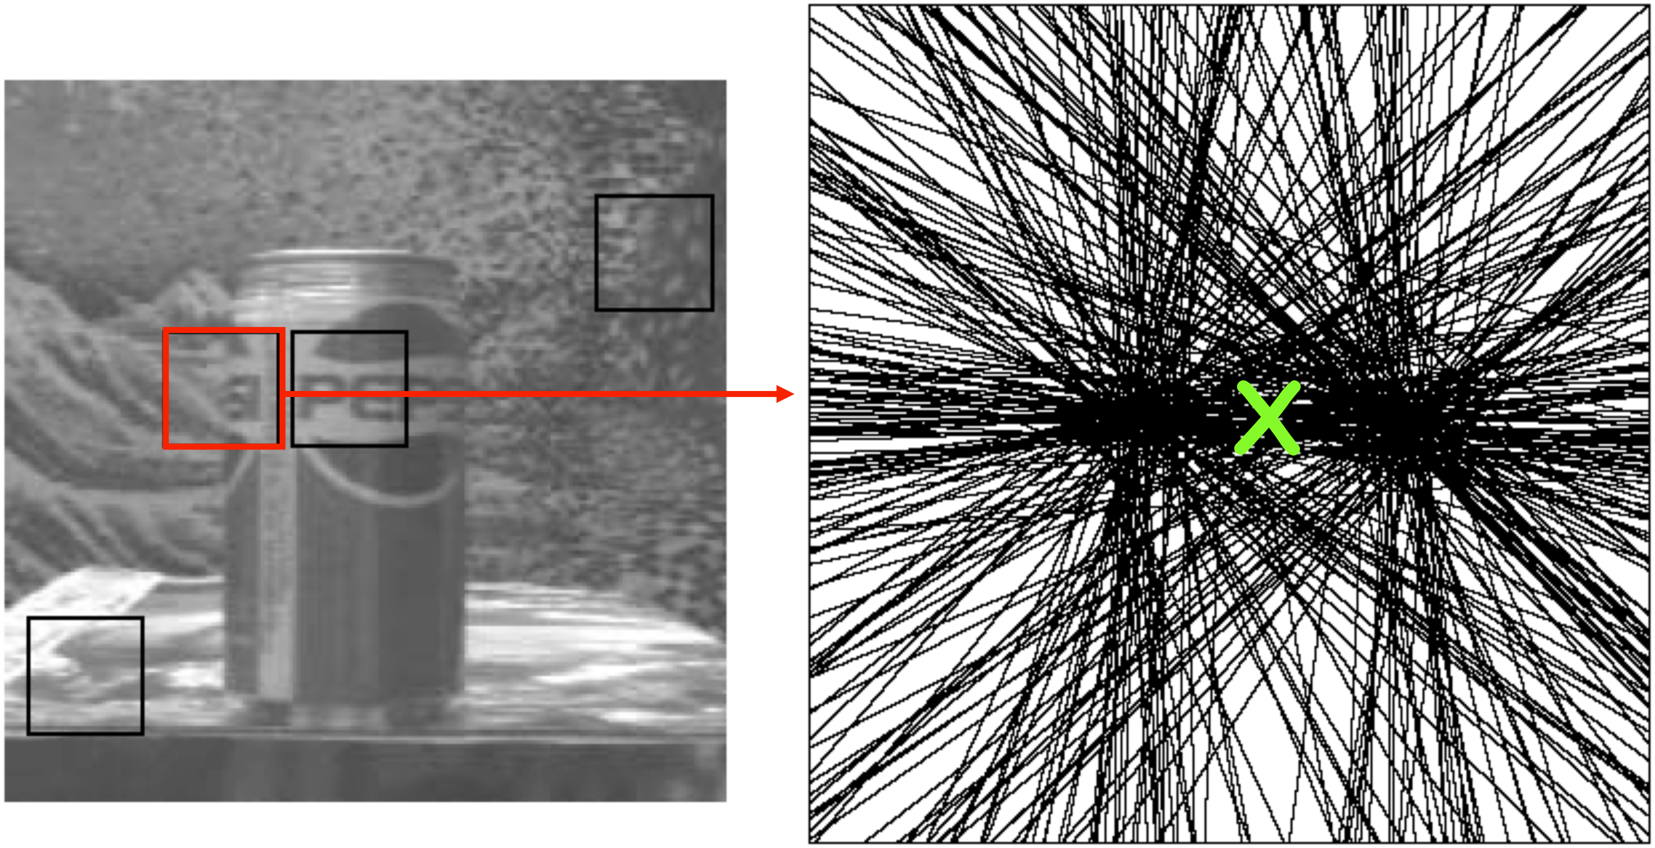

Least squares fit.

Robust Statistics

- Recover the best fit to the majority of the data.
- Detect and reject outliers.

History.

Estimating the mean

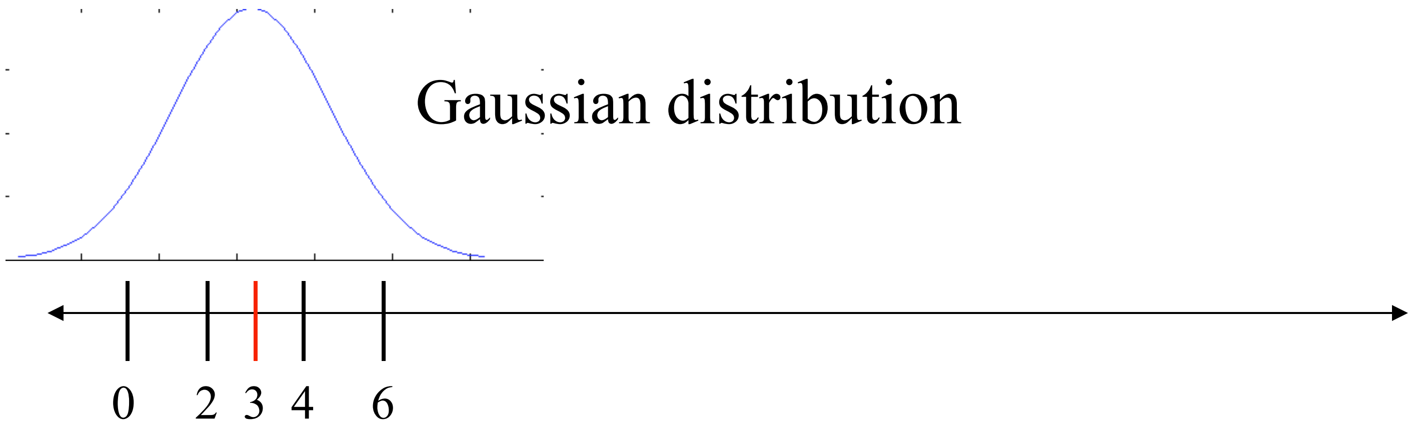

Gaussian distribution

Mean is the optimal solution to:

$$\mu = \frac{1}{N}\sum_{i=1}^{N} d_i$$

$$\min_{\mu} \sum_{i=1}^{N} \underbrace{(d_i - \mu)}_{\text{residual}}{}^2$$

Estimating the Mean

The mean maximizes this likelihood:

$$\max_{\mu} p(d_i \mid \mu) = \frac{1}{\sqrt{2\pi}\sigma} \prod_{i=1}^{N} \exp(-\frac{1}{2}(d_i - \mu)^2 / \sigma^2)$$

The negative log gives (with sigma=1):

$$\min_{\mu} \sum_{i=1}^{N} (d_i - \mu)^2$$

"least squares" estimate

Estimating the mean

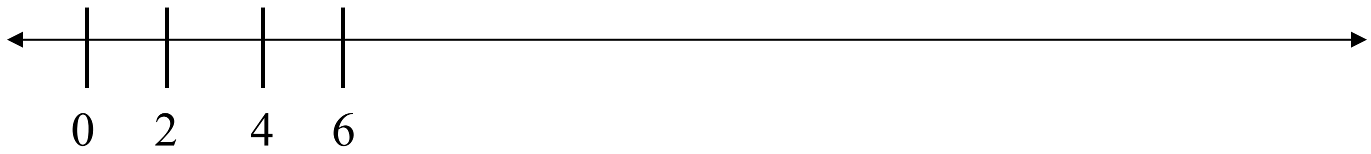

Estimating the mean

What happens if we change just one measurement?

$$\mu' = \mu + \frac{\Delta}{N}$$

With a single "bad" data point I can move the mean arbitrarily far.

