Introduction to Computer Vision

Michael J. Black

Nov 2009

Non-linear optimization and regularization

November 11, 200912-1pm, 182 George St, Rm 110

Generative Models for Image Analysis

Lo-Bin Chang

Division of Applied Mathematics Brown University

A probabilistic grammar for the grouping and labeling of parts and objects, when taken together with pose and part-dependent appearance models, constitutes a generative scene model and a Bayesian framework for image analysis. To the extent that the generative model generates features, as opposed to pixel intensities, the posterior distribution (i.e. the conditional distribution on part and object labels given the image) is based on incomplete information; feature vectors are generally insufficient to recover the original intensities. I will propose a way to learn pixel-level models for the appearances of parts. I will demonstrate the utility of the models with some experiments in Bayesian image classification.

Goals

- Today
  - Finish project ideas
  - Finish robust stats and non-linear optimization
  - Start regularization
- Wednesday
  - Finish regularization and dense flow
  - Start tracking

Assignments

- Assignment 4 out today.
  - Problem 1 due next Monday (pretty easy) and you have everything you need.
  - Problem 2 due a week later. Tracking.
- Project proposal due Friday

Problem

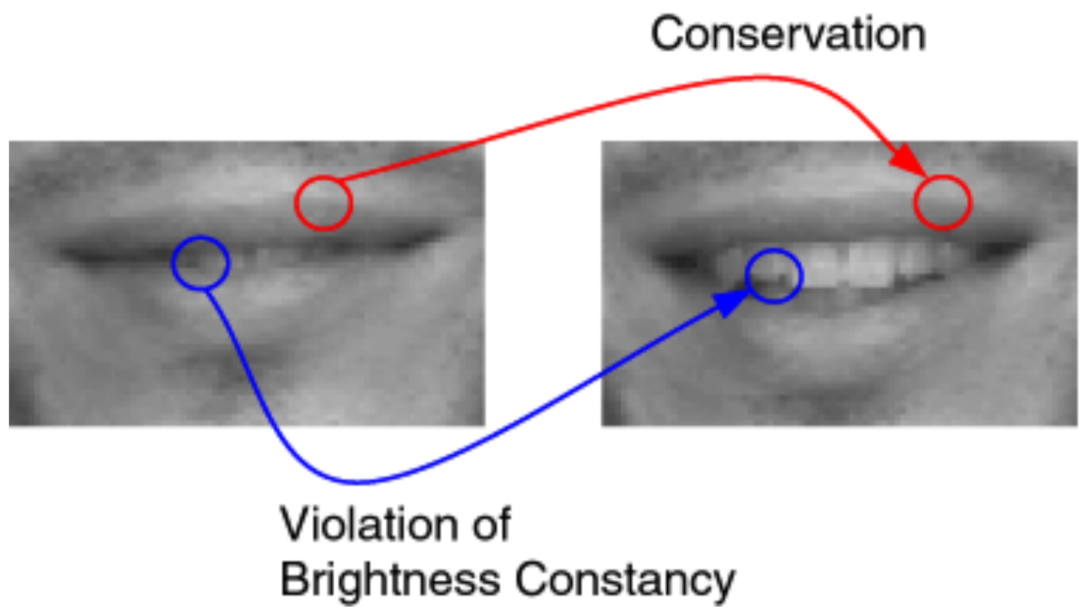

Violations of brightness constancy result in large residuals:

$$\left| \mathrm{I}_x u(\mathbf{x};\mathbf{a}) + \mathrm{I}_y v(\mathbf{x};\mathbf{a}) + \mathrm{I}_\mathrm{t} \right|$$

