Introduction to Computer Vision

Michael J. Black

Nov 2009

Regularization and dense flow

Goals

- Today
  - Finish regularization and dense flow
- Friday
  - Start tracking

Formalize this Idea

Noisy 1D signal:

Noisy measurements u(x)

Regularization

Find the "best fitting" smoothed function v(x)

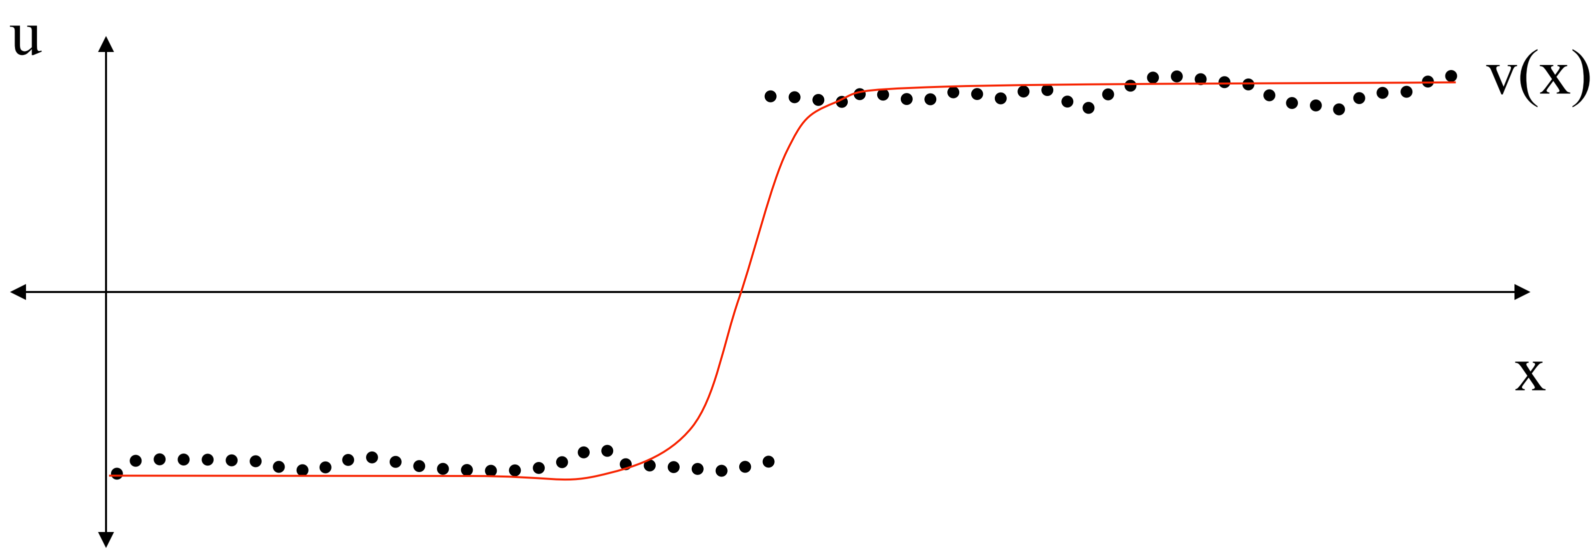

Noisy measurements u(x)

Membrane model

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

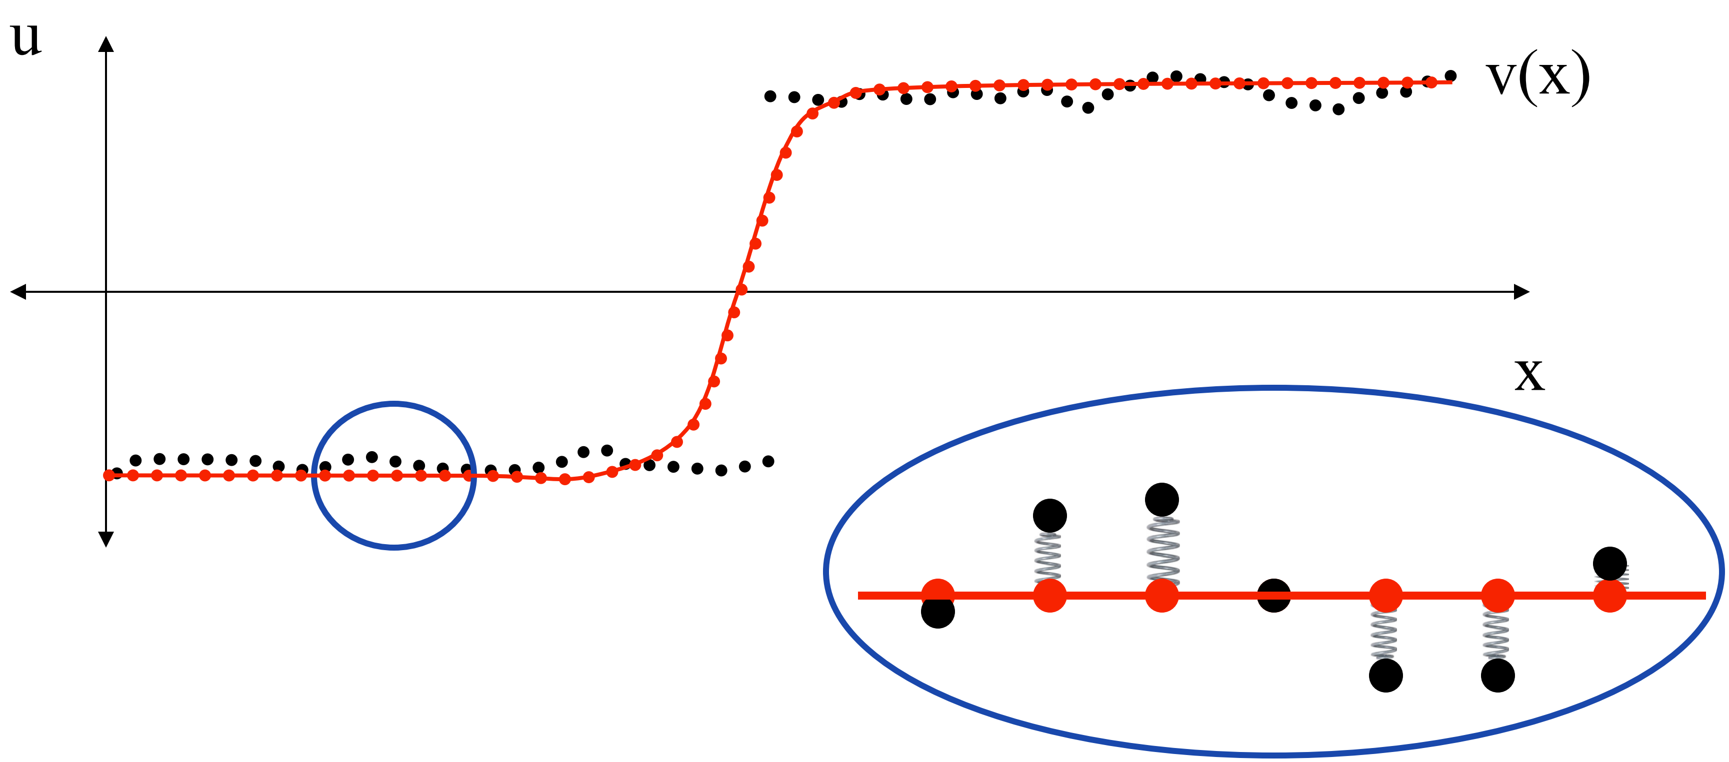

Membrane model

Find the "best fitting" smoothed function v(x)