Influence

Breakdown point

- * percentage of outliers required to make the solution arbitrarily bad.

Least squares:

- * influence of an outlier is linear (Δ/N)
- * breakdown point is 0% -- not robust!

What about the median?

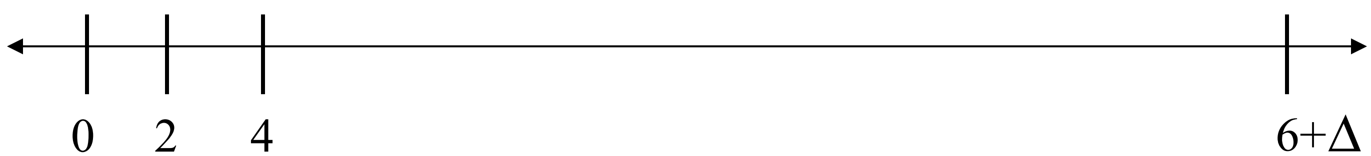

What’s Wrong?

$$\min_{\mu} \sum_{i=1}^{N} (d_i - \mu)^2$$

Outliers (large residuals) have too much influence.

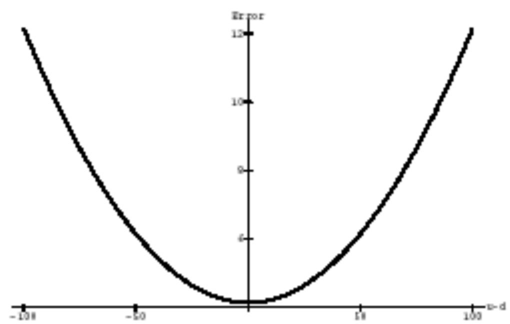

$$\rho(x) = x^2$$

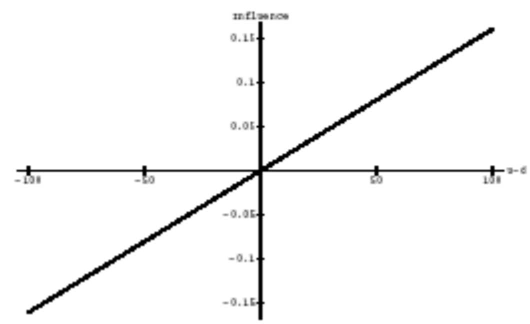

$$\psi(x) = 2x$$

Approach

Influence is proportional to the derivative of the r function.

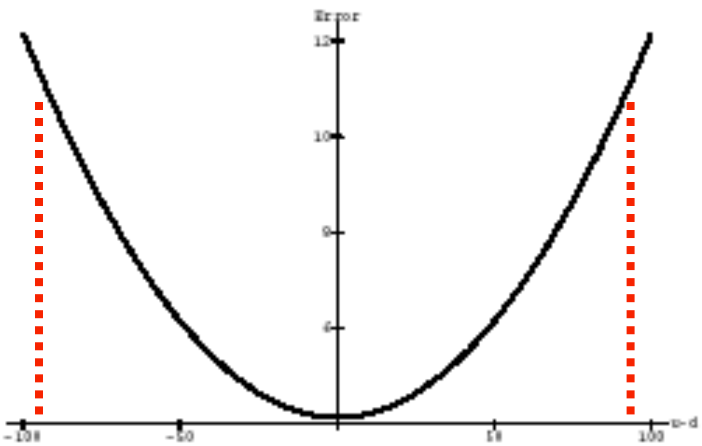

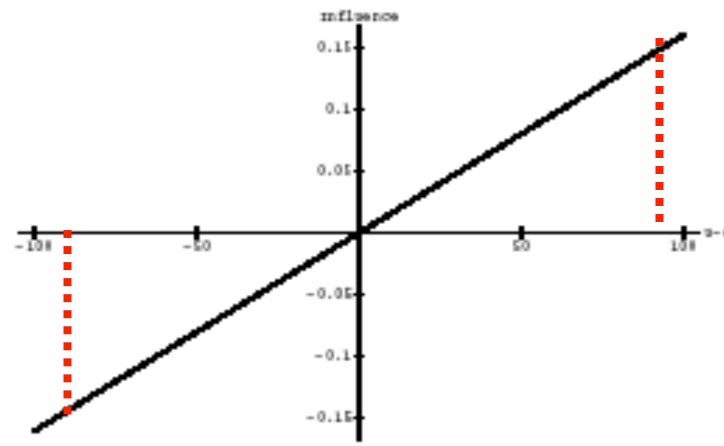

Want to give less influence to points beyond some value.

Approach

$$\min_{\mu} \sum_{i=1}^{N} \rho(d_i - \mu, \sigma)$$

Robust error function

Scale parameter

Replace

$$\rho(x,\sigma) = \left(\frac{x}{\sigma}\right)^2$$

with something that gives less influence to outliers.