* Chose ρ 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

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

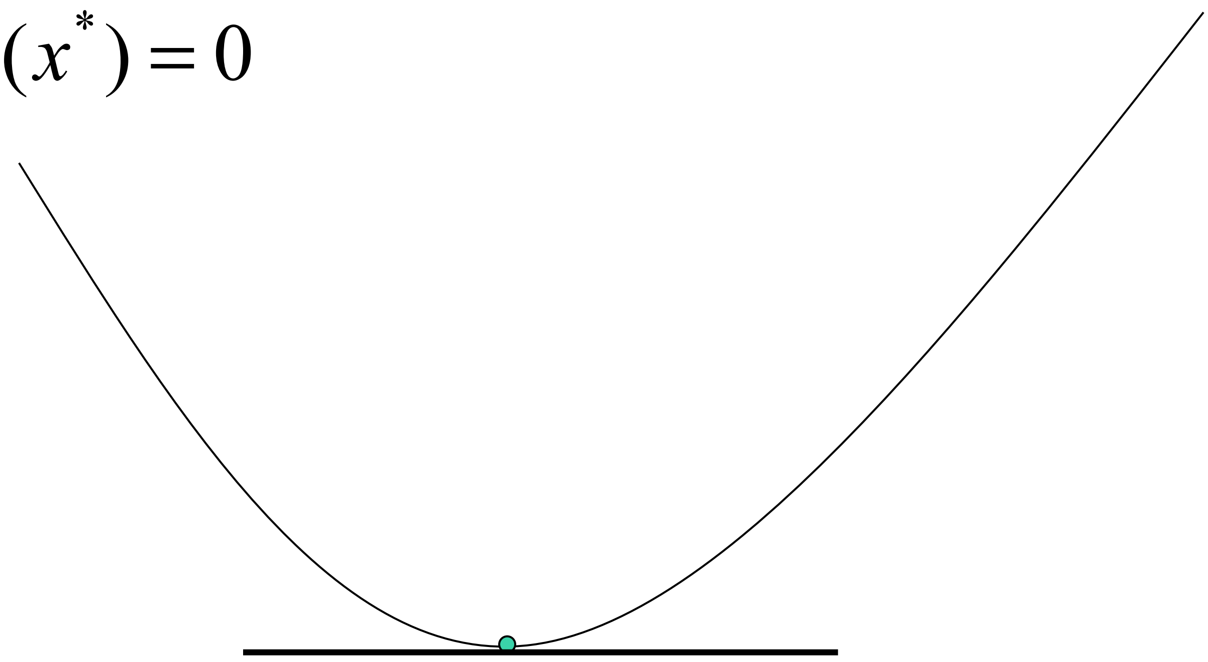

Gradient Descent

Minimize

$f(x)$

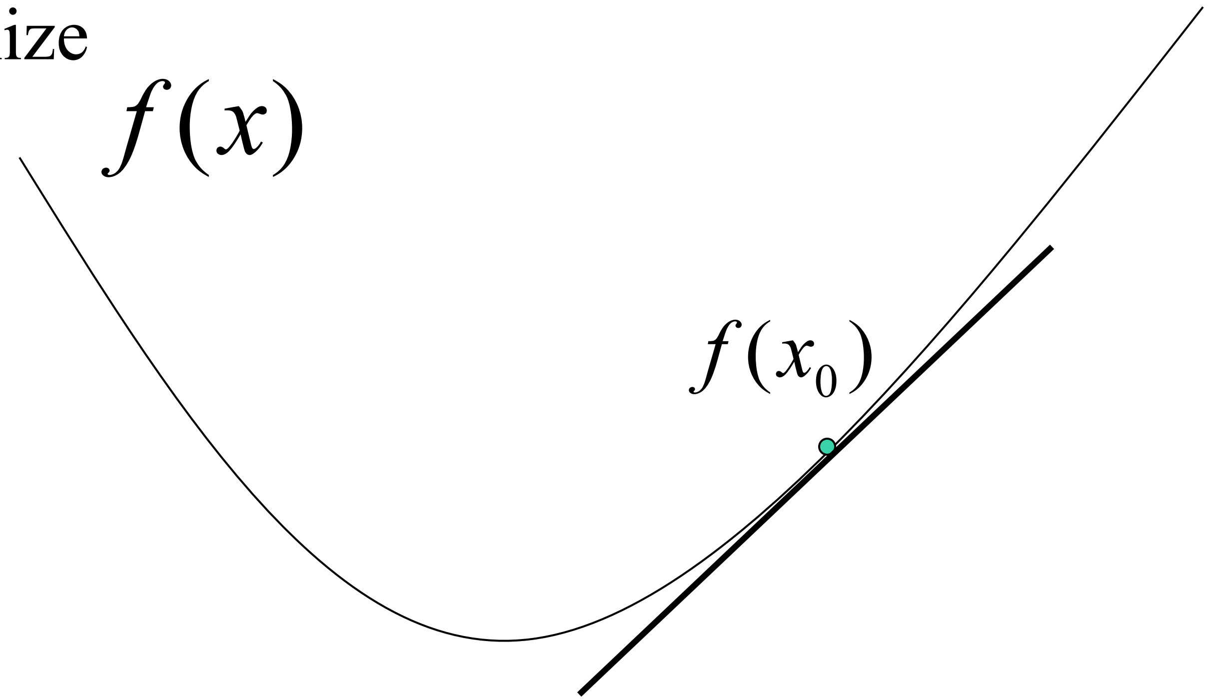

Coordinate Descent

Minimize

$$f(x)$$

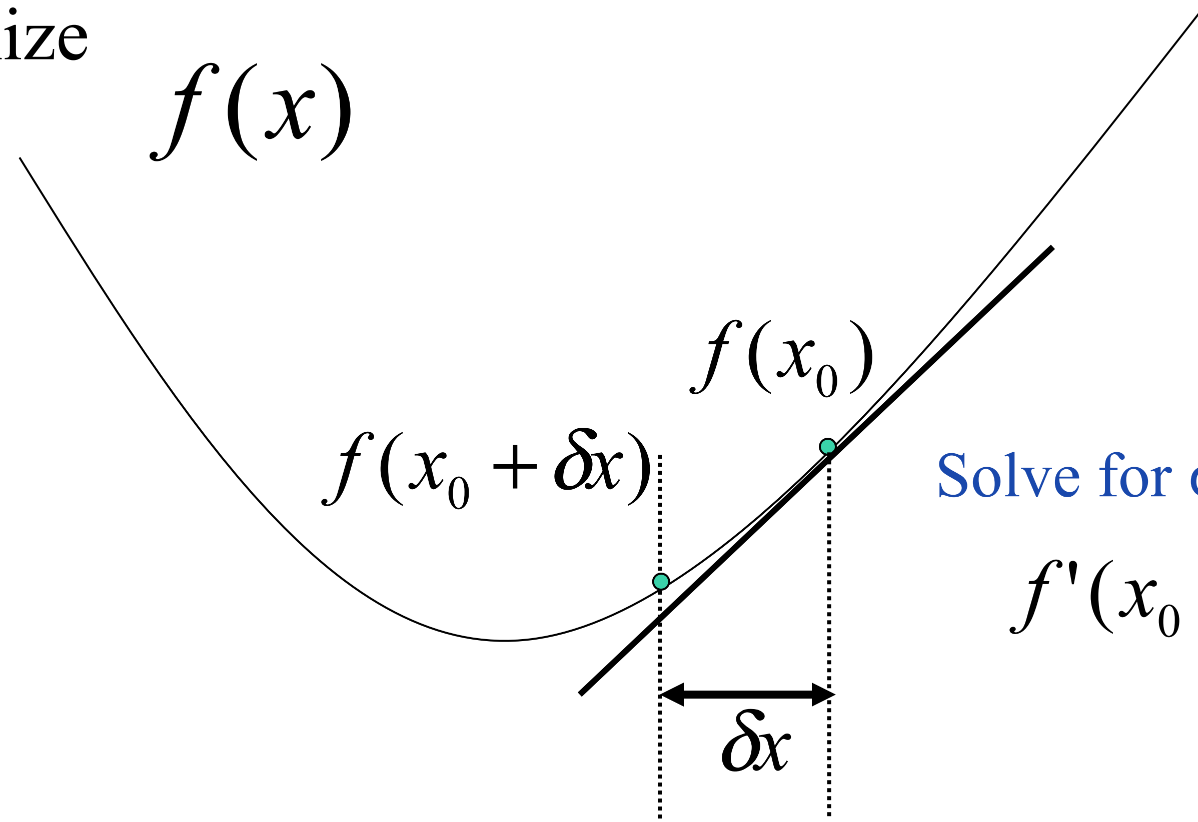

Solve for dx

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

Coordinate Descent

Minimize

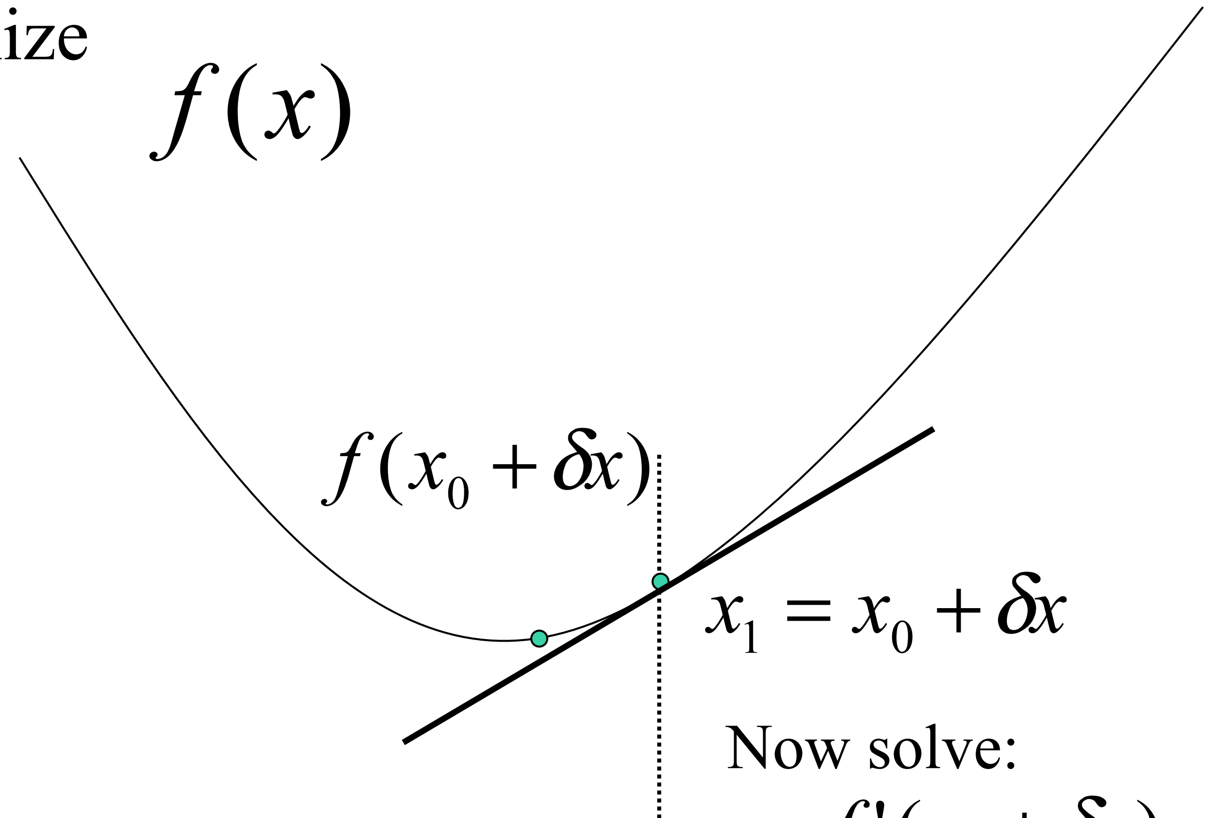

Now solve:

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

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.

Robust motion estimation

$$\delta u = w \sum_{x,y\in R} \psi(I_x u + I_y v + I_t, \sigma) I_x \Big/ \sum_{x,y\in R} \psi'(I_x u + I_y v + I_t, \sigma) I_x^2$$

$$\delta v = w \sum_{x,y\in R} \psi(I_x u + I_y v + I_t, \sigma) I_y \Big/ \sum_{x,y\in R} \psi'(I_x u + I_y v + I_t, \sigma) I_y^2$$