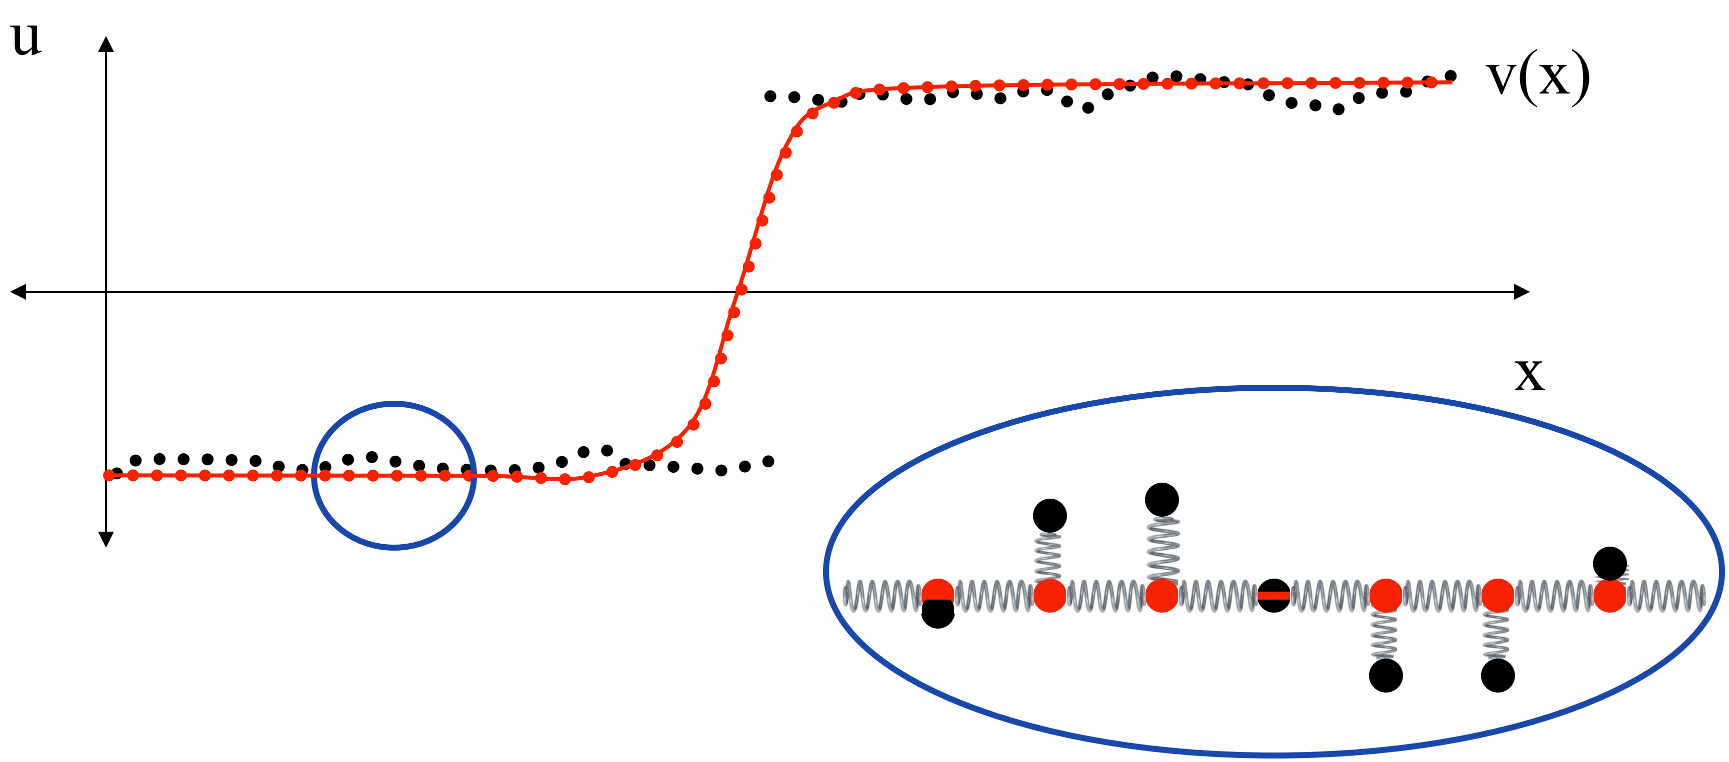

Membrane model

Find the "best fitting" smoothed function v(x)