L1 Norm

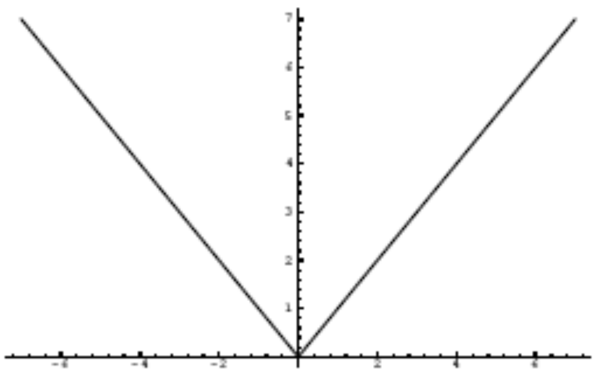

$$\rho(x) = |x|$$

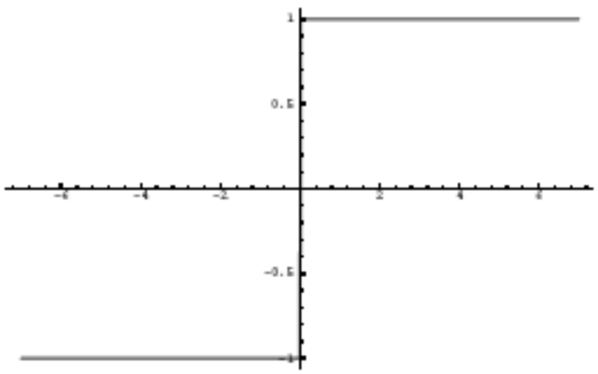

$$\psi(x) = \text{sign}(x)$$

L1 Norm

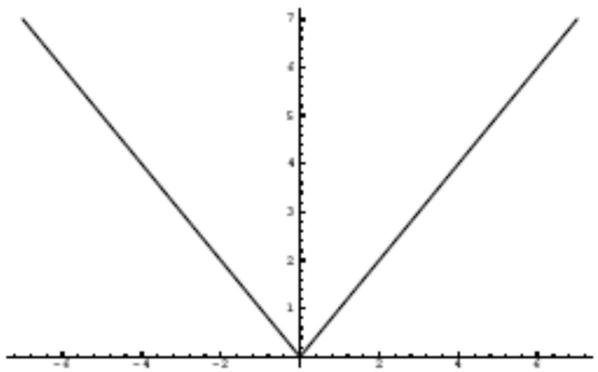

$$\rho(x) = |x|$$

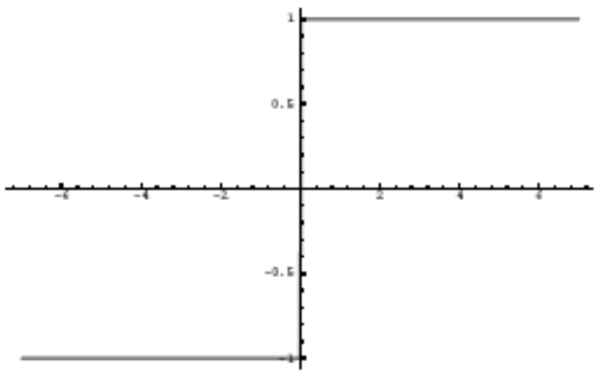

$$\psi(x) = \text{sign}(x)$$

Estimating the median

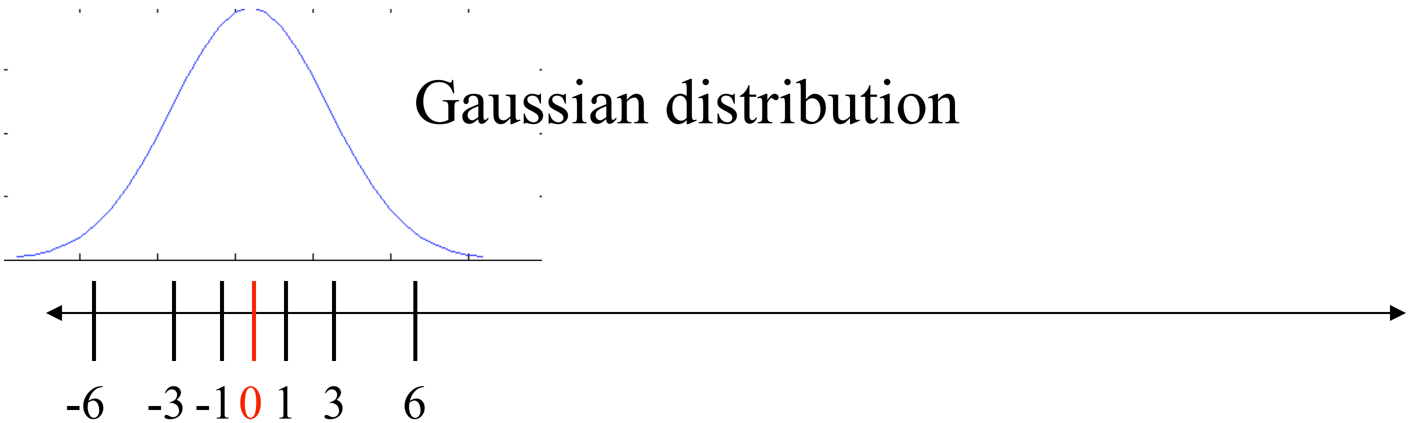

0: 6+3+1+0+1+3+6=20

$$\min_{\mu} \sum_{i=1}^{N} | \underbrace{d_i - \mu}_{\text{residual}} |$$

Estimating the median

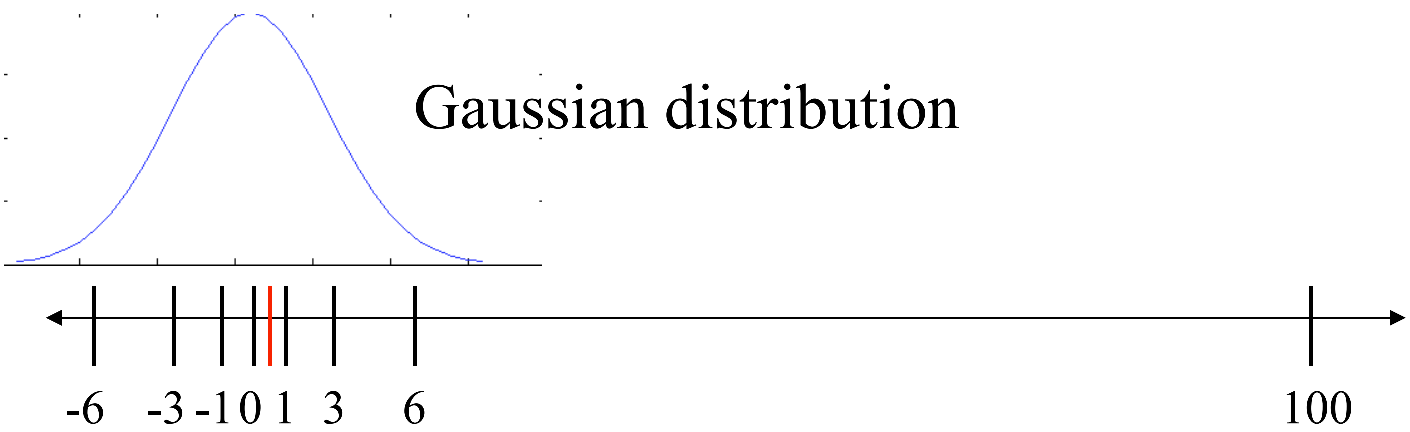

0: 6+3+1+0+1+3+6+100=120

1: 7+4+2+1+0+2+5+99=120

$$\min_{\mu} \sum_{i=1}^{N} | \underbrace{d_i - \mu}_{\text{residual}} |$$

Estimating the median

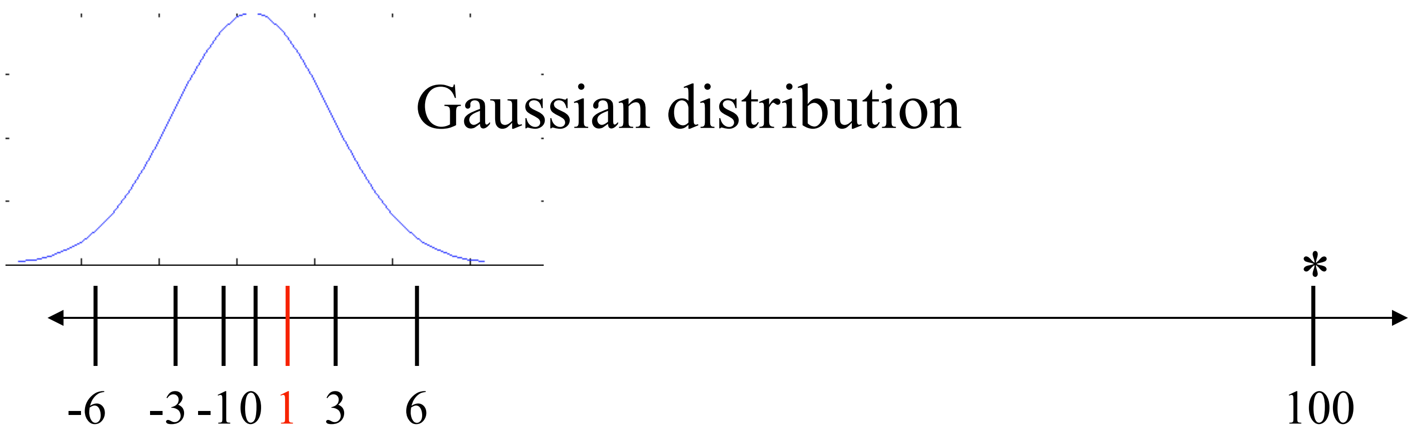

0: 6+3+1+0+1+3+6+100=120+100=220

1: 7+4+2+1+0+2+5+99=120+99=219