$$u_{n+1} = u_n + \delta u$$

$$v_{n+1} = v_n + \delta v$$

I take the max of the second derivative.
Guaranteed to converge.
But not guaranteed to find a global minimum!

Estimating Flow

Minimize:

$$E(\mathbf{a}) = \sum_{\mathbf{x} \in R} \rho(\mathrm{I}_x u(\mathbf{x};\mathbf{a}) + \mathrm{I}_u v(\mathbf{x};\mathbf{a}) + \mathrm{I}_t, \sigma)$$

Parameterized models provide strong constraints:

* Hundreds, or thousands, of constraints.
* Handful (e.g. six) unknowns.

Can be very accurate (when the model is good)!

Robust Estimation

Geman-McClure function works well.
Twice differentiable, redescending.

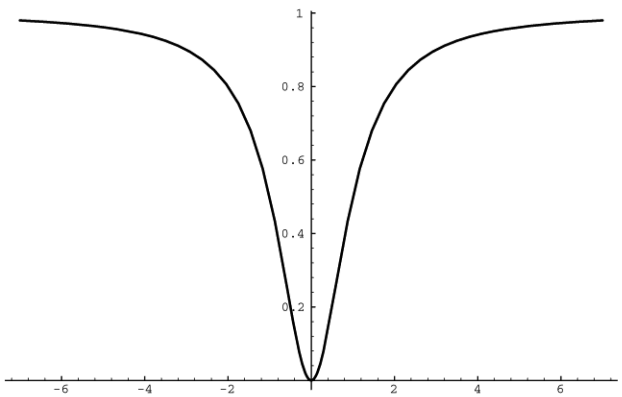

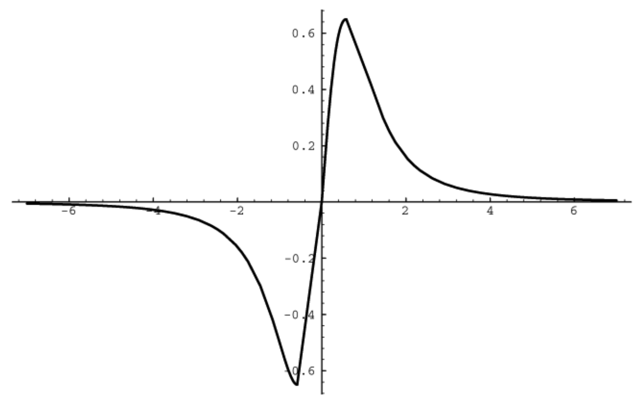

$$\rho(r,\sigma) = \frac{r^2}{\sigma^2 + r^2} \qquad \psi(r,\sigma) = \frac{2r\sigma^2}{(\sigma^2 + r^2)^2}$$