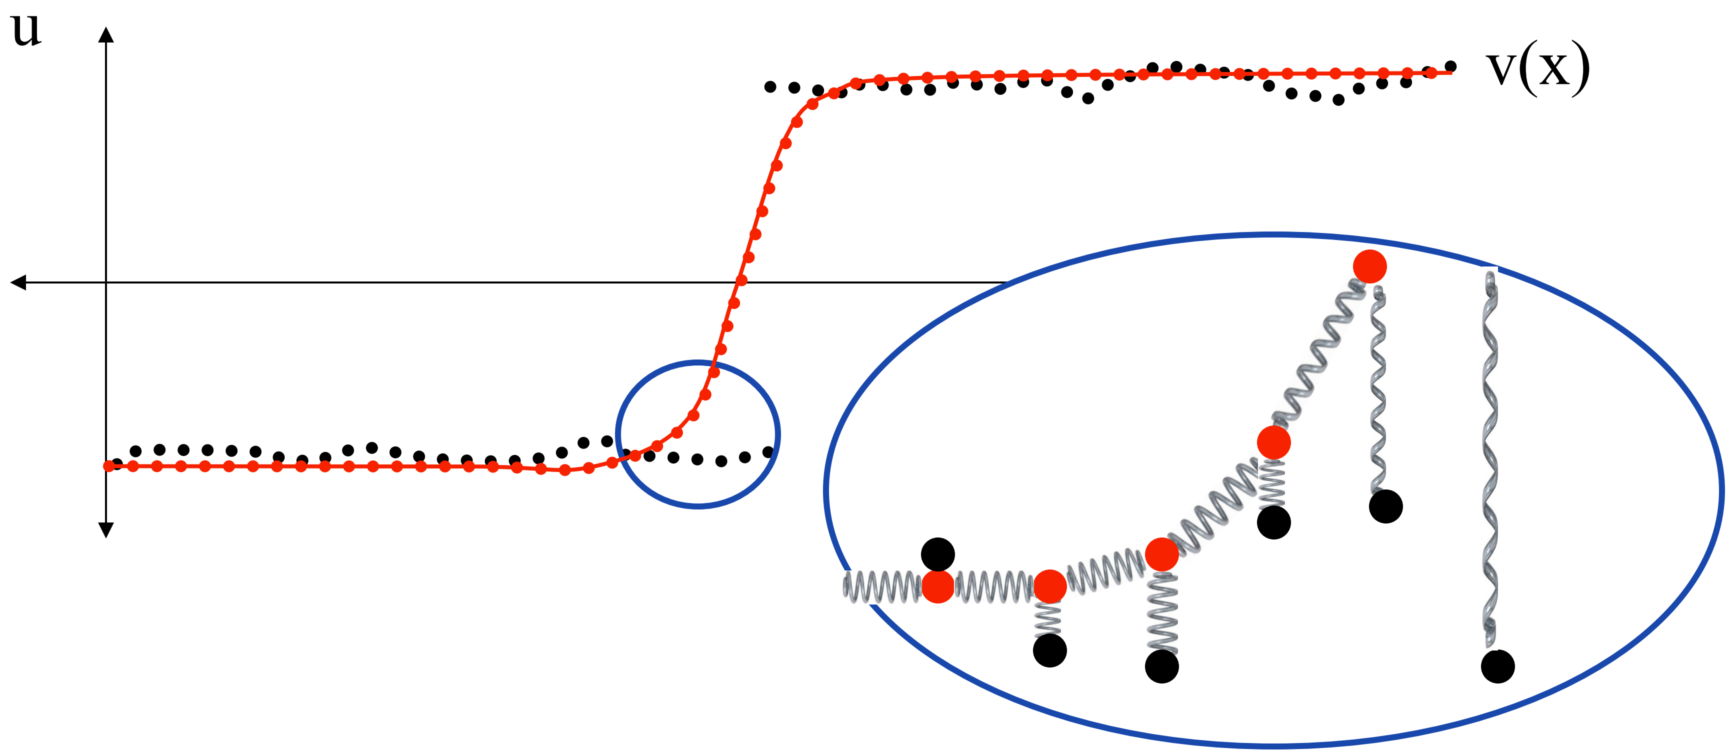

Regularization

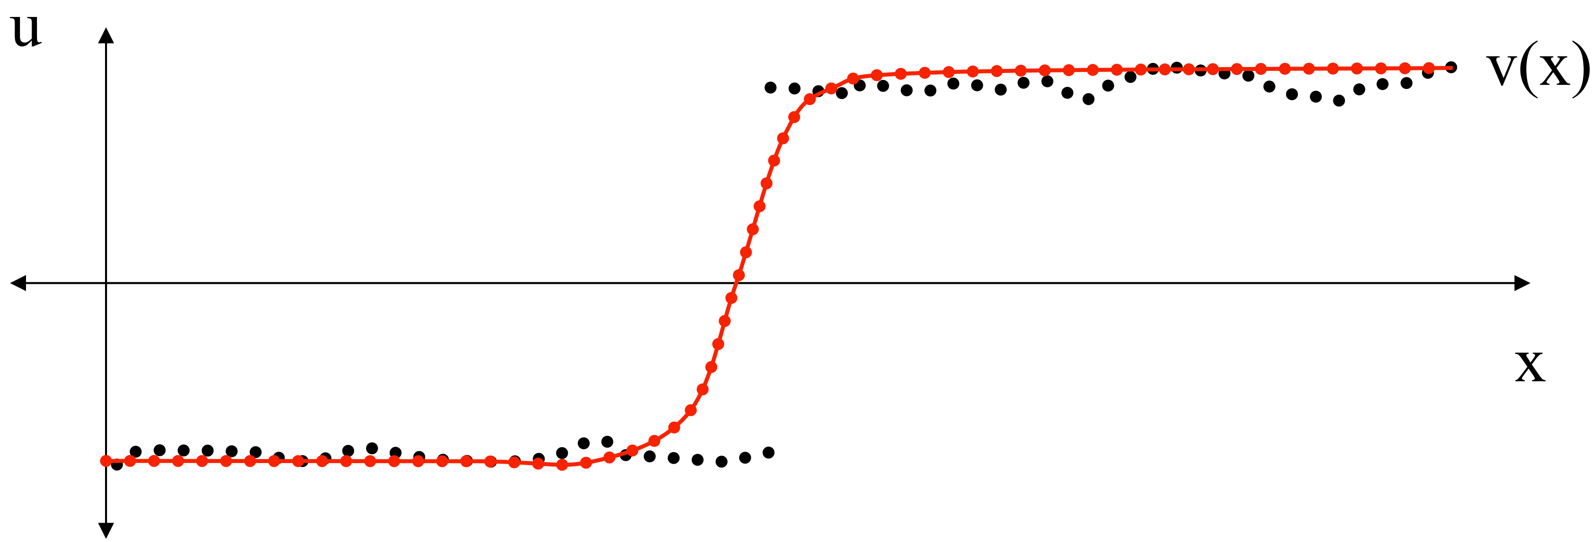

Minimize:

Faithful to the data

Spatial smoothness assumption

$$E(v) = \sum_{x=1}^{N} (v(x) - u(x))^2 + \lambda \sum_{x=1}^{N-1} (v(x+1) - v(x))^2$$

Bayesian Interpretation

$$E(v) = \sum_{x=1}^{N} (v(x) - u(x))^2 + \lambda \sum_{x=1}^{N-1} (v(x+1) - v(x))^2$$

$$p(v \mid u) \propto p(u \mid v)\, p(v)$$

Generative Models

$$u(x) = v(x) + \eta$$

$$\eta \sim N(0, \sigma_1)$$

$$v(x) = v(x+1) + \eta_2$$

$$\eta_2 \sim N(0, \sigma_2)$$

Likelihood and Prior

$$p(v \mid u) \propto p(u \mid v)\, p(v)$$

$$p(u \mid v) = \prod_{x=1}^{N} \frac{1}{\sqrt{2\pi}\sigma_1} \exp\left(-\frac{1}{2}(u(x) - v(x))^2 / \sigma_1^2\right)$$

$$p(v) = \prod_{x=1}^{N-1} \frac{1}{\sqrt{2\pi}\sigma_2} \exp\left(-\frac{1}{2}(v_x(x))^2 / \sigma_2^2\right)^{\lambda}$$

Discontinuities

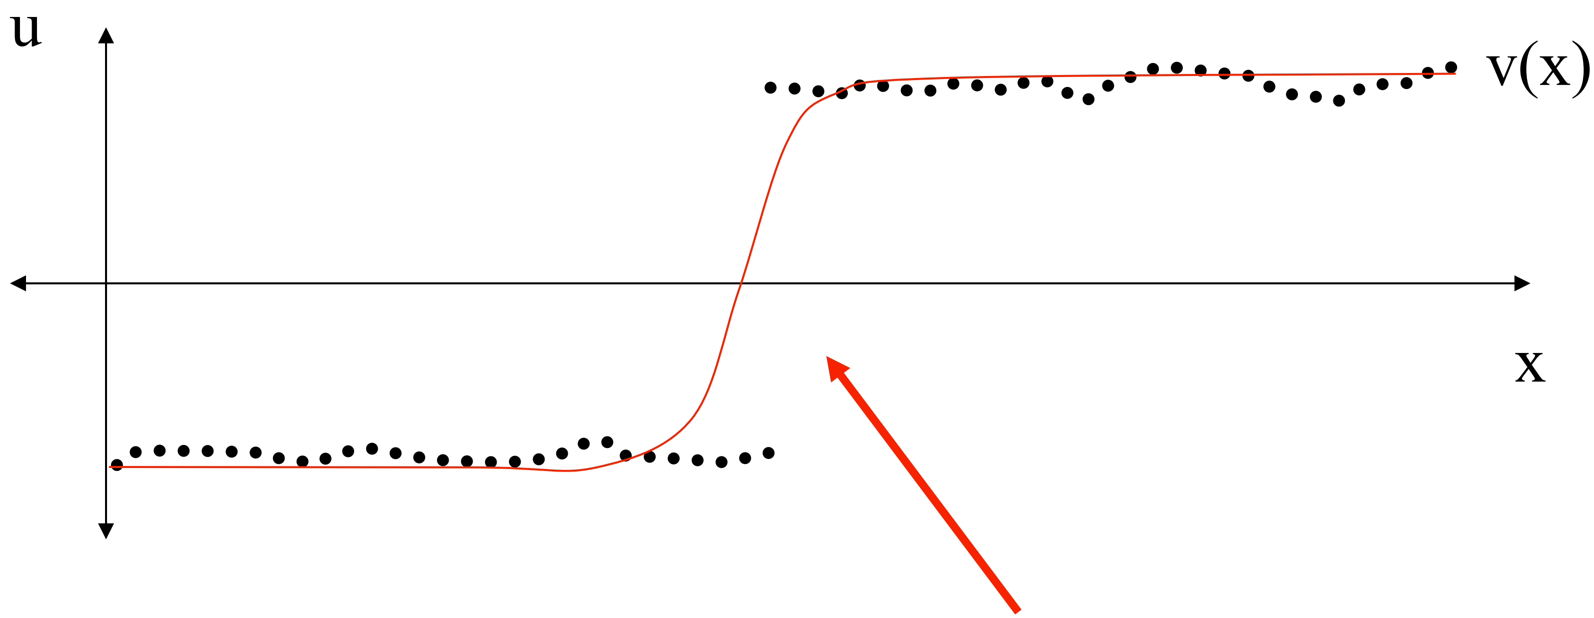

What about this discontinuity?
What is happening here?
What can we do?