$$\min_{\mu} \sum_{i=1}^{N} |\underbrace{d_i - \mu}_{\text{residual}}|$$

Estimating the median

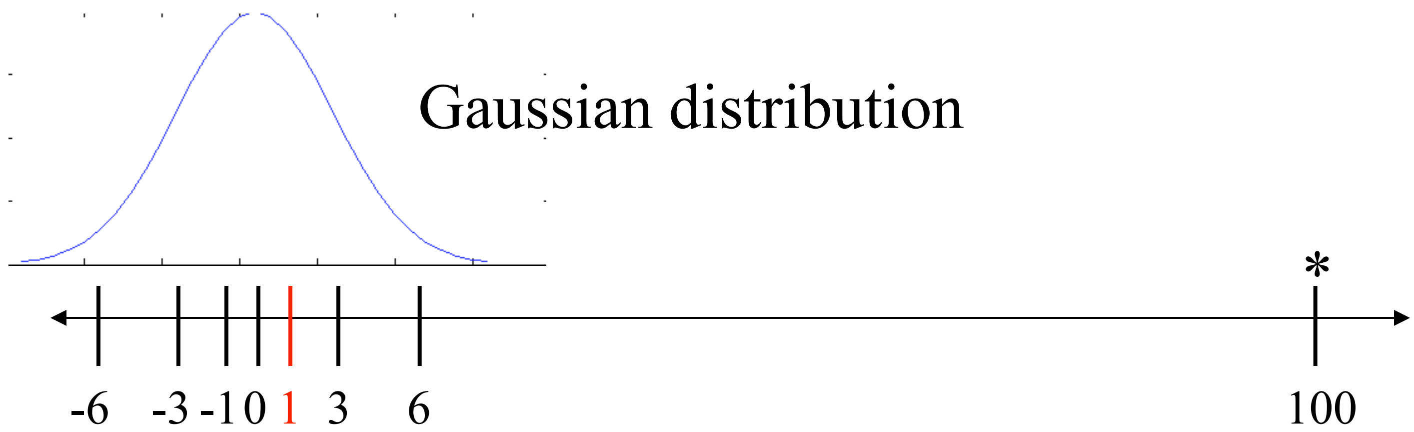

0: 6+3+1+0+1+3+6+100=120+100=220

1: 7+4+2+1+0+2+5+99=120+99=219

2: 8+5+3+2+1+1+4+98=122+98=220

Estimating the median

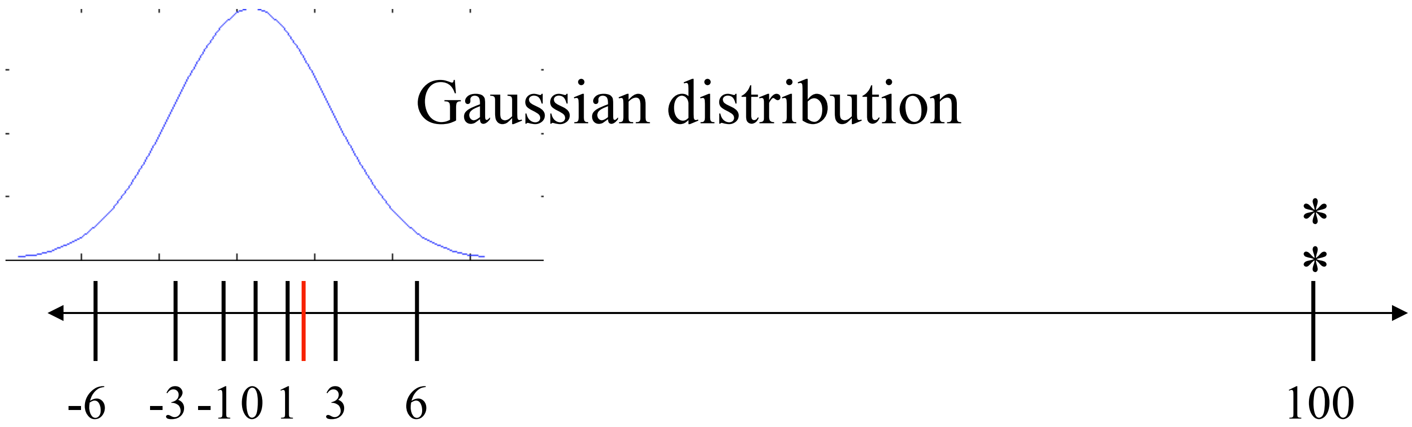

0: 6+3+1+0+1+3+6+100=120+100=220+100=320

1: 7+4+2+1+0+2+5+99=120+99=219+99=318

2: 8+5+3+2+1+1+4+98=122+98=220+98=318

Estimating the median

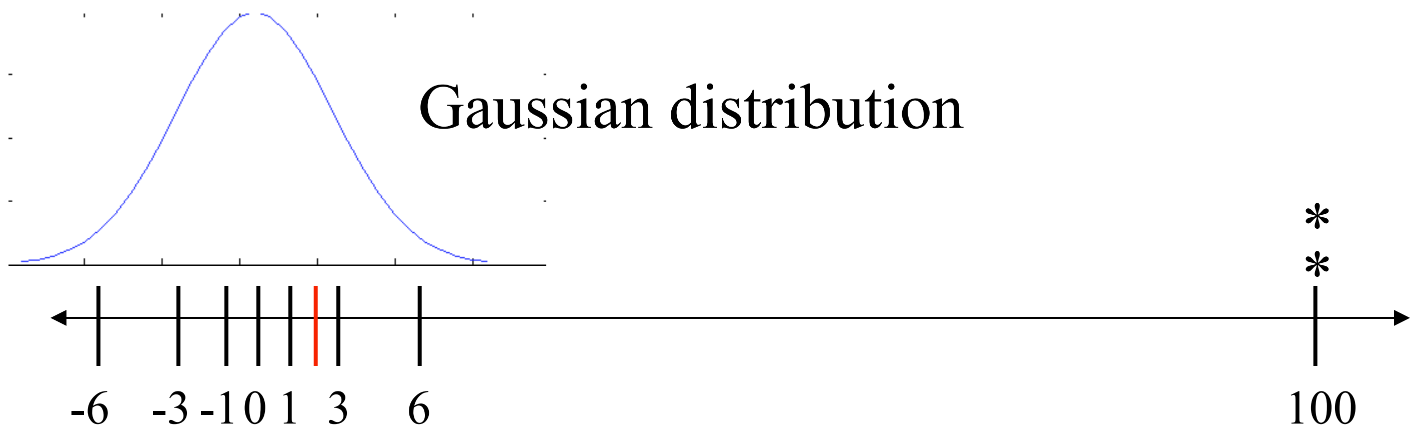

0: 6+3+1+0+1+3+6+100=120+100=220+100=320

1: 7+4+2+1+0+2+5+99=120+99=219+99=318

2: 8+5+3+2+1+1+4+98=122+98=220+98=318

3: 9+6+4+3+2+0+3+97=124+97=221+97=318

Estimating the median

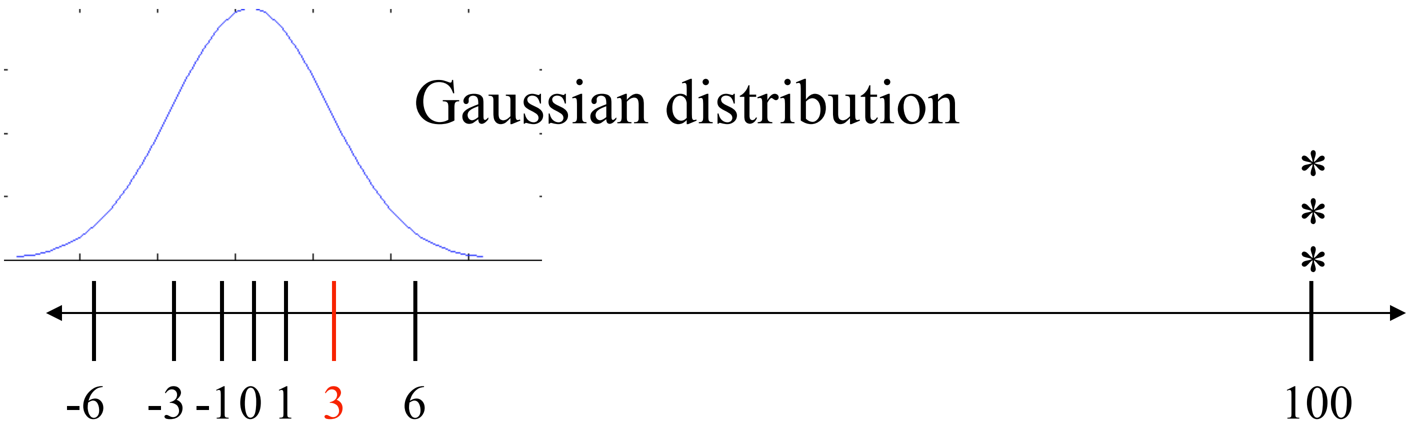

0: 6+3+1+0+1+3+6+100=120+100=220+100=320+100=420

1: 7+4+2+1+0+2+5+99=120+99=219+99=318+99=417