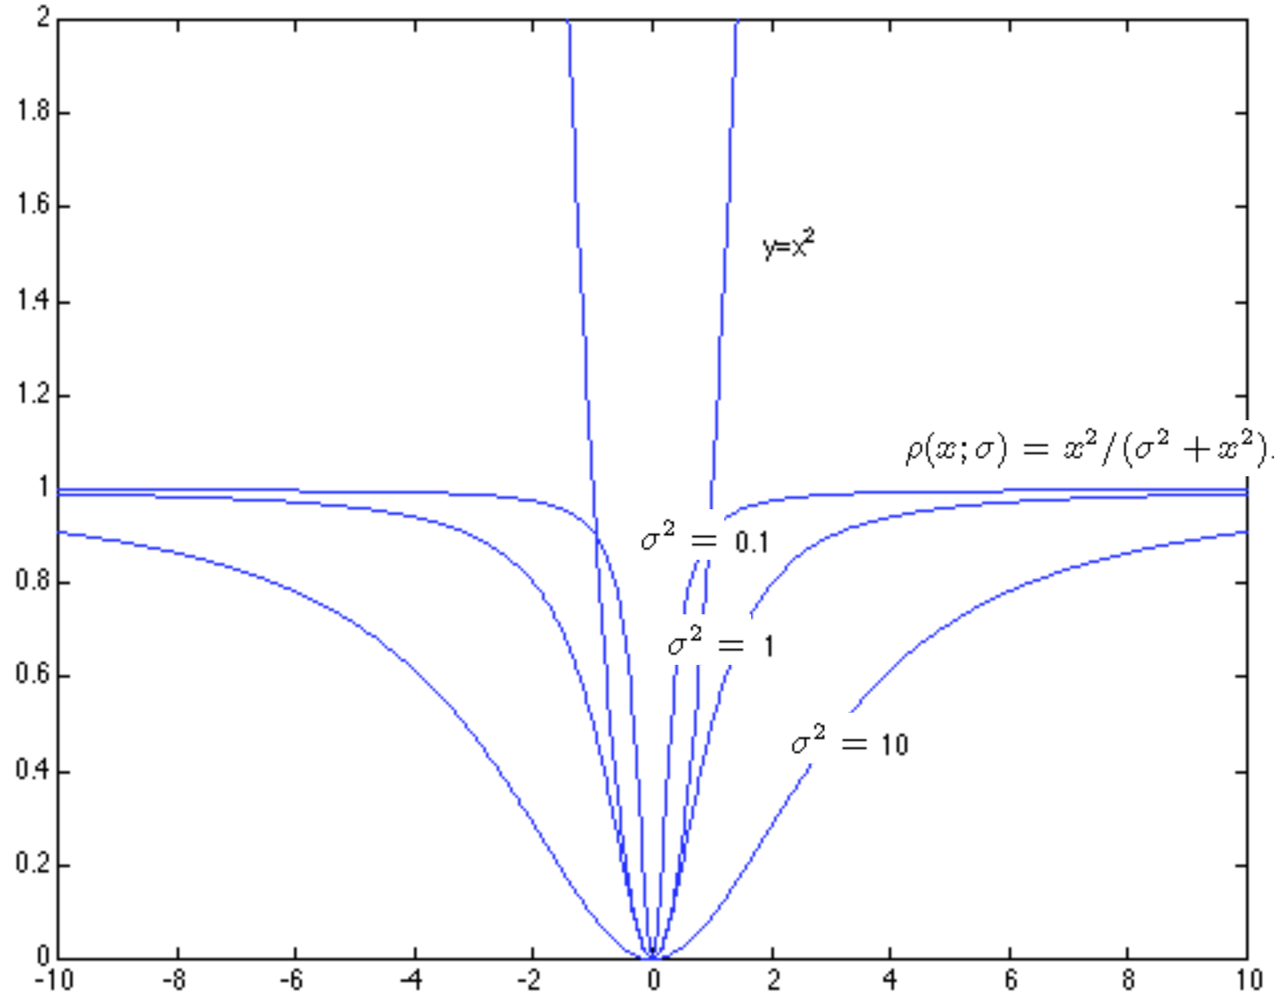
2
1.8
1.6
1.4
1.2
1
0.8
0.6
0.4
0.2
0
-10
-8
-6
-4
-2
0
2
4
6
8
10
y=x2
$\rho(x;\sigma) = x^2/(\sigma^2 + x^2)$.
$\sigma^2 = 0.1$
$\sigma^2 = 1$
$\sigma^2 = 10$

Deterministic Annealing

Start with a "quadratic" optimization problem and gradually reduce outliers.

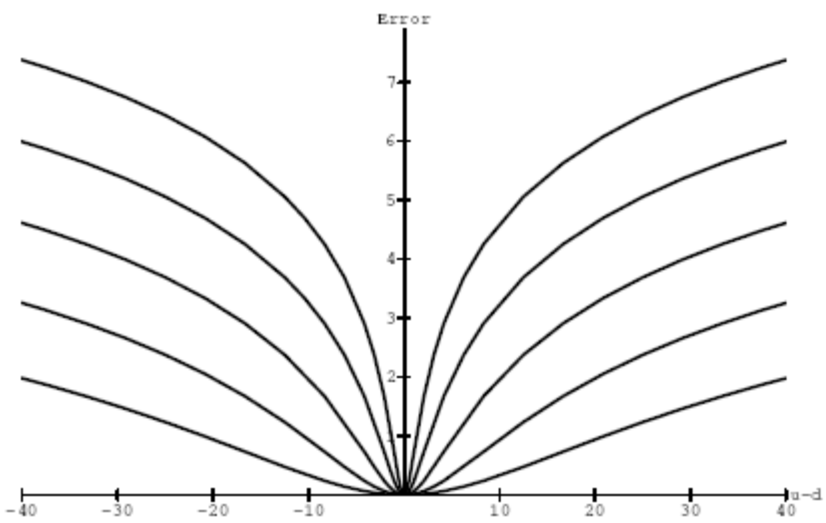

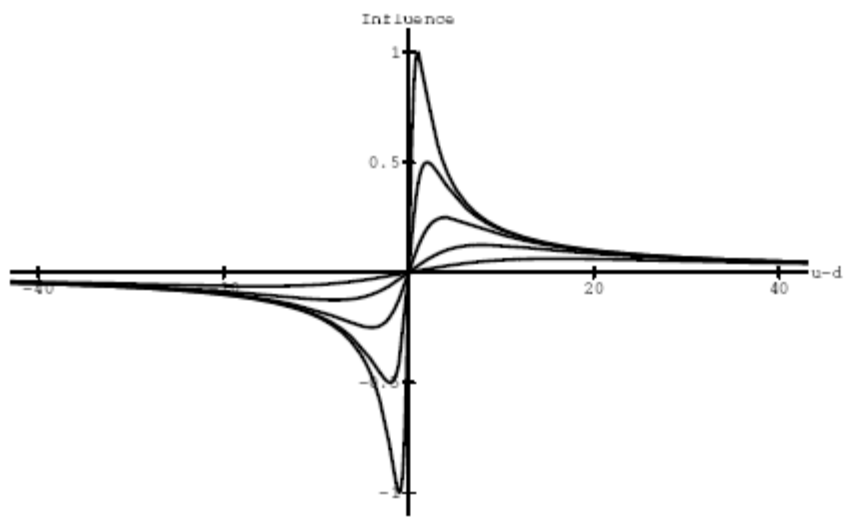

Continuation method

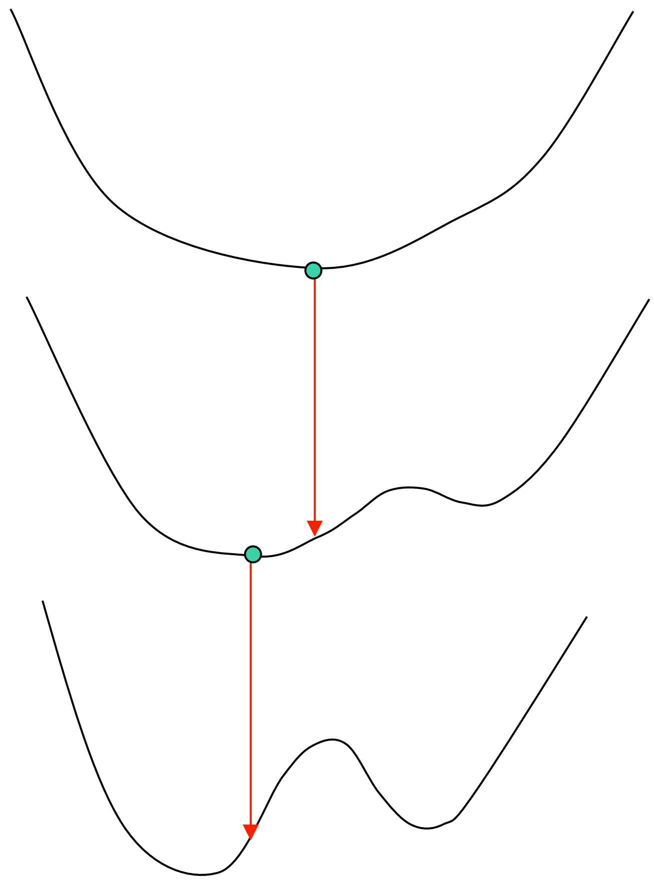

GNC: Graduated Non-Convexity

Probabilistic Interpretation

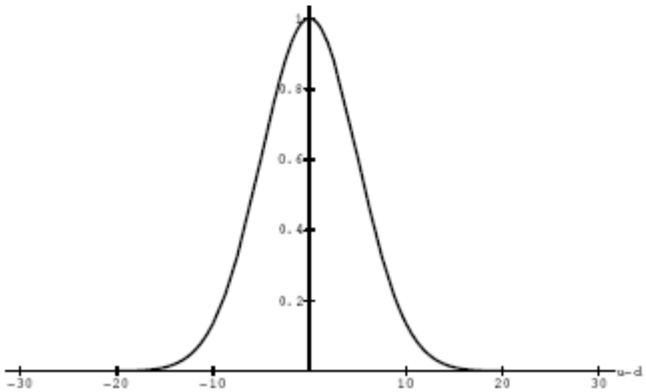

Gaussian

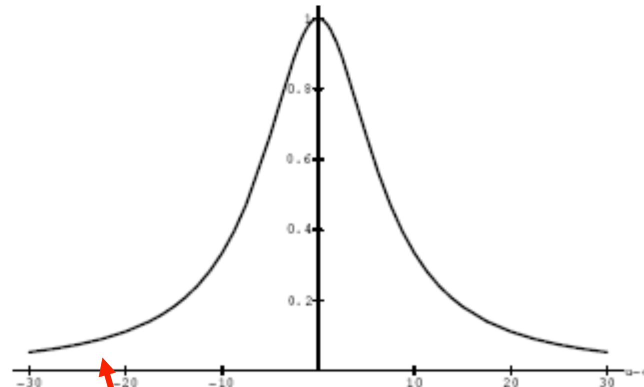

Cauchy

"heavy" tails

Fragmented Occlusion

Results

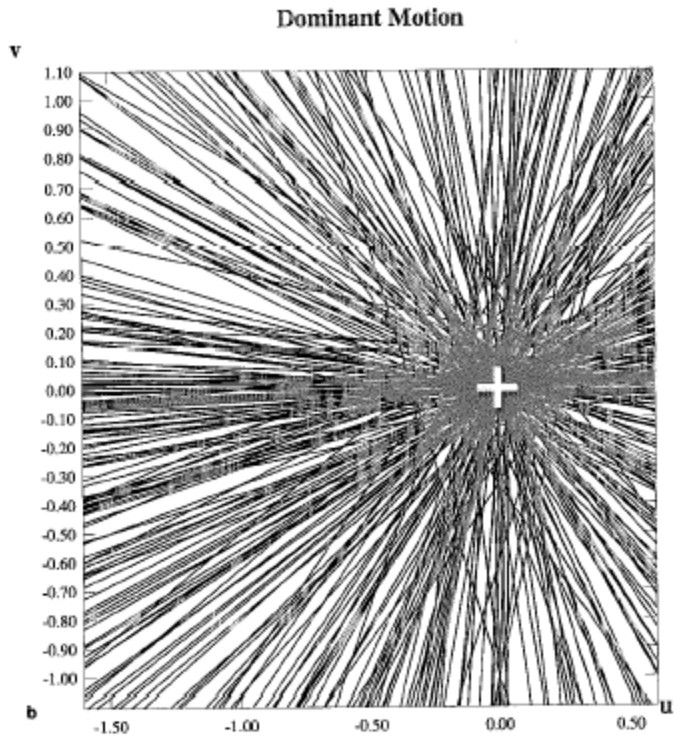
Dominant Motion
v
1.10
1.00
0.90
0.80
0.70
0.60
0.50
0.40
0.30
0.20
0.10
0.00
-0.10
-0.20
-0.30
-0.40
-0.50
-0.60
-0.70
-0.80
-0.90
-1.00
b
-1.50
-1.00
-0.50
0.00
0.50
u


a

Results

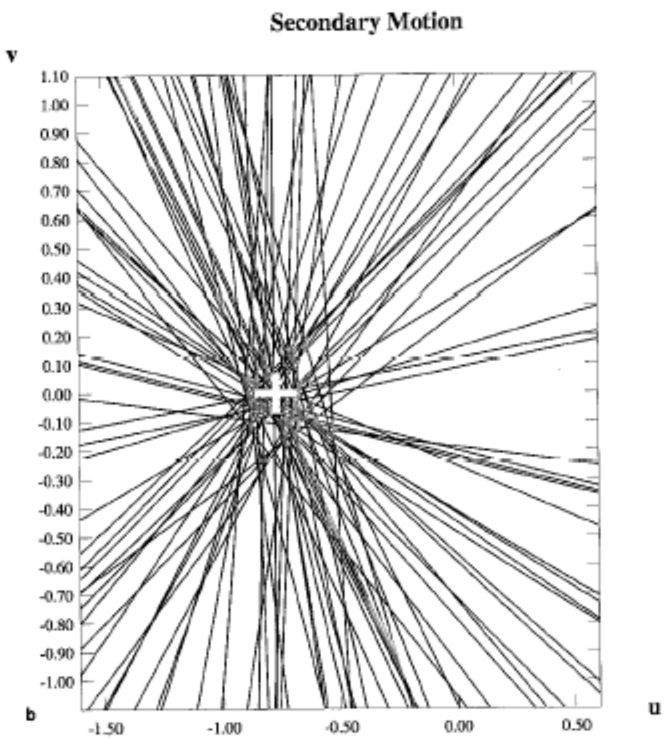
Secondary Motion
v
1.10
1.00
0.90
0.80
0.70
0.60
0.50
0.40
0.30
0.20
0.10
0.00
-0.10
-0.20
-0.30
-0.40
-0.50
-0.60
-0.70
-0.80
-0.90
-1.00
b
-1.50
-1.00
-0.50
0.00
0.50
u

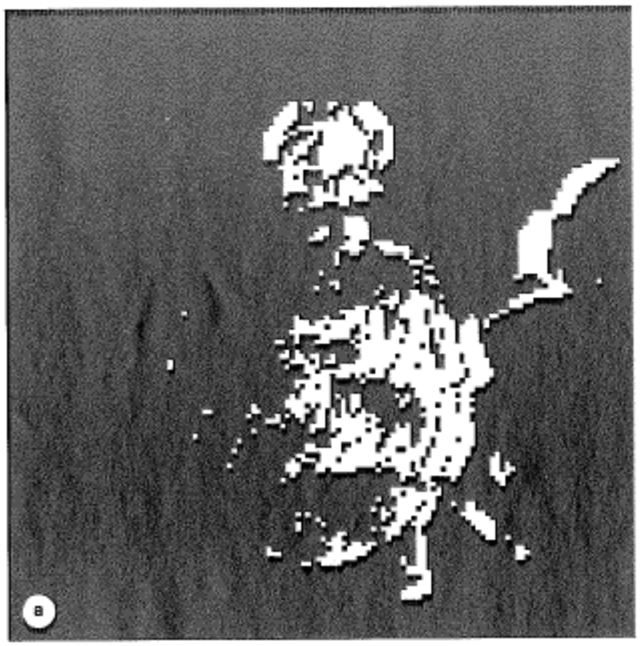
a

What about small regions?

What if the region is a single pixel?

What is the problem with this?

How might we fix it?

Spatial Coherence

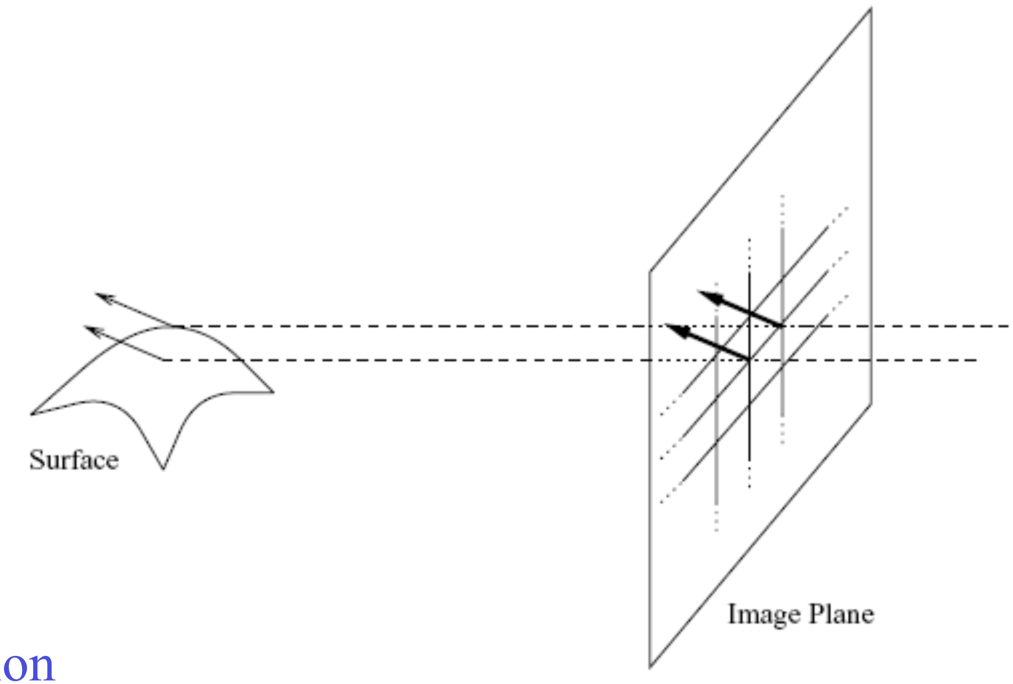

Assumption

* Neighboring points in the scene typically belong to the same surface and hence typically have similar motions.
* Since they also project to nearby points in the image, we expect spatial coherence in image flow.

Formalize this Idea

Noisy 1D signal:

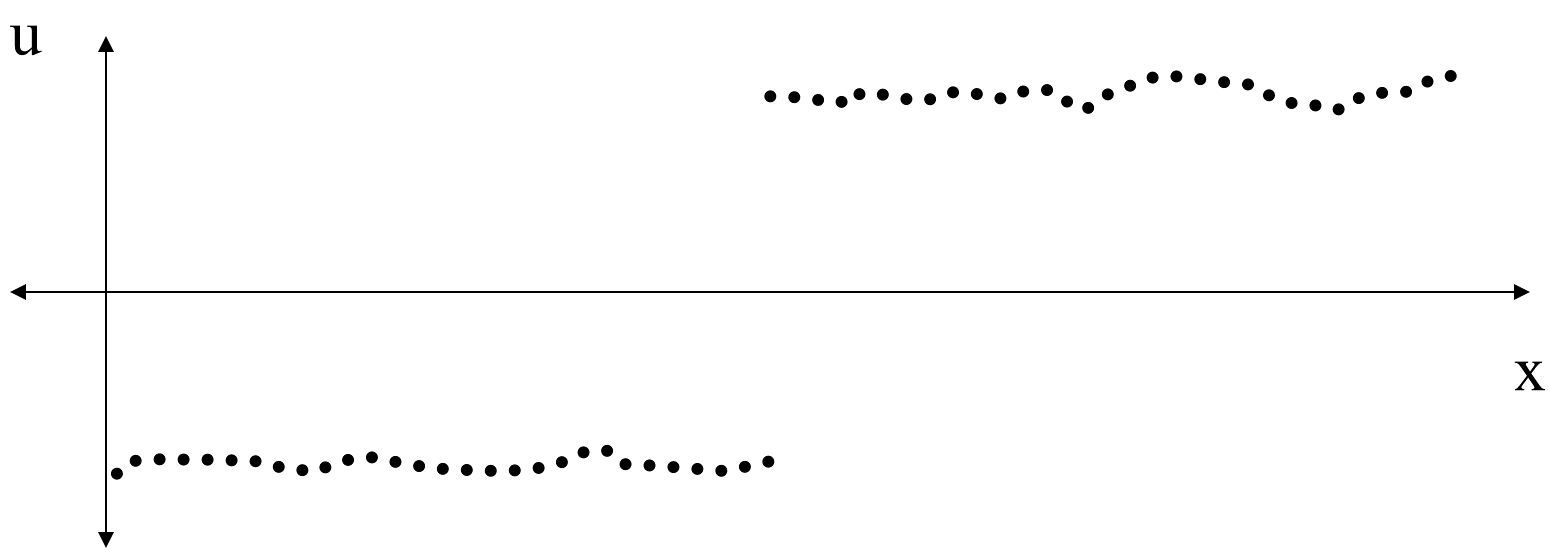

Noisy measurements u(x)

Regularization

Find the “best fitting” smoothed function v(x)

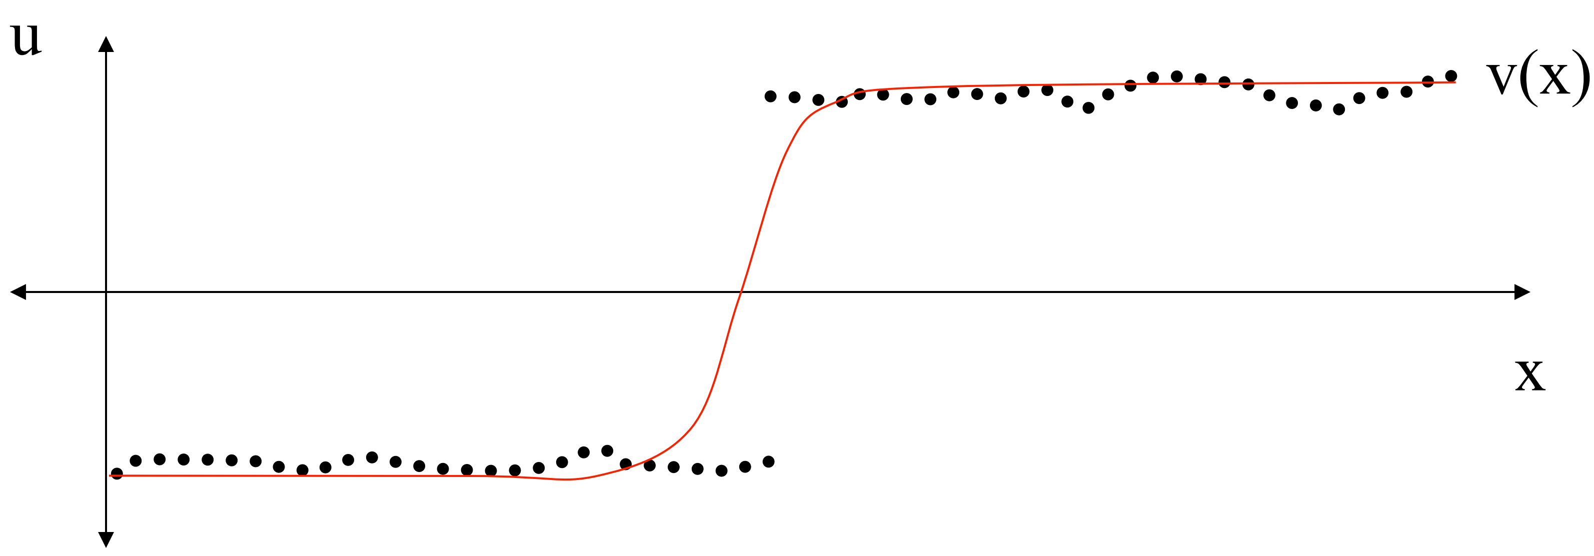

Noisy measurements u(x)

Membrane model

Find the “best fitting” smoothed function v(x)

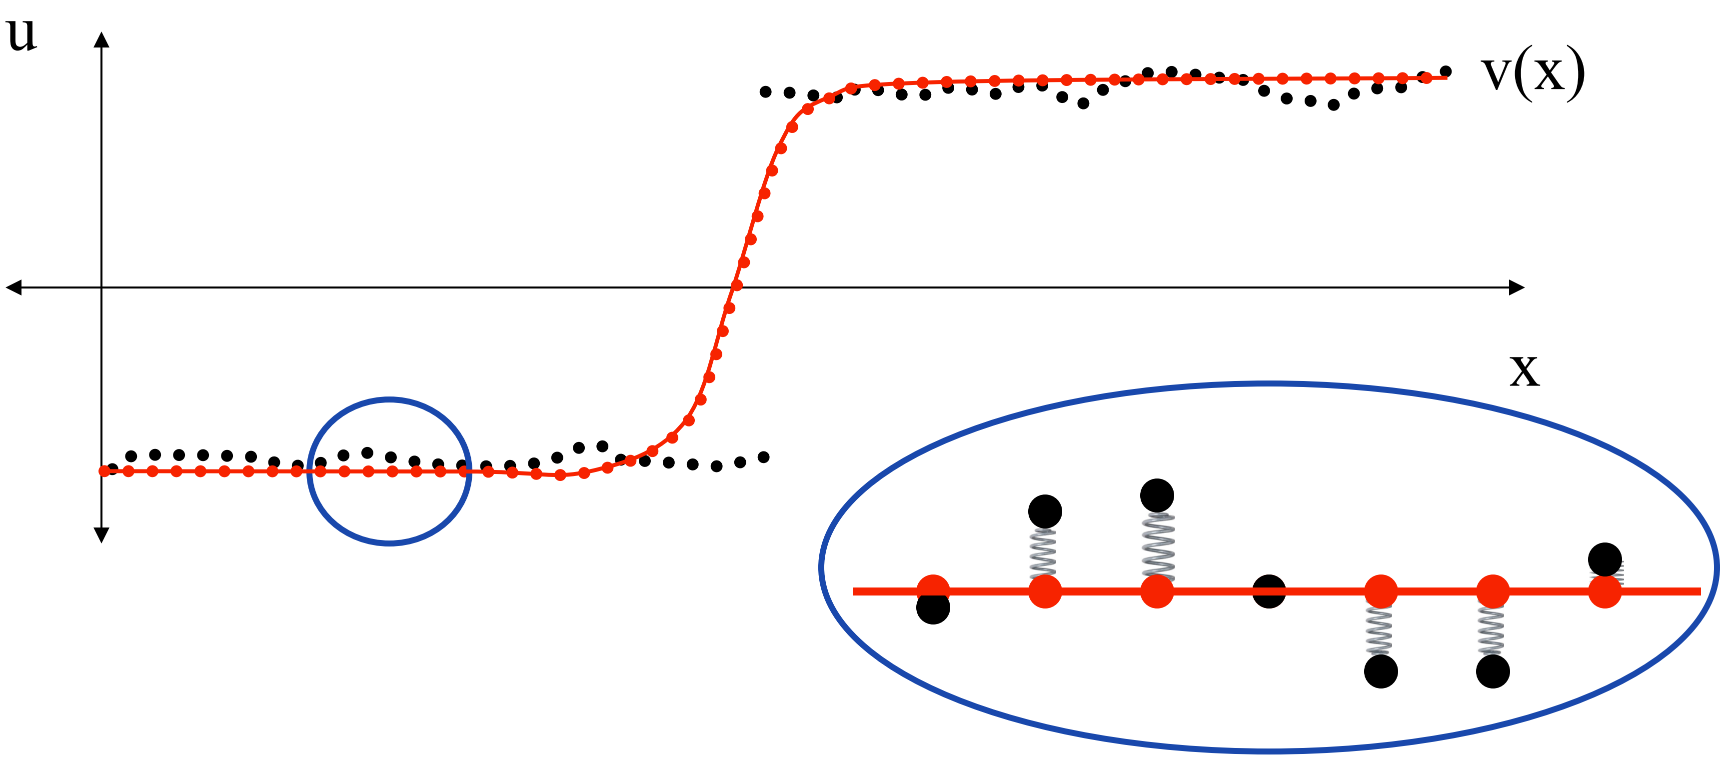

Membrane model

Find the “best fitting” smoothed function v(x)

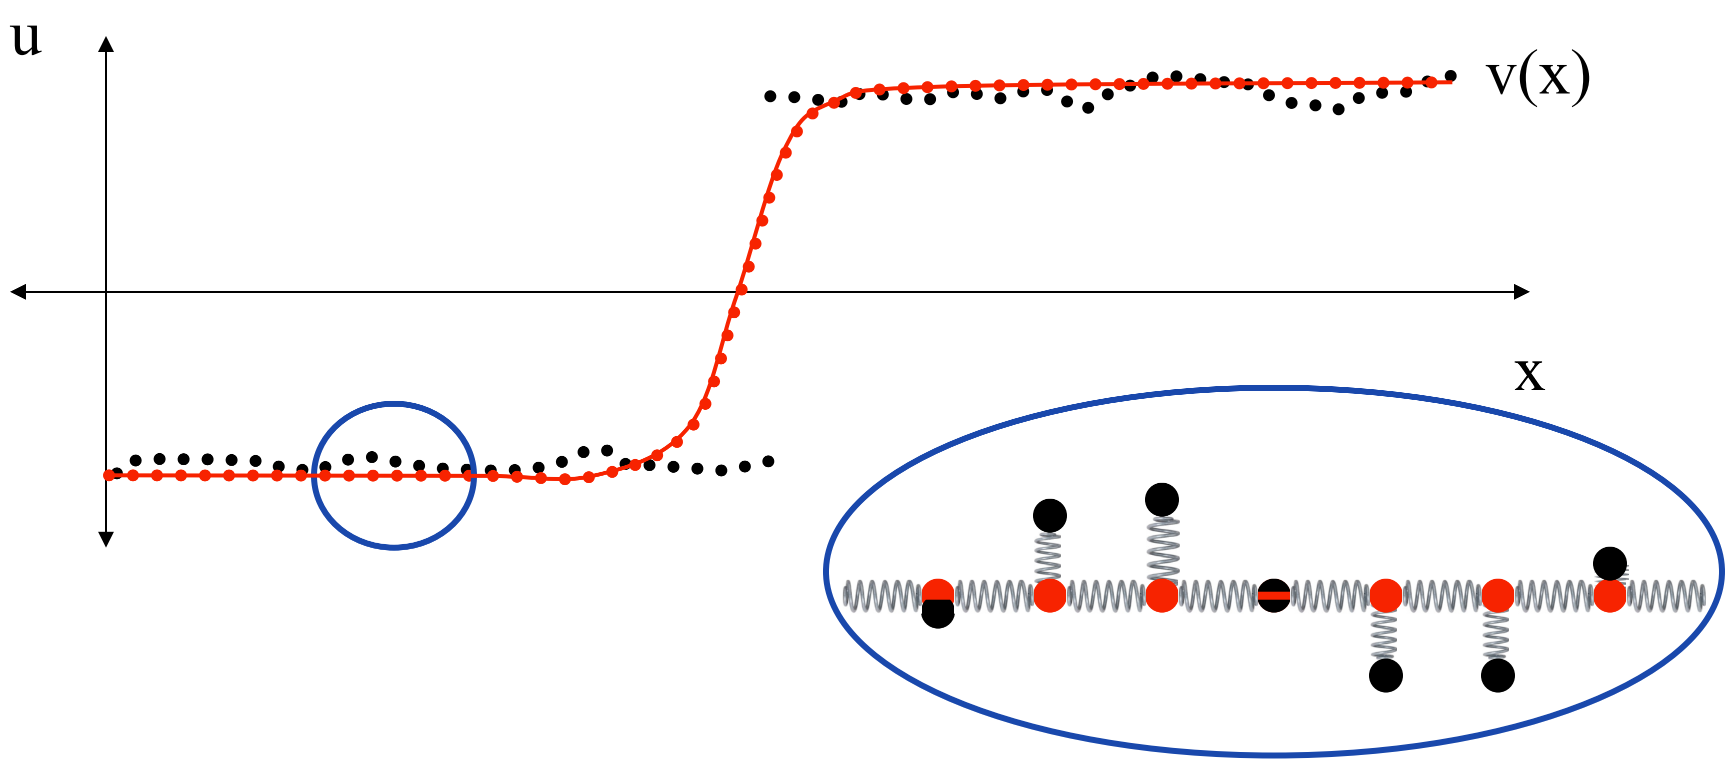

Membrane model

Find the “best fitting” smoothed function v(x)