Weak membrane model

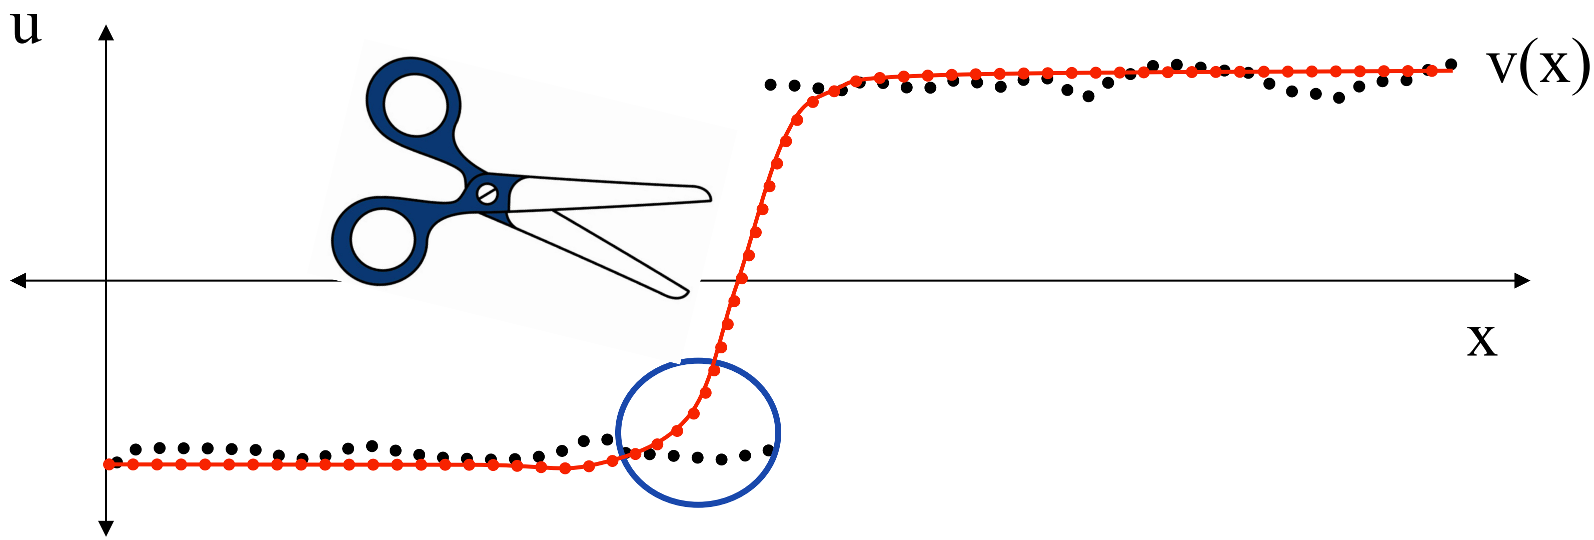

$$E(v,l) = \sum_{x=1}^{N} (v(x) - u(x))^2 + \lambda \sum_{x=1}^{N-1} \left[l(x)(v(x+1) - v(x))^2 + \beta (1 - l(x)) \right]$$

$$l(x) \in \{0,1\}$$

Analog line process

Penalty function

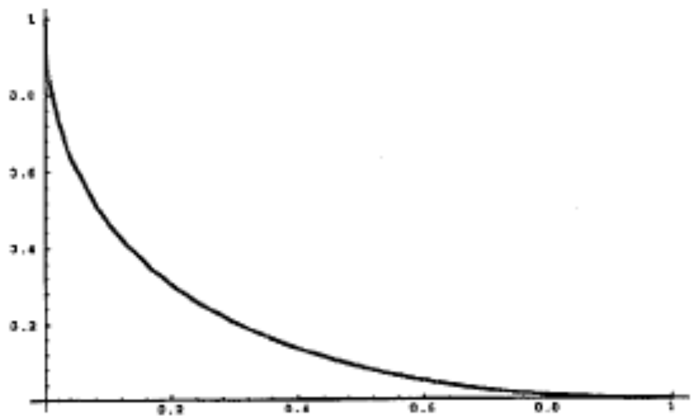

Family of quadratics

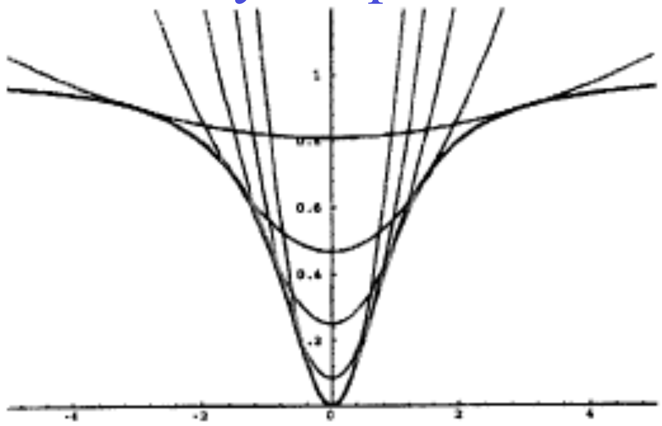

$$E(v,l)=\sum_{x=1}^{N}(v(x)-u(x))^2+\lambda\sum_{x=1}^{N-1}\left[l(x)(v(x+1)-v(x))^2+\Psi(l(x))\right]$$

$$0\le l(x)\le 1$$

Analog line process

Infimum defines a robust error function.

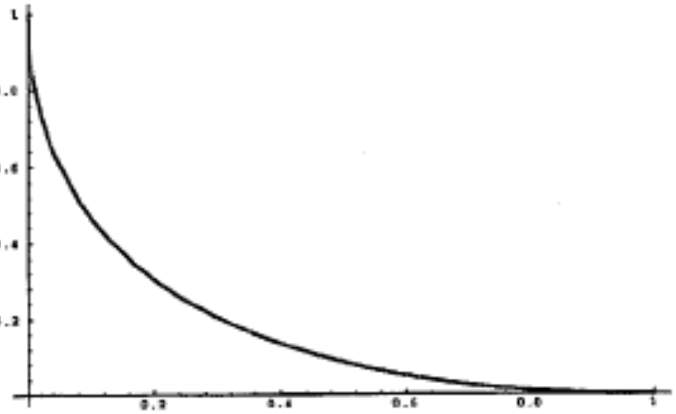

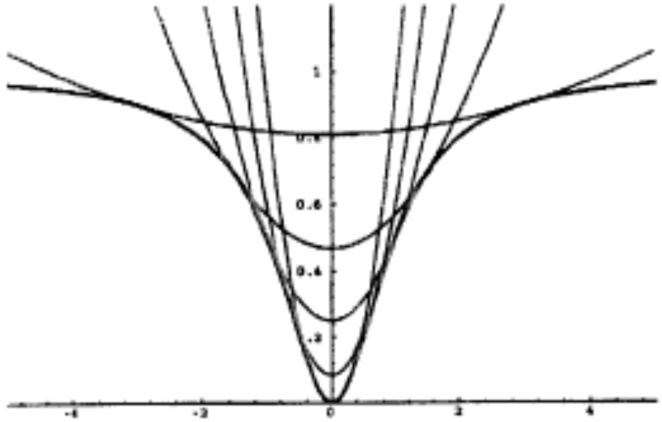

Minima are the same:

$$E(v,l) = \sum_{x=1}^{N} (v(x) - u(x))^2 + \lambda \sum_{x=1}^{N-1} \left[l(x)(v(x+1) - v(x))^2 + \Psi(l(x)) \right]$$

$$E(v) = \sum_{x=1}^{N} (v(x) - u(x))^2 + \lambda \sum_{x=1}^{N-1} \rho(v(x+1) - v(x), \sigma_2)$$