2: 8+5+3+2+1+1+4+98=122+98=220+98=318+98=416

3: 9+6+4+3+2+0+3+97=124+97=221+97=318+97=415

Estimating the median

7 inliers out of 13 data points

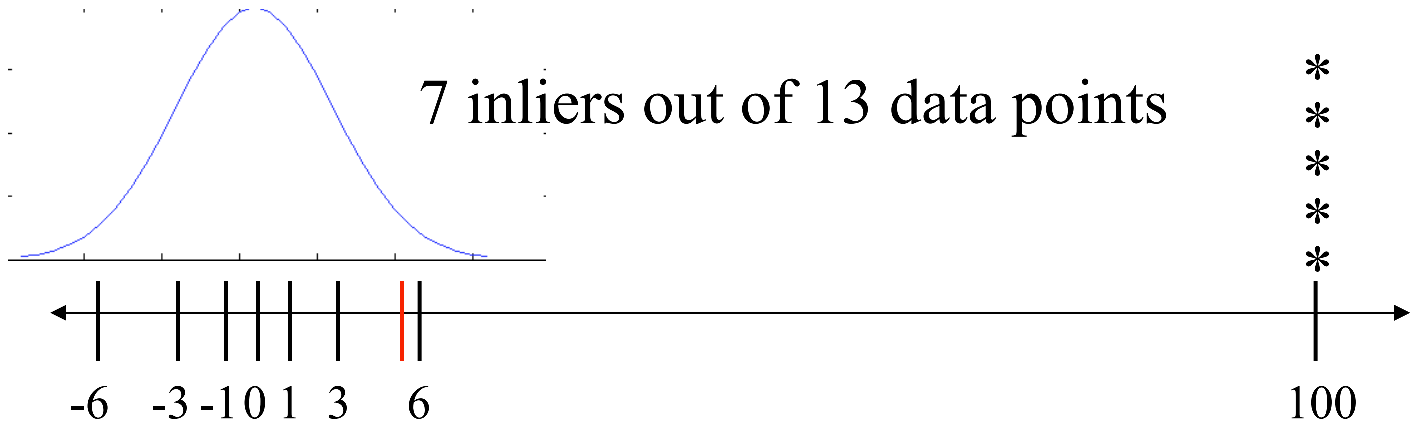

3: 9+6+4+3+2+0+3=27+(7*97)=706

4: 10+7+5+4+3+1+2=32+(7*96)=704

5: 11+8+6+5+4+2+1=37+(7*95)=702

6: 12+9+7+6+5+3+0=42+(7*95)=707

Estimating the median

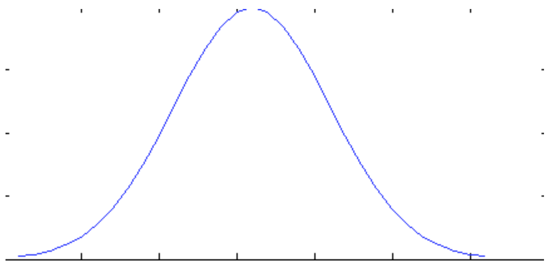

L1 can tolerate up to 50% outliers

Influence of each outlier is the same.
1000 is the same as 100

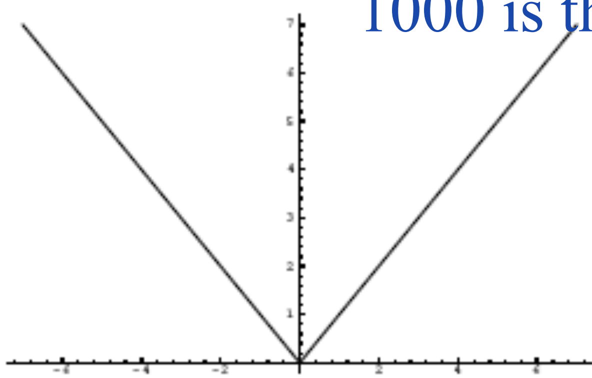

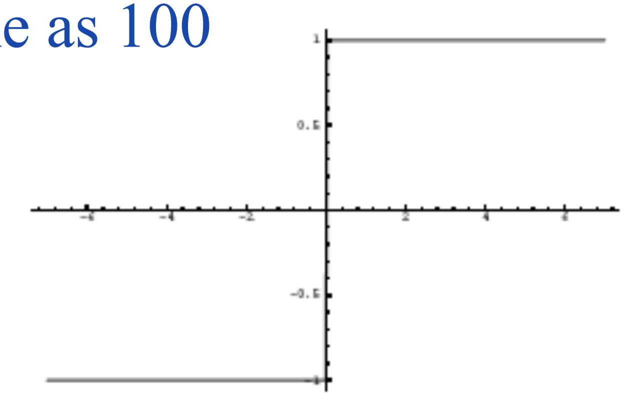

The problems get worse as the number of parameters increases.
For linear regression problems, L1 can tolerate approx 25% outliers.

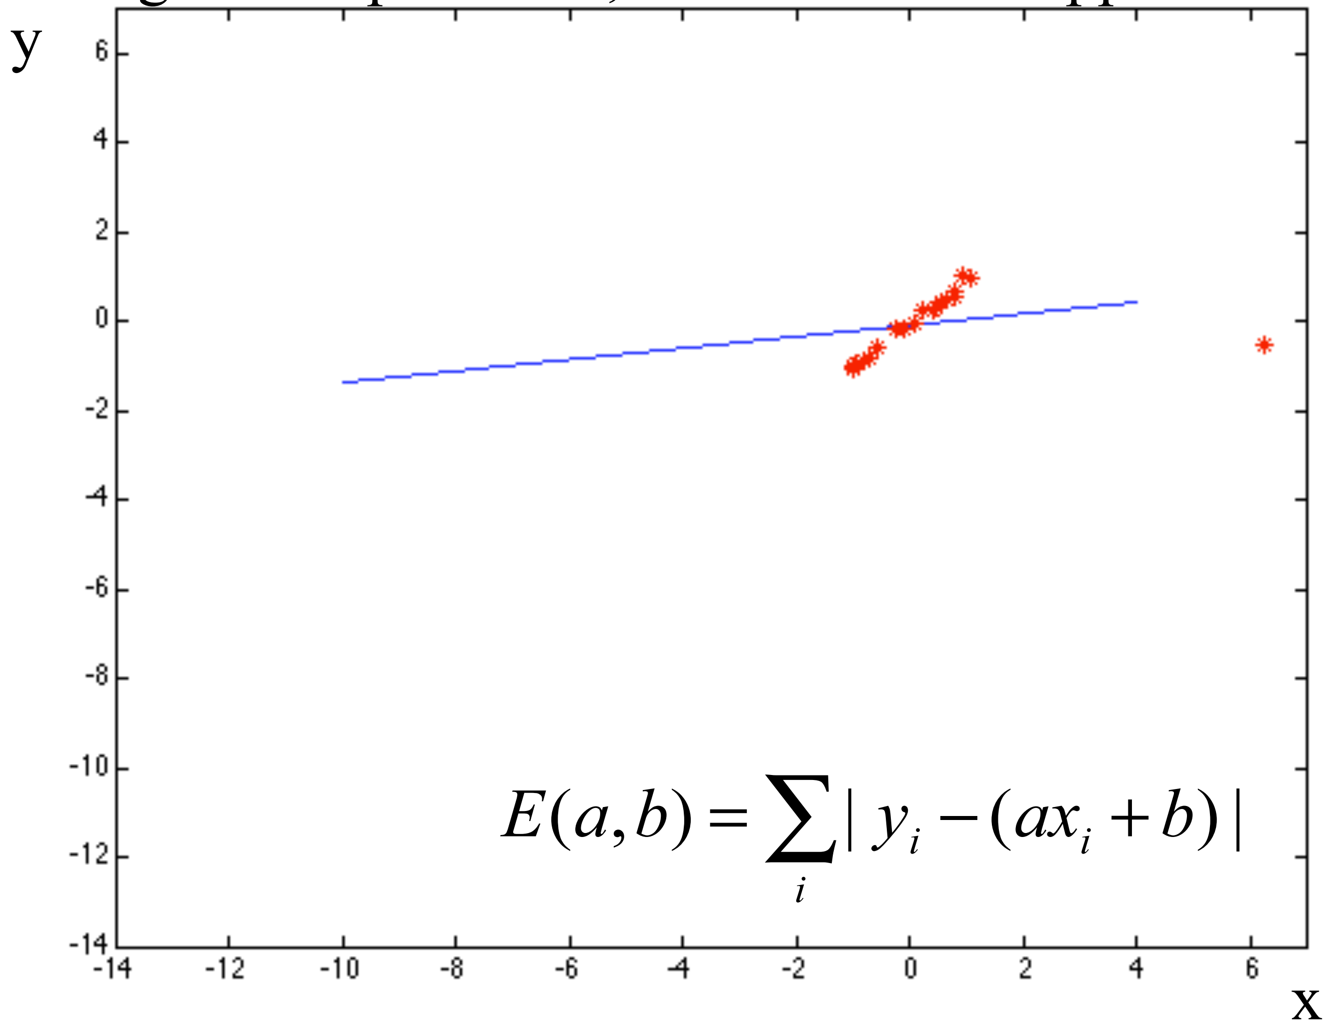

$$E(a,b) = \sum_i | y_i - (ax_i + b) |$$

What’s Wrong?