Robust Regularization

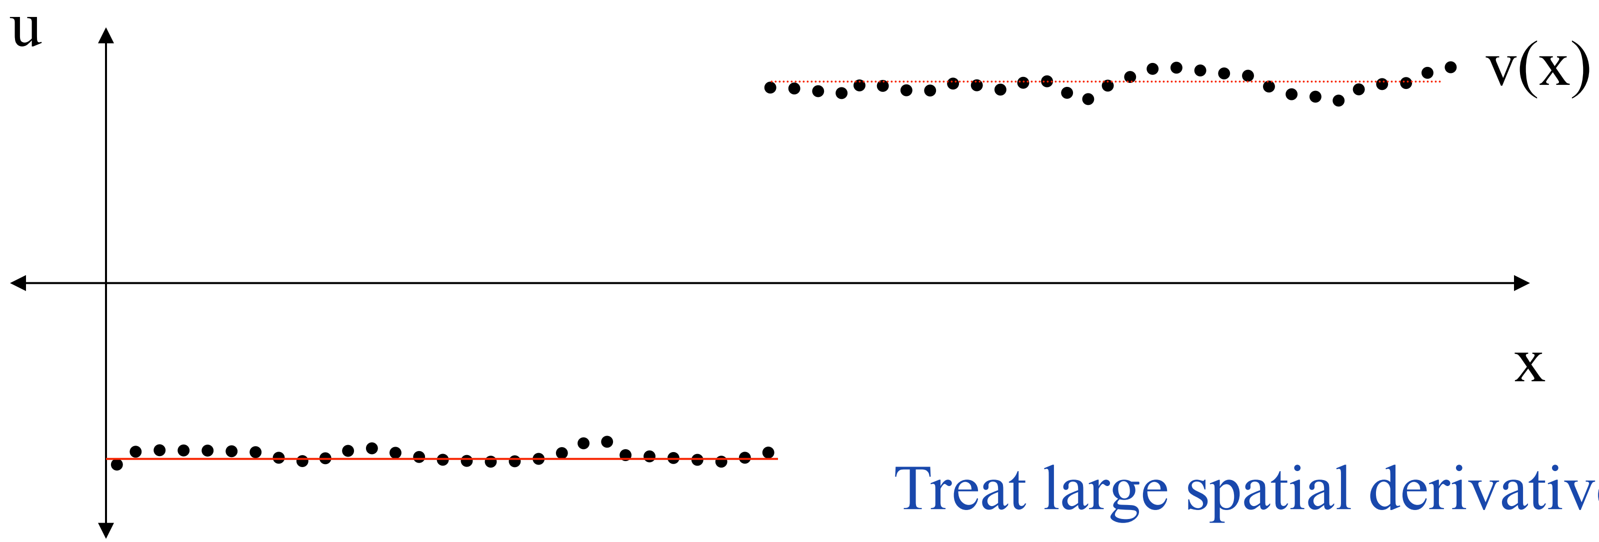

Treat large spatial derivatives as outliers.

Minimize:

$$E(v) = \sum_{x=1}^{N} \rho(v(x) - u(x), \sigma_1) + \lambda \sum_{x=1}^{N-1} \rho(v(x+1) - v(x), \sigma_2)$$

Optical flow

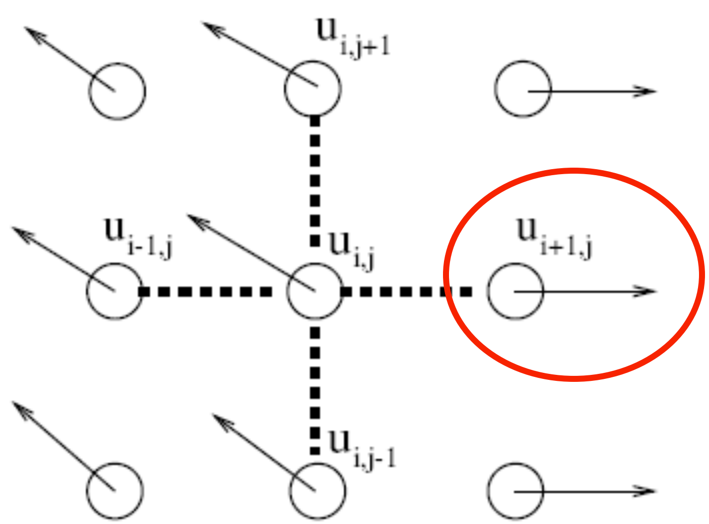

Outlier with respect to neighbors.

Robust formulation of spatial coherence term

$$E_S(u,v) = \rho(u_x) + \rho(u_y) + \rho(v_x) + \rho(v_y)$$

Standard Bayesian formulation

$$p(\mathbf{u}, \mathbf{v} | \mathbf{I}_1, \mathbf{I}_2) \propto p(\mathbf{I}_2 | \mathbf{u}, \mathbf{v}, \mathbf{I}_1) p(\mathbf{u}, \mathbf{v})$$

Data term
How second image can be generated from first image and flow fields

Spatial term
Prior knowledge of flow field

$$E(\mathbf{u}, \mathbf{v}) = E_{\mathrm{D}}(\mathbf{u}, \mathbf{v}) + \lambda E_{\mathrm{S}}(\mathbf{u}, \mathbf{v})$$

“Dense” Optical Flow

$$E_D(\mathbf{u}(\mathbf{x})) = \rho(I_x(\mathbf{x})u(\mathbf{x}) + I_y(\mathbf{x})v(\mathbf{x}) + I_t(\mathbf{x}), \sigma_D)$$

$$E_S(u,v) = \sum_{\mathbf{y}\in G(\mathbf{x})} [\rho(u(\mathbf{x}) - u(\mathbf{y}), \sigma_S) + \rho(v(\mathbf{x}) - v(\mathbf{y}), \sigma_S)]$$

Objective function:

$$E(\mathbf{u}) = \sum_{\mathbf{x}} E_D(\mathbf{u}(\mathbf{x})) + \lambda E_S(\mathbf{u}(\mathbf{x}))$$

When ρ is quadratic = “Horn and Schunck”

Optimization

$$u^{(n+1)} = u^{(n)} - \omega \frac{1}{T(u)} \frac{\partial E}{\partial u}$$

$$v^{(n+1)} = v^{(n)} - \omega \frac{1}{T(v)} \frac{\partial E}{\partial v}$$

$$\frac{\partial E}{\partial u_s} = \psi(I_x u_s + I_u v_s + I_t, \sigma_D) I_x + \lambda \sum_{n \in G(s)} \psi(u_s - u_n, \sigma_S)$$

$T(u)$ = max of second derivative

Optimization

- Gradient descent
- Coarse-to-fine (pyramid)
- Deterministic annealing

Example

Example

Quadratic:

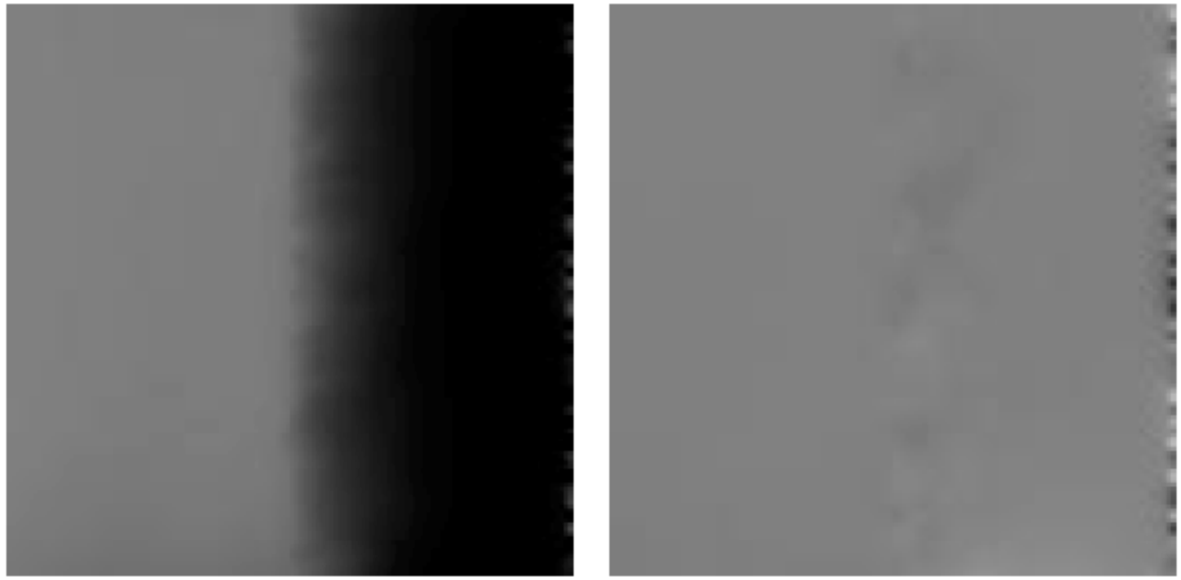

Robust:

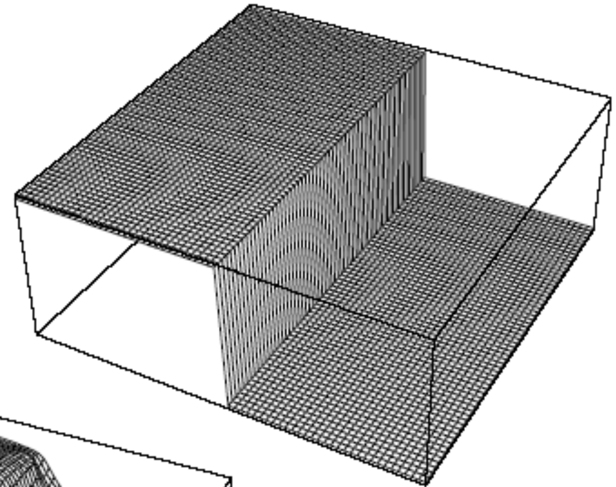

a

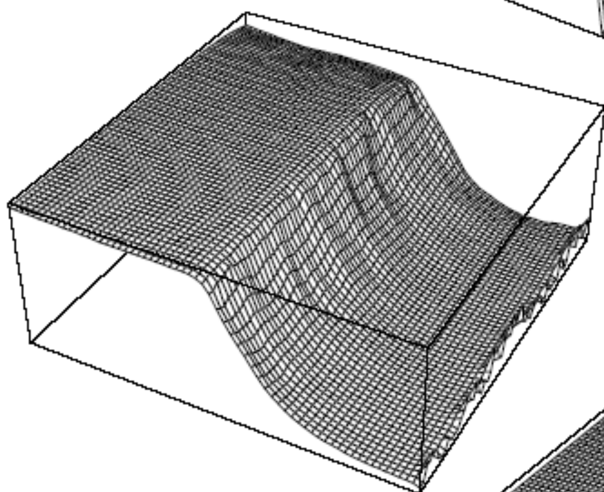

b

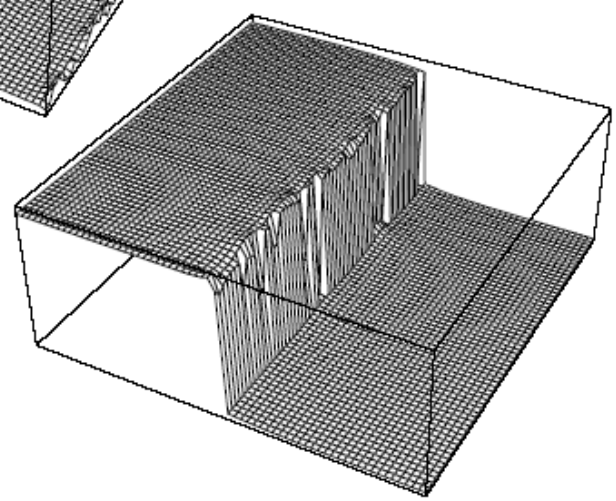

c

Magnitude of horizontal flow

Outliers

Points where the influence is reduced

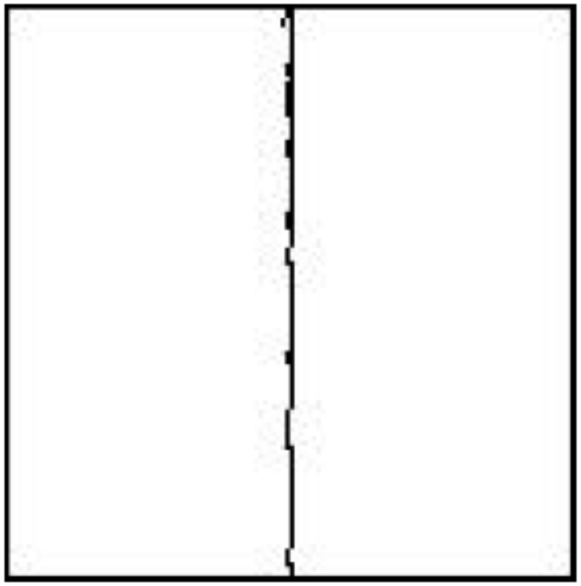

Spatial term

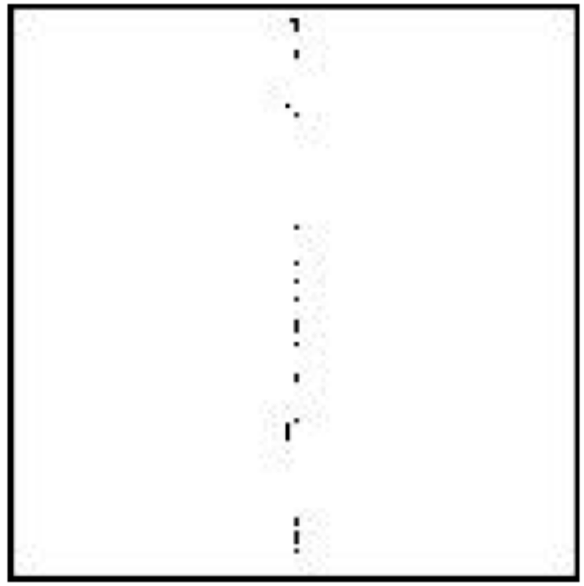

Data term

With 5% uniform random noise added to the images.