$$\min_{\mu} \sum_{i=1}^{N} (d_i - \mu)^2$$

Outliers (large residuals) have too much influence.

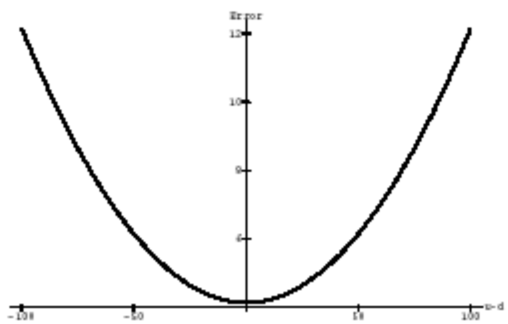

$$\rho(x) = x^2$$

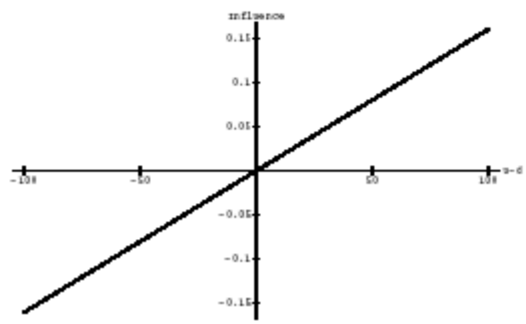

$$\psi(x) = 2x$$

Approach

$$\min_{\mu} \sum_{i=1}^{N} \rho(d_i - \mu, \sigma)$$

Robust error function

Scale parameter

Replace

$$\rho(x, \sigma) = \left(\frac{x}{\sigma} \right)^2$$

with something that gives less influence to outliers.

Redescending Function

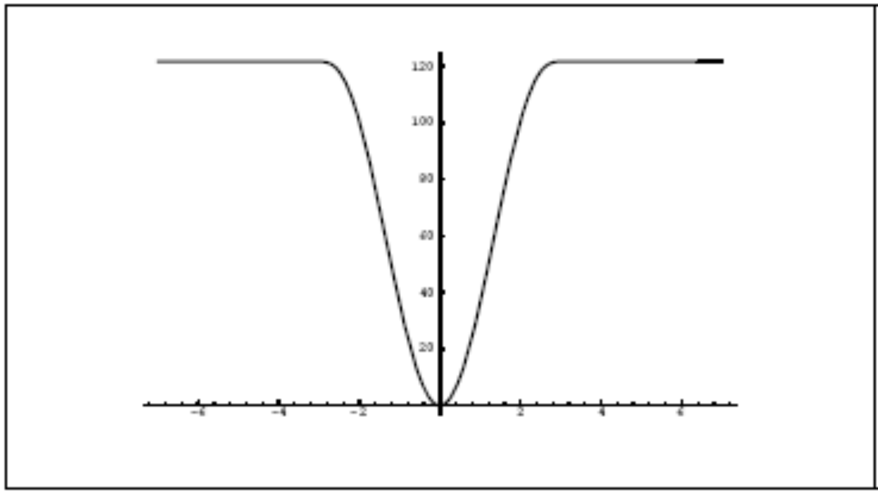

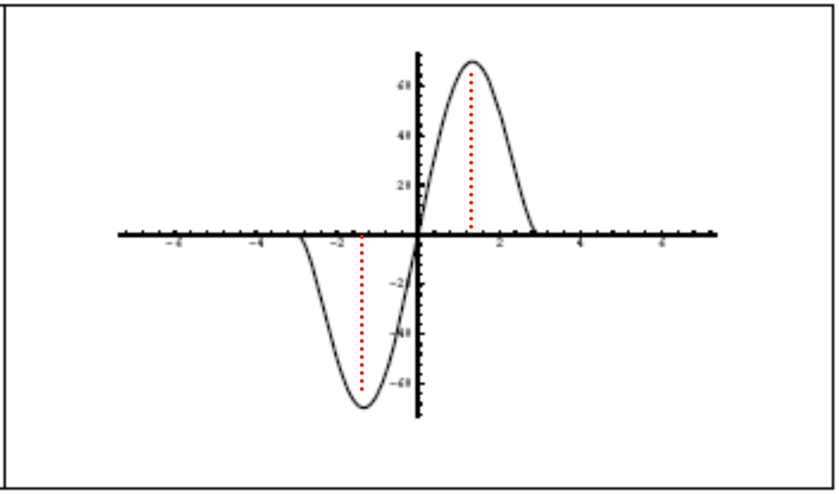

Tukey's biweight.

Beyond a point, the influence begins to decrease.

Beyond where the second derivative is zero – outlier points

Dectecting “Outliers”

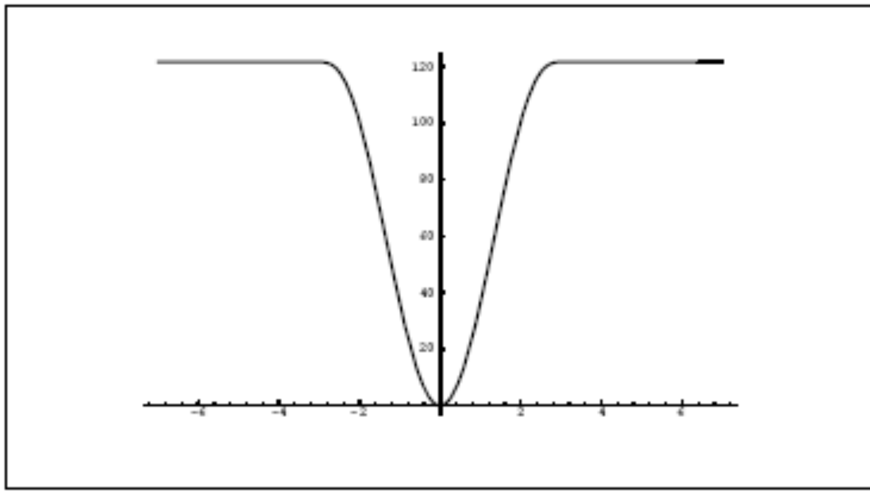

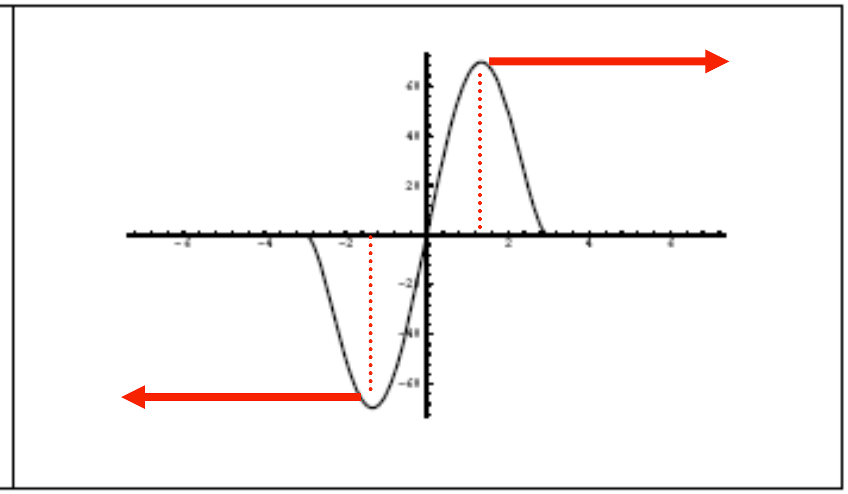

Tukey’s biweight.

Beyond a point, the influence begins to decrease.

Beyond where the second derivative is zero – outlier points

Problem

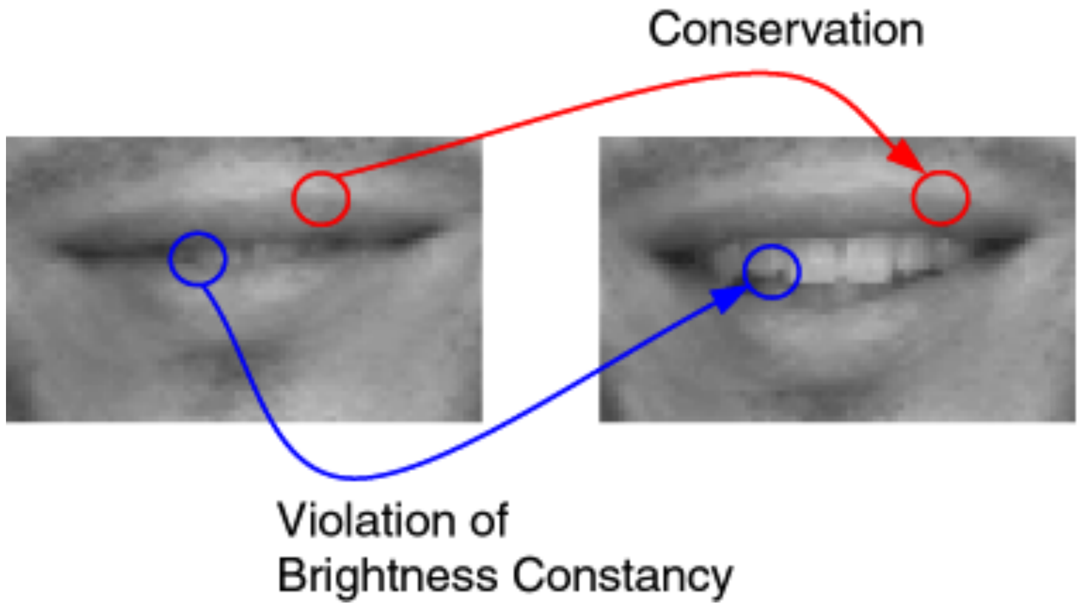

Violations of brightness constancy result in large residuals:

$$\left| I_x u(\mathbf{x};\mathbf{a}) + I_y v(\mathbf{x};\mathbf{a}) + I_t \right|$$

* chose r to be insensitive to outliers

Robust Estimation

$$E(\mathbf{a}) = \sum_{x,y \in R} \rho(I_x u + I_y v + I_t, \sigma)$$

How do we minimize this?

Robust Estimation

$$E(\mathbf{a}) = \sum_{x,y \in R} \rho(I_x u + I_y v + I_t, \sigma)$$

Minimize: differentiate and set equal to zero:

$$\frac{\partial E}{\partial u} = \sum_{x,y \in R} \psi(I_x u + I_y v + I_t, \sigma) I_x = 0$$

$$\frac{\partial E}{\partial v} = \sum_{x,y \in R} \psi(I_x u + I_y v + I_t, \sigma) I_y = 0$$

No closed form solution!

Gradient Descent

Minimize $f(x)$

The root x* corresponds to the point

$$f'(x^*) = 0$$

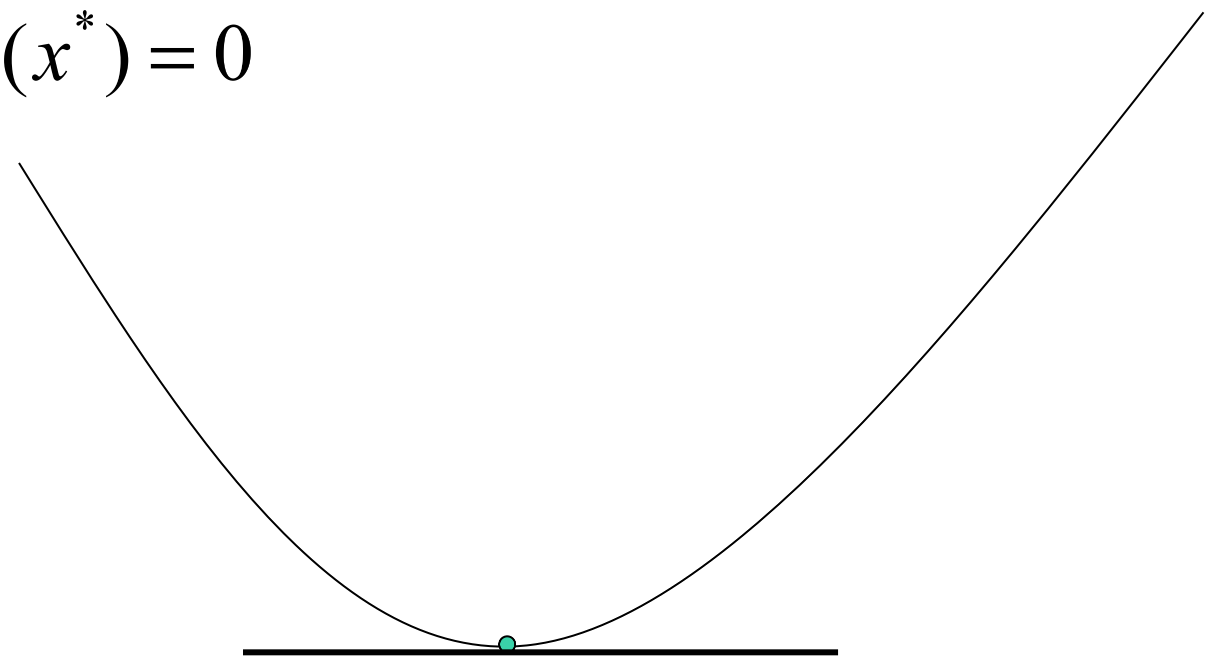

Gradient Descent

Minimize

$f(x)$

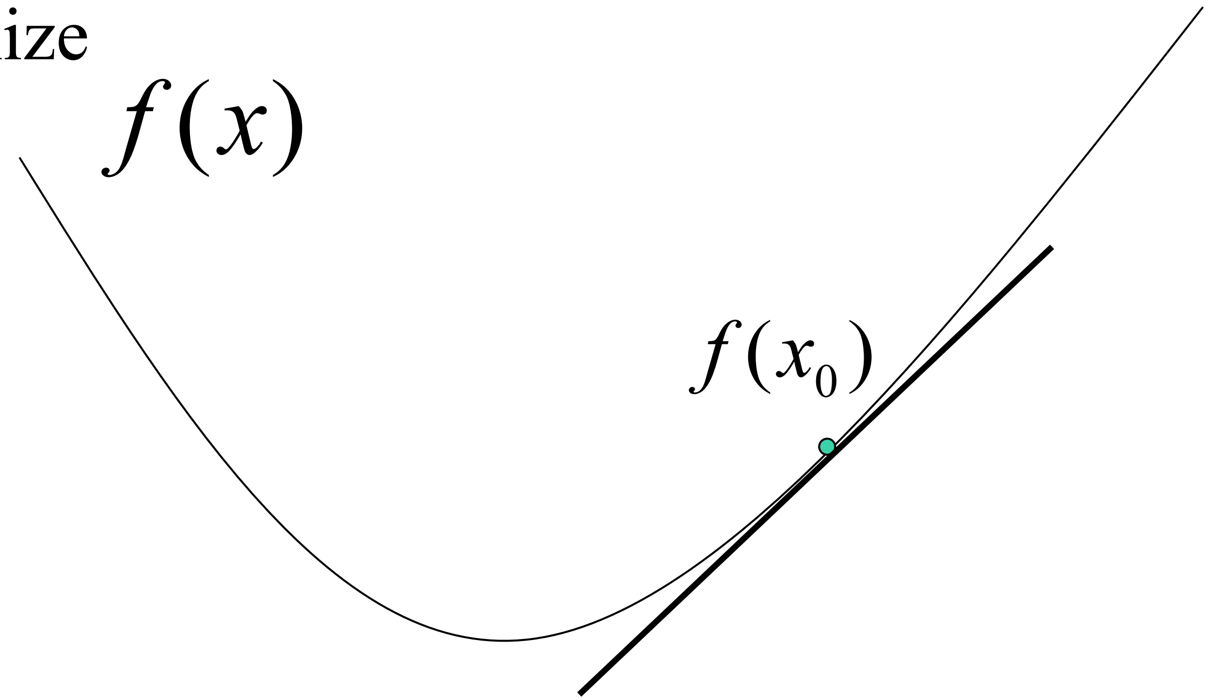

Coordinate Descent

Minimize

$f(x)$

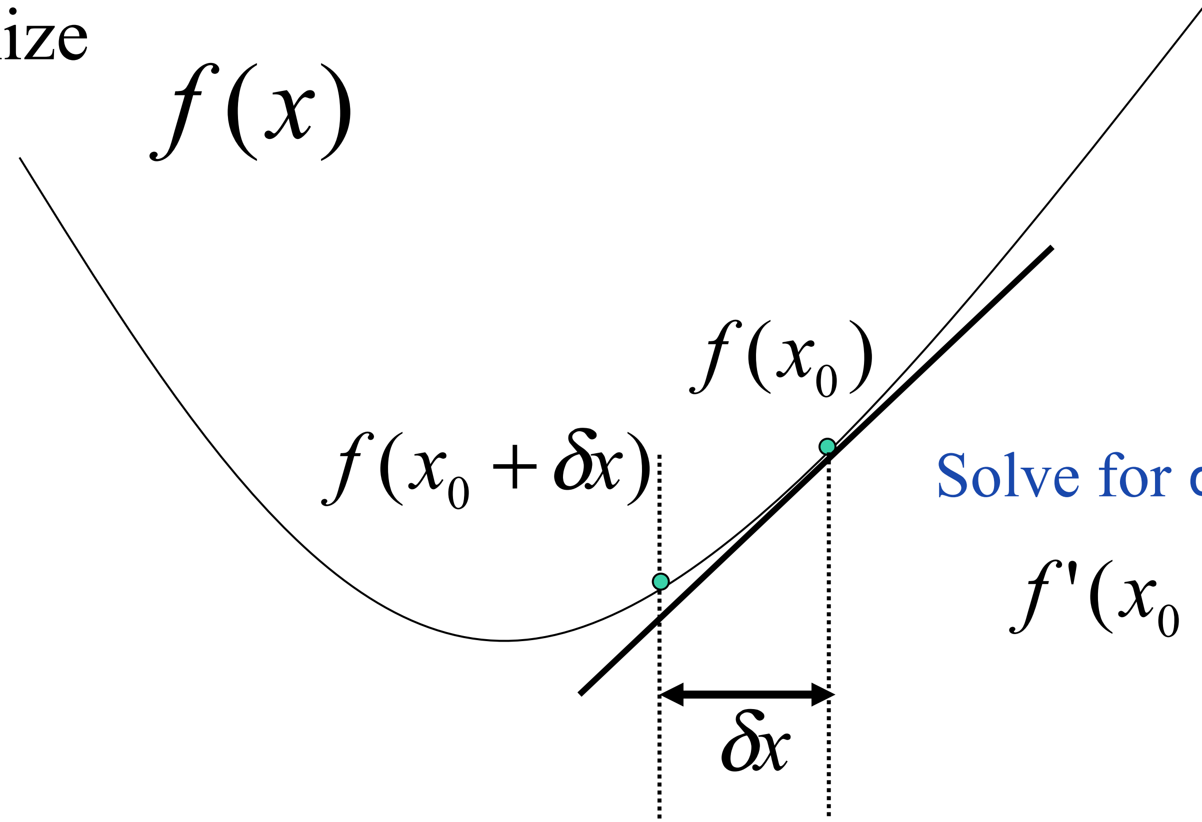

Solve for dx

$$f'(x_0 + \delta x) = 0$$

Coordinate Descent

Minimize

$$f(x)$$

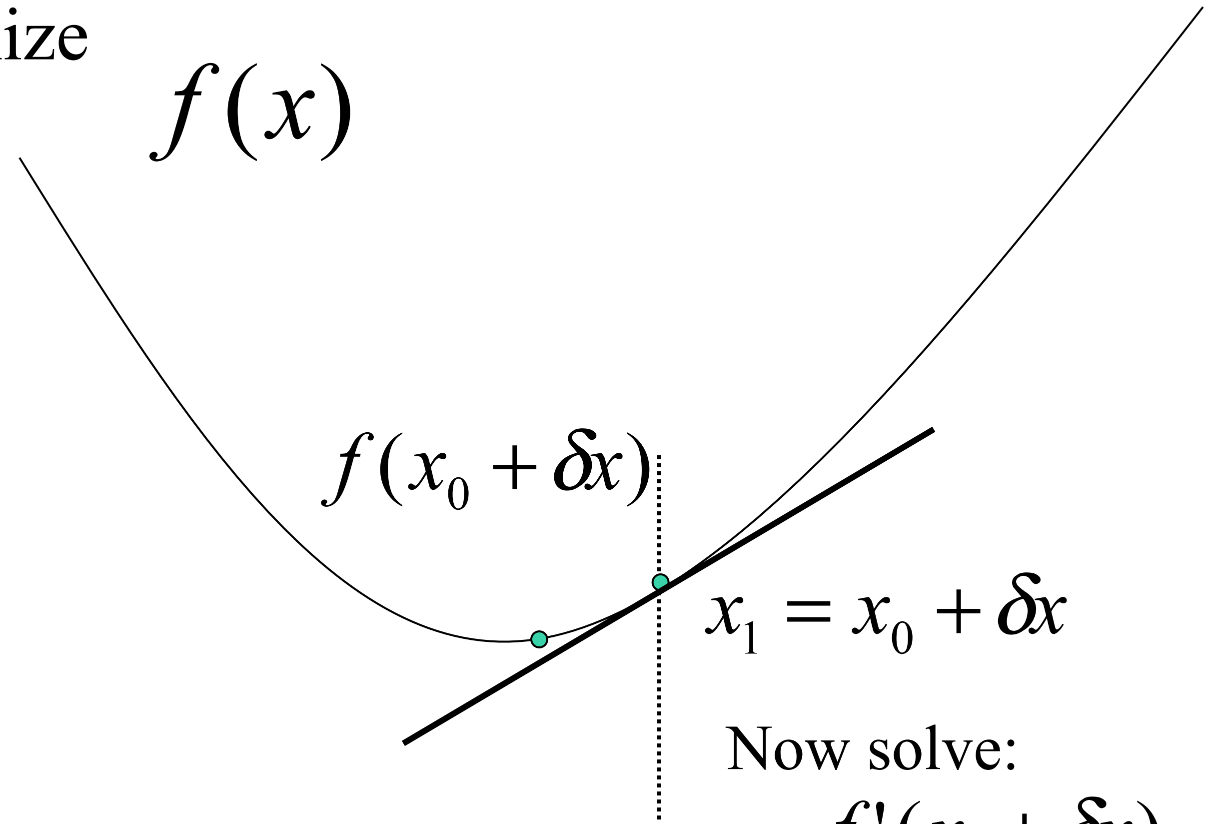

Now solve:

$$f'(x_1 + \delta x) = 0$$

Newton's Method

$$f'(x+\delta x) \approx f'(x) + f''(x)\delta x = 0$$

↓

$$\delta x = -f'(x) / f''(x)$$

↓

$$x^{(n+1)} = x^{(n)} + \delta x$$

Incrementally update the estimate of x

In multiple dimensions the second derivative corresponds to the Hessian matrix – we'll just do coordinate descent.