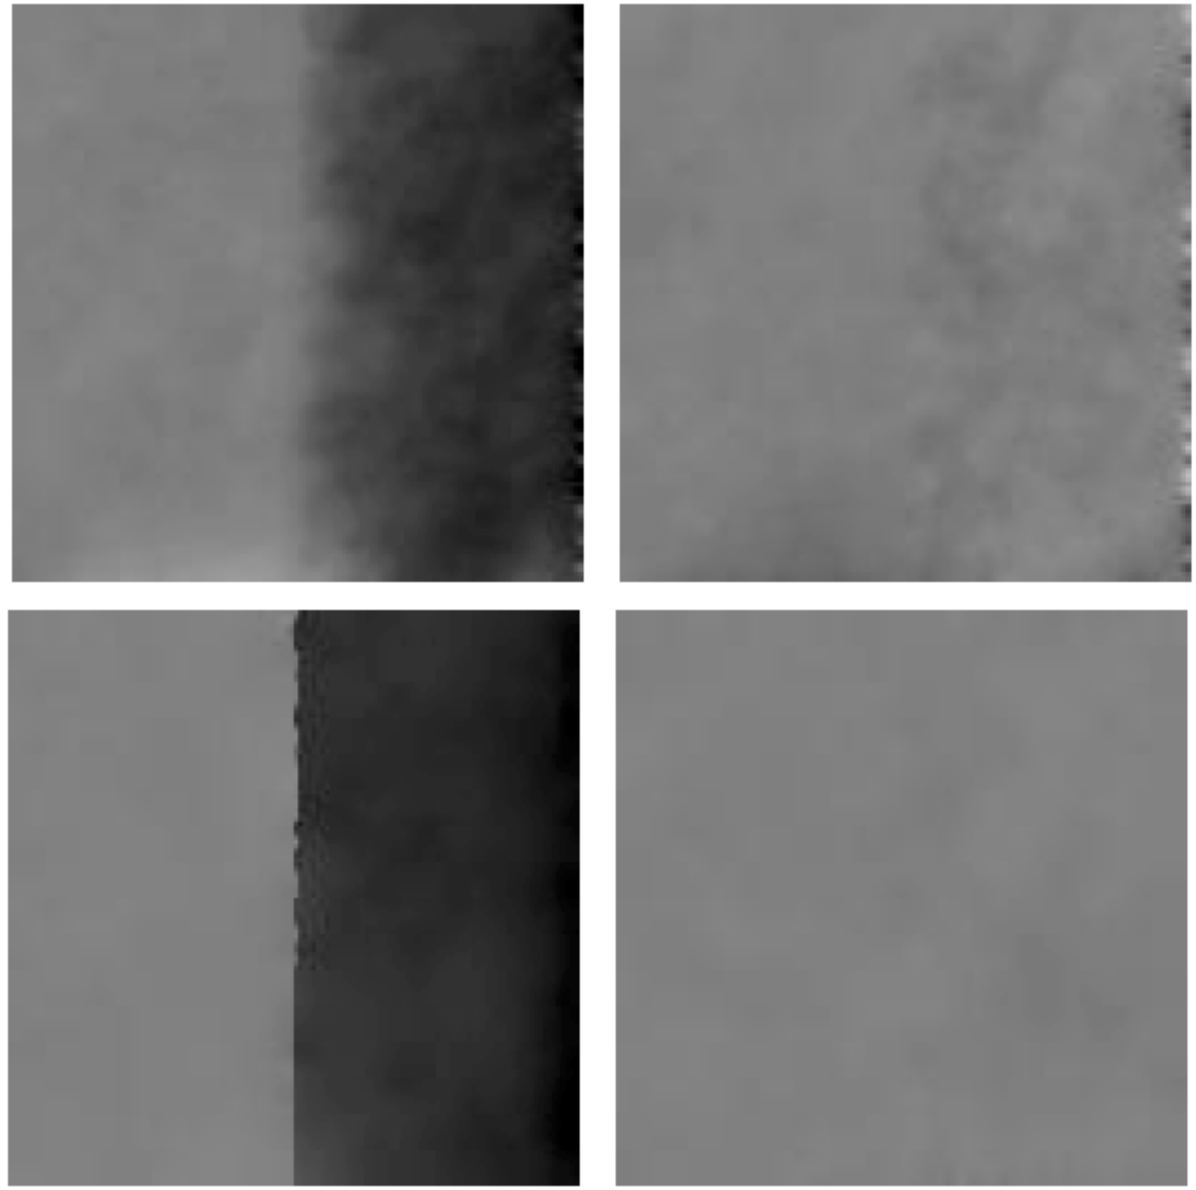

Horizontal Component

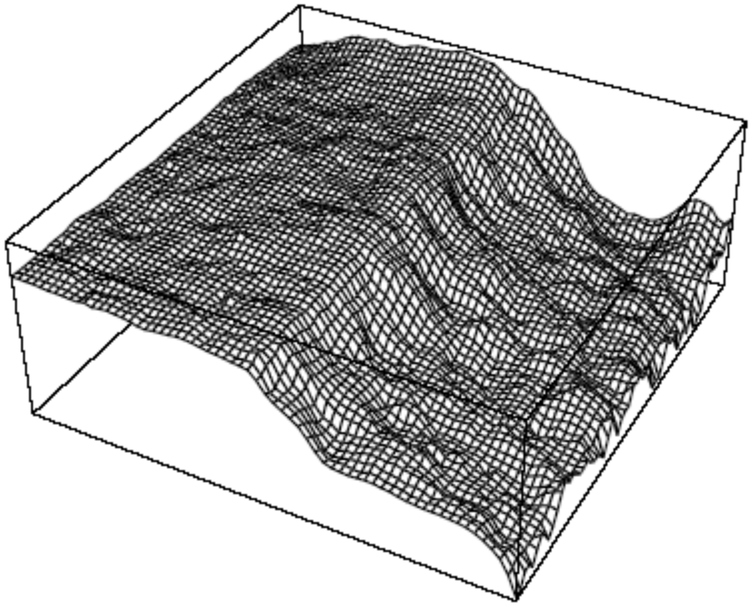

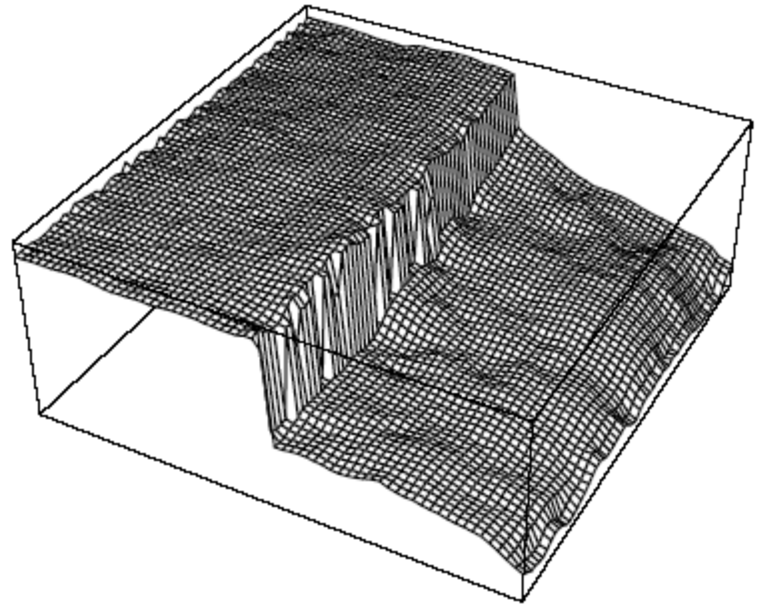

More Noise

Quadratic:

Quadratic data term,
robust spatial term:

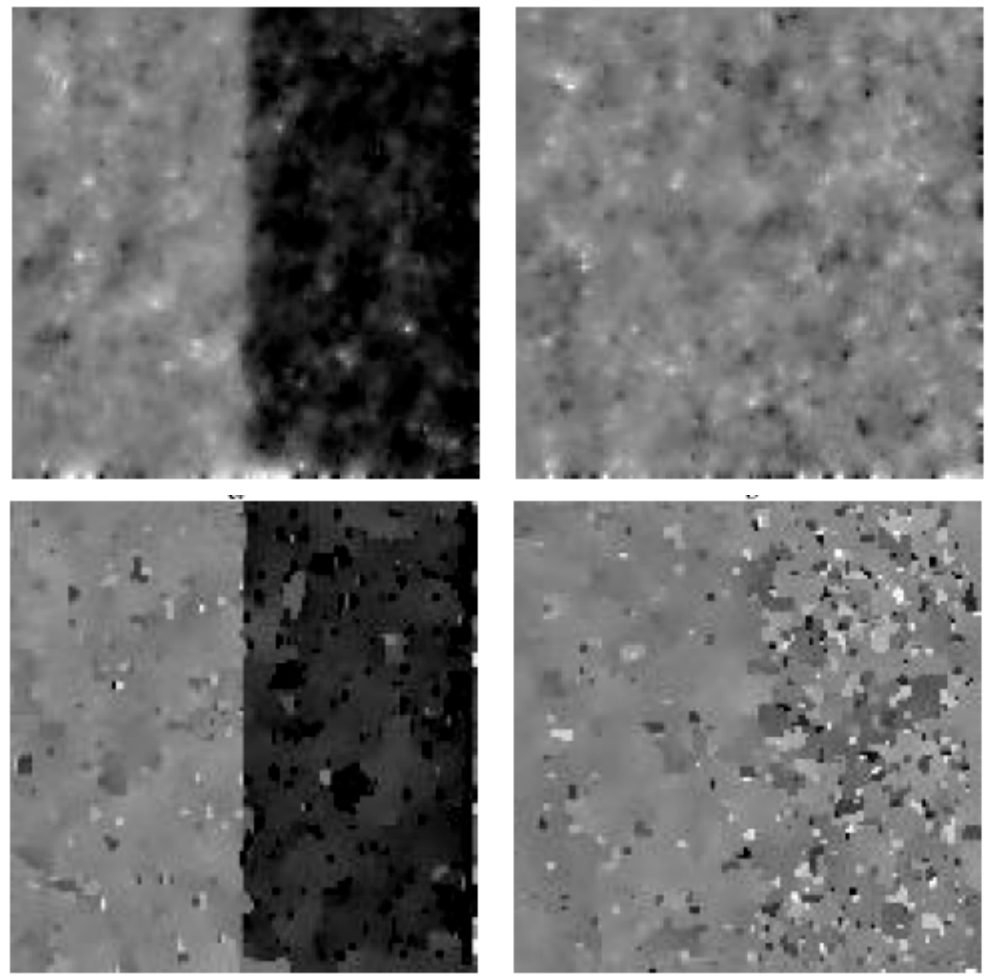

Both Terms Robust

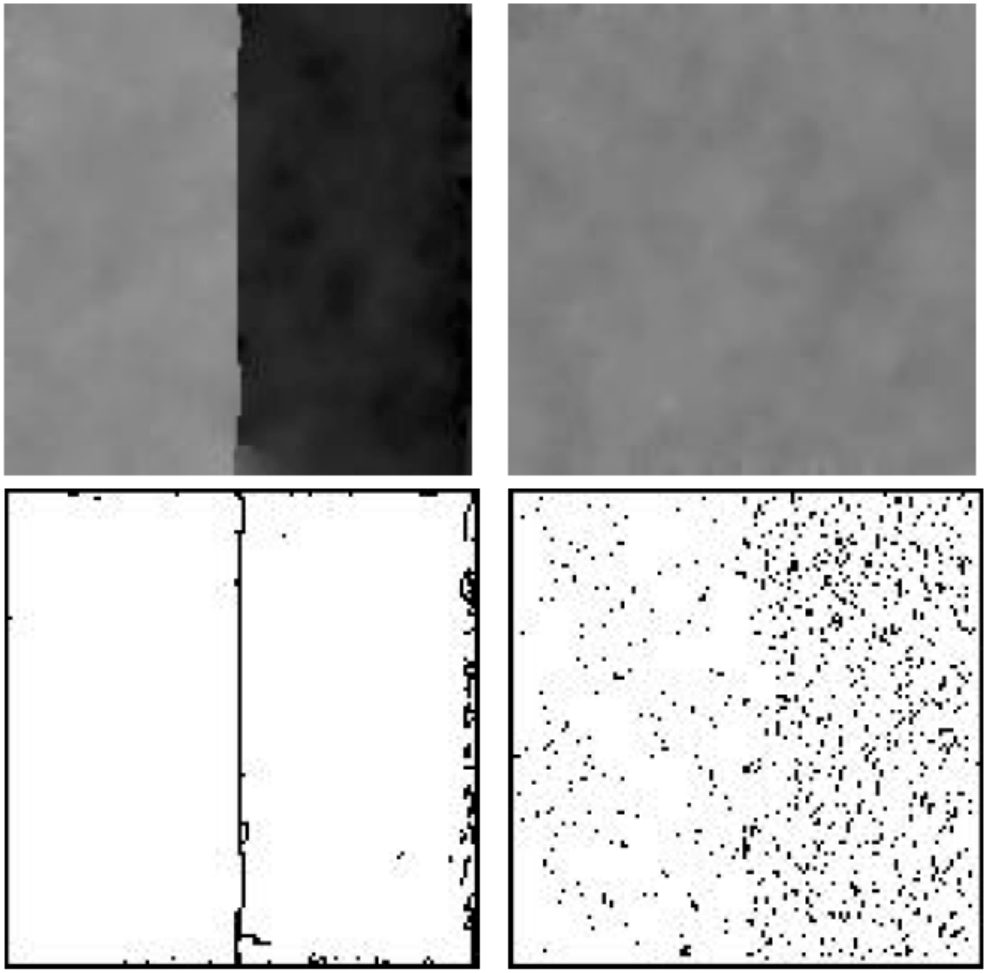

Spatial and data outliers:

Pepsi sequence

Real Sequence


Deterministic annealing.
First stage (large s):

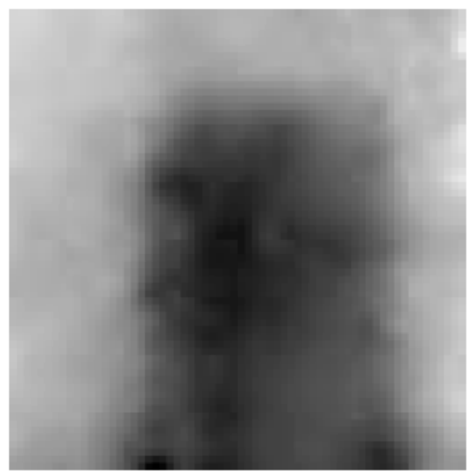
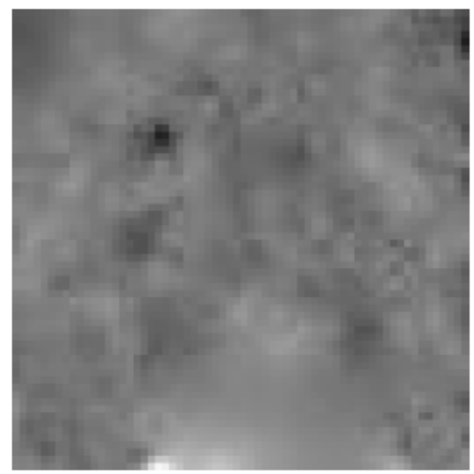

Real Sequence

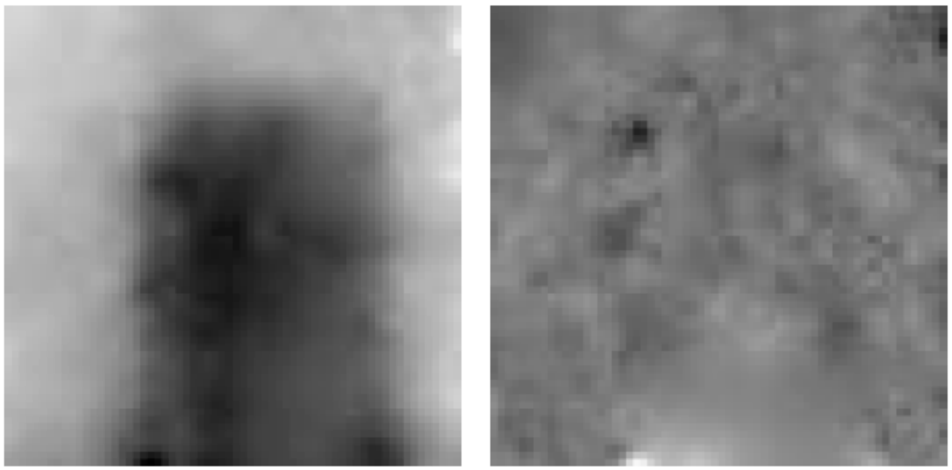

Final result after annealing: