Bayesian Modeling of Uncertainty in Low-Level Vision

Paper by Richard Szeliski

International Journal of Computer Vision

(1990)

Presented by Deqing Sun

Low-level vision

Extract dense fields (intrinsic images)

- Surface interpolation: value of every image pixel
- Optical flow: motion of every image pixel

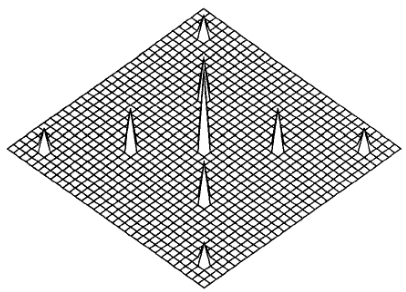

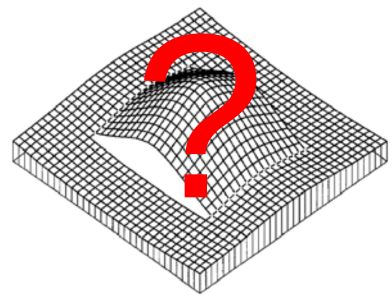

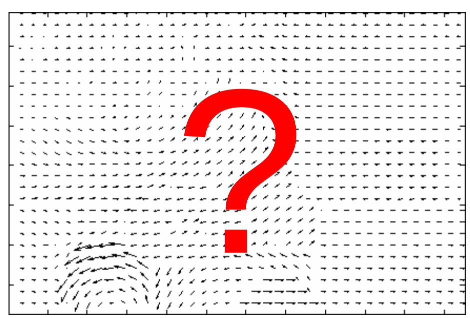

Low-level vision

Generally ill-posed: data insufficiently constraint solution

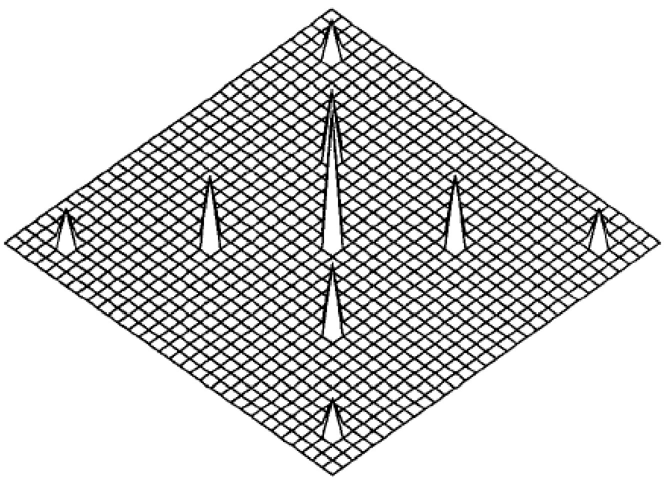

d

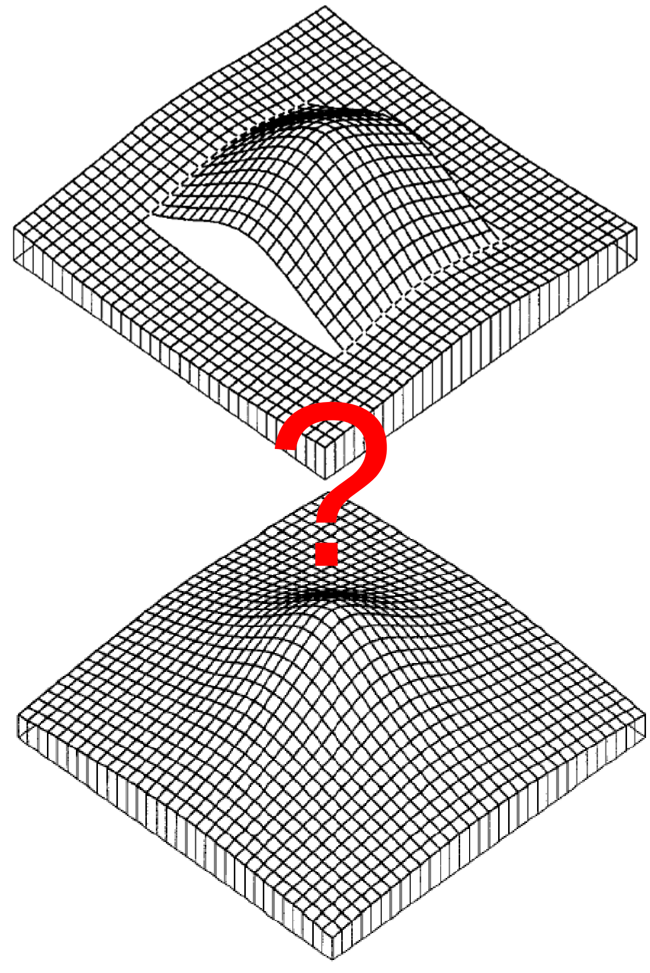

u

Regularization

- Energy minimization

$$E(u) = (1-\lambda)E_d(u) + \lambda E_p(u)$$

- Data compatibility

$$E_d(u) = \frac{1}{2}\sum_i c_i[u(x_i, y_i) - d_i]^2$$

- Smoothness (1st order)

$$E_p(u) = \frac{1}{2}\int\int(u_x^2 + u_y^2)dxdy$$

Discretization

$$E_p(\mathbf{u}) = \frac{1}{2}\sum_{(i,j)}[(u_{i+1,j} - u_{i,j})^2 + (u_{i,j+1} - u_{i,j})^2]$$

Bayesian modeling

- Bayes rule

$$p(\mathbf{u}|\mathbf{d}) = \frac{p(\mathbf{d}|\mathbf{u})p(\mathbf{u})}{p(\mathbf{d})}$$

- Sensor model $p(\mathbf{d}|\mathbf{u})$

  How to generate data from unknown

- Prior model $p(\mathbf{u})$

  Prior information about unknown

Prior model

- MRF: conditional independence

$$p(u_i | \mathbf{u}) = p(u_i | \{u_j\}), \qquad j \in N_i$$

- Gibbs distribution

$$p(\mathbf{u}) = \frac{1}{Z_p} \exp(-E_p(\mathbf{u})/T_p)$$

$$E_p(\mathbf{u}) = \sum_{c \in C} E_c(\mathbf{u})$$

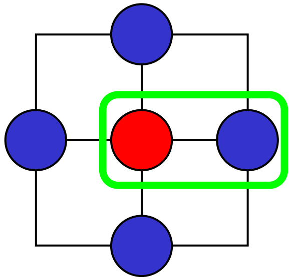

- 1^{st} order smoothness

$$E_{\{(i,j),(i+1,j)\}}(\mathbf{u}) = (u_{i+1,j} - u_{i,j})^2$$

Sensor model

- Single-clique MRF

$$p(\mathbf{d}|\mathbf{u}) = \frac{1}{Z_d} \exp(-E_d(\mathbf{u}, \mathbf{d}))$$

$$E_d(\mathbf{u}, \mathbf{d}) = \sum_i E_d^i(u_i, d_i)$$

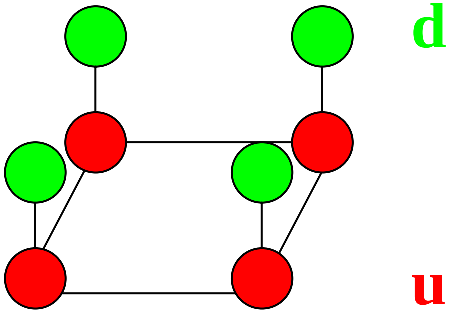

- Quadratic penalty implies Gaussian noise

$$p(d_i|\mathbf{u}) = \frac{1}{\sqrt{2\pi}\sigma_i} \exp\left(-\frac{(u(x_i, y_i) - d_i)^2}{2\sigma_i^2}\right)$$

Sensor model

Contaminated Gaussian

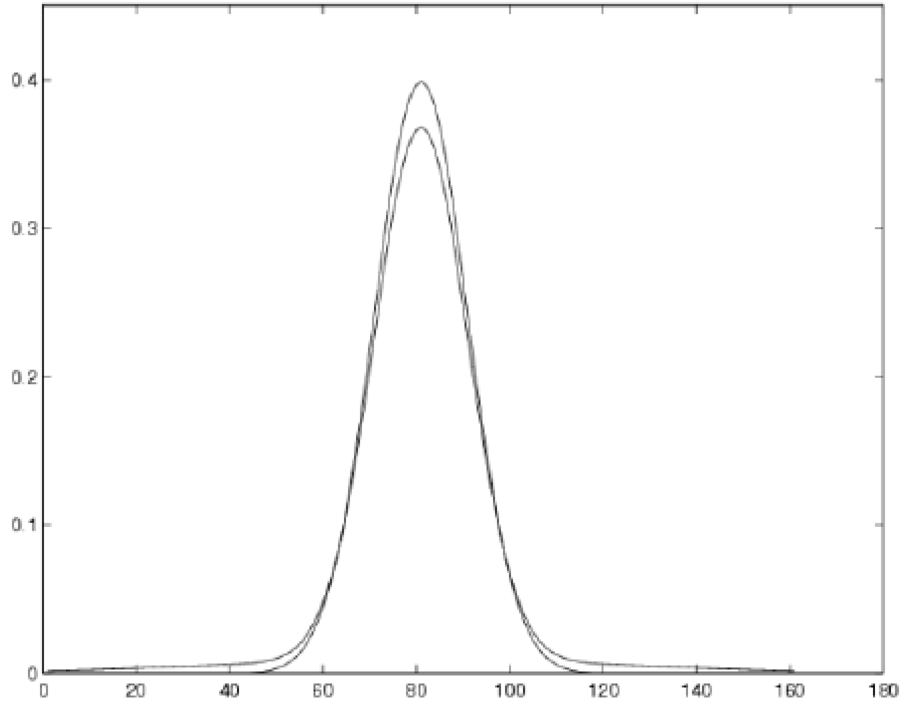

$$p(d_i | u) = \frac{1-\varepsilon}{\sqrt{2\pi}\sigma_1} \exp\left(\frac{-(u_i - d_i)^2}{2\sigma_1^2}\right) + \frac{\varepsilon}{\sqrt{2\pi}\sigma_2} \exp\left(\frac{-(u_i - d_i)^2}{2\sigma_2^2}\right)$$

Sensor model

Contaminated Gaussian: energy function

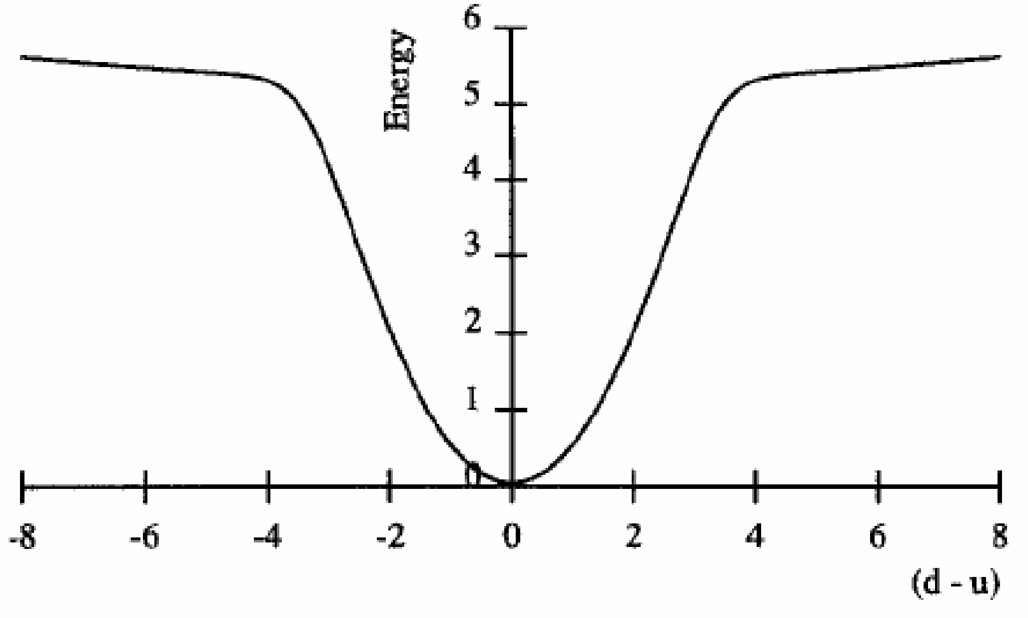

Fig. 5. Constraint energy for contaminated Gaussian.

Posterior distribution

- MRF too

$$p(\mathbf{u}|\mathbf{d}) = \frac{p(\mathbf{d}|\mathbf{u})p(\mathbf{u})}{p(\mathbf{u})} = \frac{1}{Z}\exp(-E(\mathbf{u}))$$

$$E(\mathbf{u}) = E_{\mathrm{p}}(\mathbf{u})/T_{\mathrm{p}} + E_{\mathrm{d}}(\mathbf{u}, \mathbf{d})$$

- MAP estimate same as energy minimization

- Bayesian modeling provides more

Optimal estimation

- Minimize loss function

$$\langle L \rangle = \int L(u, u')p(u \mid d)du$$

- MAP estimation (negative delta)

$$\langle L \rangle = \int -\delta(u, u')p(u \mid d)du$$

Wrong estimates equally bad

Uncertainty

Variance at each point (2nd order statistics)

$$\text{var}\ (u_i) = \sigma_i^2 = \int (u_i - u_i^*)^2\ p(\mathbf{u}|\mathbf{d})\ d\mathbf{u}$$

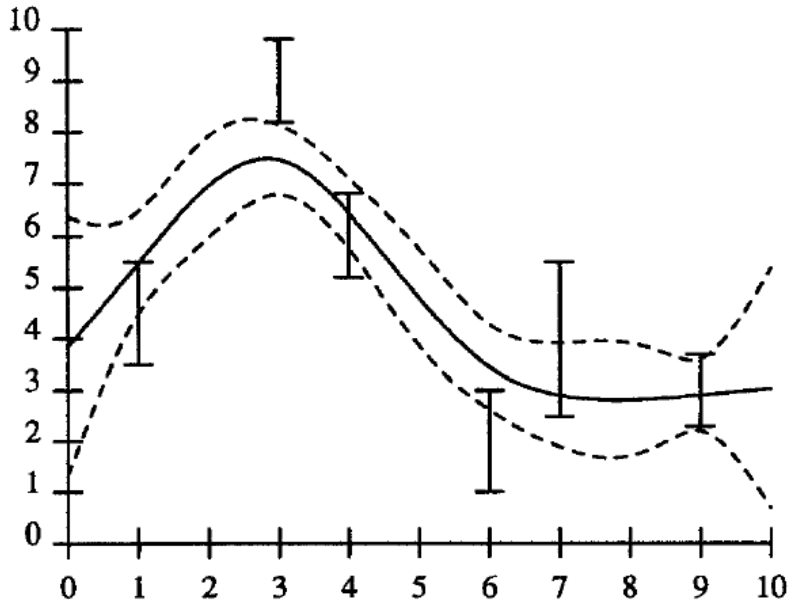

Fig. 8. Cubic spline with confidence interval. The vertical lines are the error bars (1 standard deviation) around the data points, and the dashed lines are the confidence interval for the whole curve.

Uncertainty

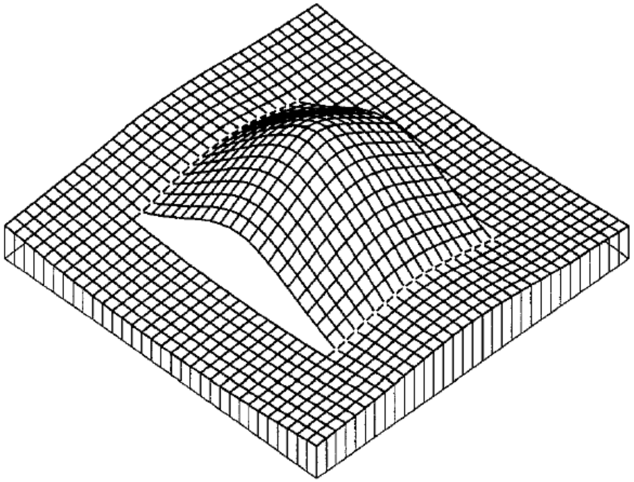

MAP estimate

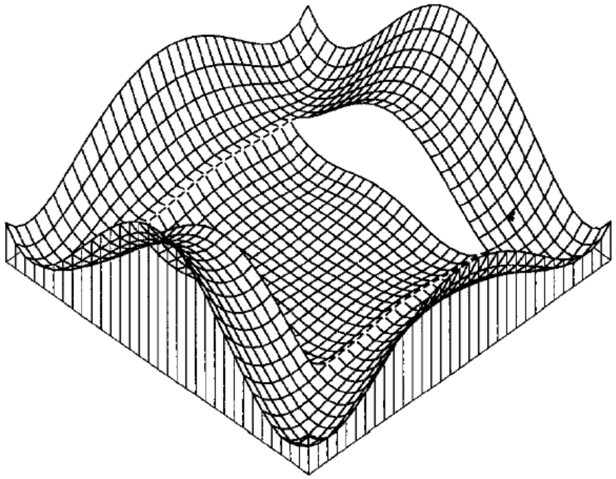

Variance field

Sample

- Analyze models by generating samples
- Gibbs sampler
  - Update each state variable by randomly picking a value from local Gibbs distribution

$$p(u_i|\mathbf{u}) = \frac{1}{Z_i} \exp(-E_p(u_i|\mathbf{u})/T_p)$$

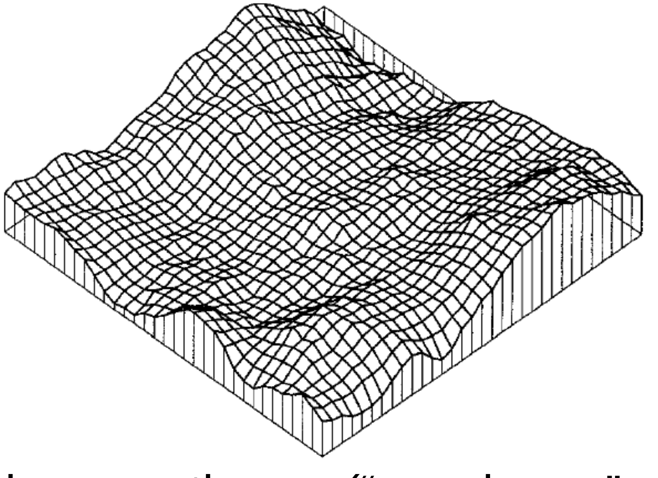

1st order smoothness ("membrane" model)

Sample

Use sample to estimate uncertainty by Monte Carlo method

$$\text{var}(u_i) = \sigma_i^2 = \int (u_i - u_i^*)^2 \, p(\mathbf{u}|\mathbf{d}) \, d\mathbf{u}$$

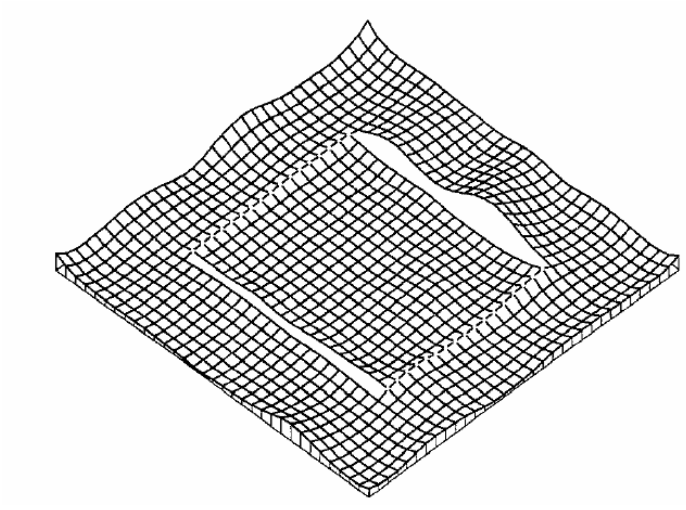

Single level, 1000 iterations

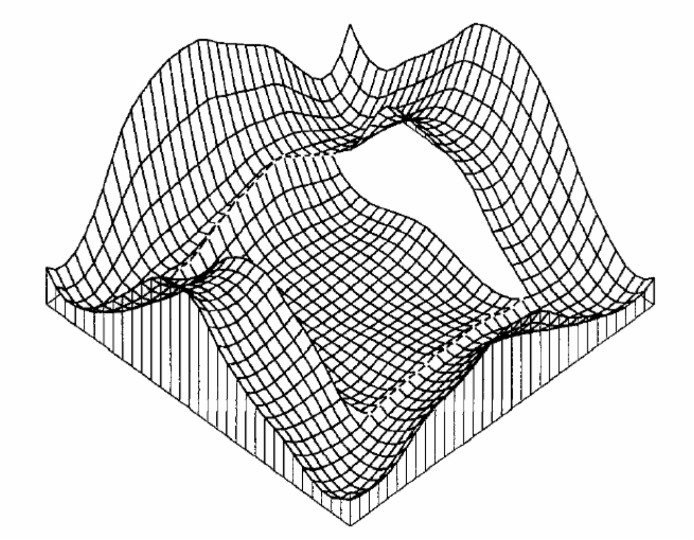

Multiresolution, 100 iterations

Dynamic models

Linear dynamic system …

Regularization parameter estimation

- How to balance data and prior term?

$$p(\mathbf{u}|\mathbf{d}) \propto \exp - [E_{\mathrm{d}}(\mathbf{u}, \mathbf{d}) + E_{\mathrm{p}}(\mathbf{u})/\sigma_{\mathrm{p}}^2]$$

- ML estimate

$$p(\mathbf{d}) = |2\pi(\mathbf{H}\mathbf{P}_0\mathbf{H}^T + \mathbf{R})|^{-1/2} \exp\left(-\frac{1}{2}\mathbf{d}^T(\mathbf{H}\mathbf{P}_0\mathbf{H}^T + \mathbf{R})^{-1}\mathbf{d}\right)$$

$$\mathbf{P}_0^{-1} = \sigma_{\mathrm{p}}^{-2}\mathbf{A}_{\mathrm{p}}$$

Learning model parameters

e.g. optical flow

Optical flow

Motion (displacement) of image pixels

“Army”

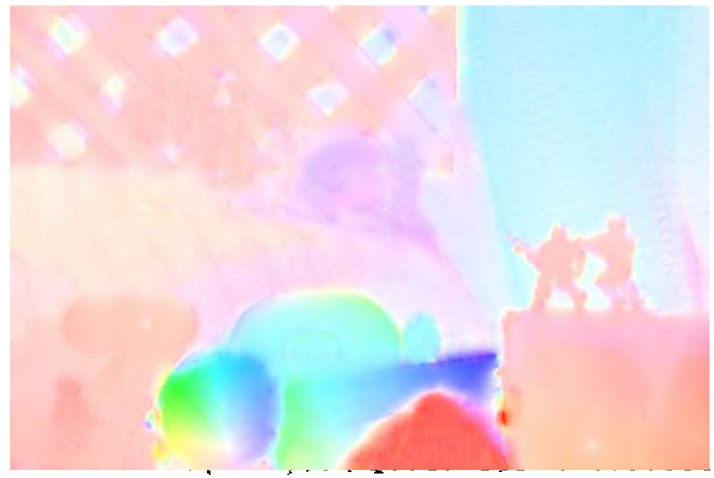

Horn & Schunck 1981

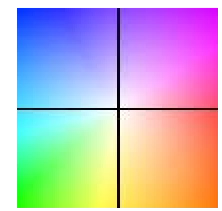

Key

Problems

- Ad-hoc choices, hand-tuned parameters
- Can we learn models from training data?

Image pair

Ground truth flow

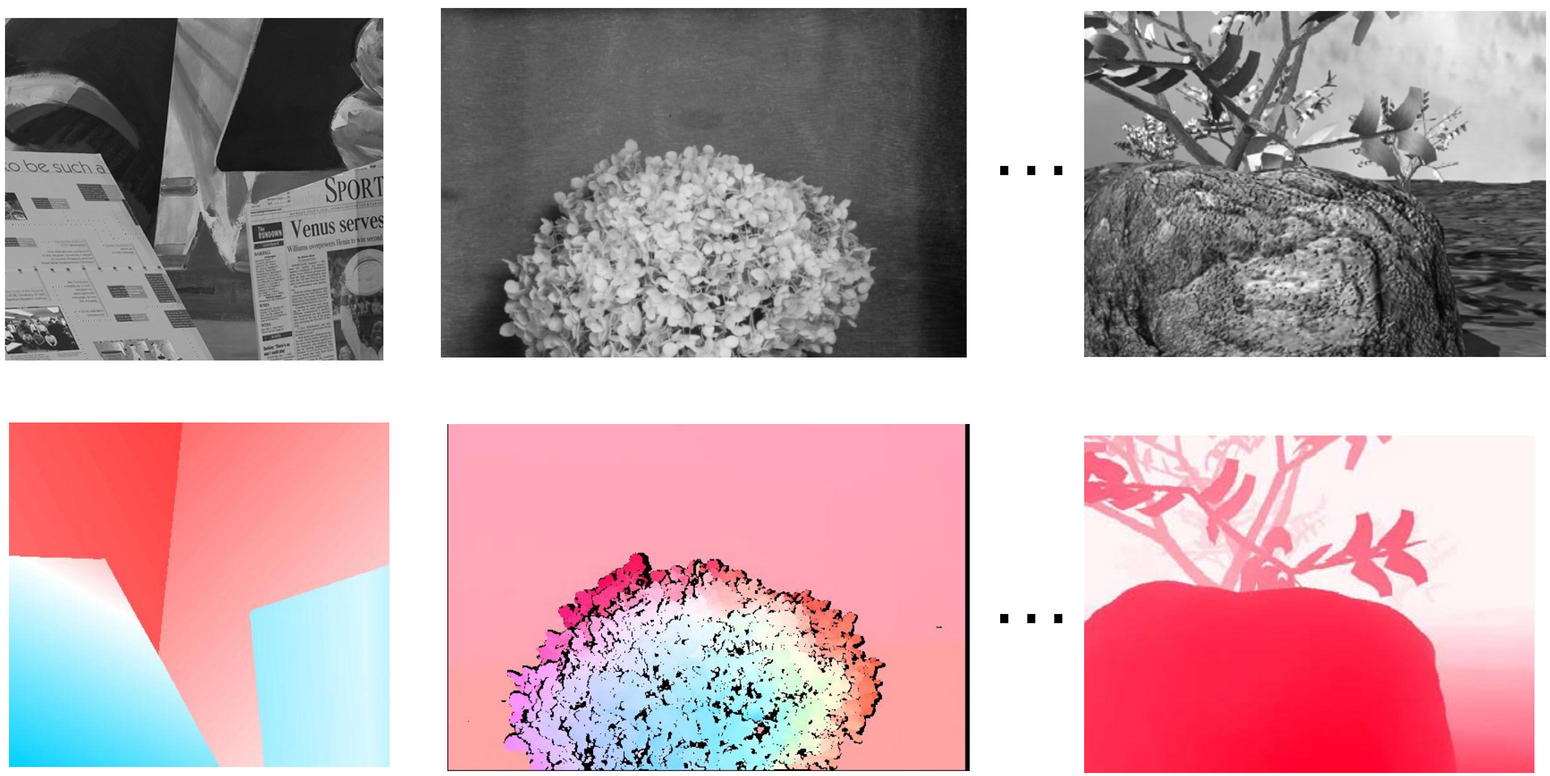

Middlebury optical flow benchmark (Baker et al. 2007)

Standard Bayesian formulation

$\mathbf{u}$: Horizontal flow $\mathbf{v}$: Vertical flow $\mathbf{I}_1$: First image $\mathbf{I}_2$: Second Image

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

Data term

How second image can be generated from first image and flow field

Spatial term

Prior knowledge of flow field

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

Brightness constancy (BC)

$$E_{\mathrm{D}}(\mathbf{u}, \mathbf{v}) = \sum_{(i,j)} \rho(I_1(i,j) - I_2(i + u_{ij}, j + v_{ij}))$$

- Penalty functions ρ (for linearized BC)
  - Quadratic (Horn & Schunck)
  - Charbonnier (Bruhn et al.)
  - Lorentzian (Black & Anandan)
- Which one should we use?

Brightness constancy

$$E_{\mathrm{D}}(\mathbf{u},\mathbf{v}) = \sum_{(i,j)} \rho(I_1(i,j) - I_2(i+u_{ij}, j+v_{ij}))$$

Probabilistic interpretation

$$p_{\mathrm{BC}}(\mathbf{I}_2|\mathbf{u},\mathbf{v},\mathbf{I}_1) \propto \exp\{-\sum_{(i,j)} \rho(I_1(i,j) - I_2(i+u_{ij}, j+v_{ij}))\}$$

$$\Downarrow$$

$$p_{\mathrm{BC}}(\mathbf{I}_2|\mathbf{u},\mathbf{v},\mathbf{I}_1) \propto \prod_{(i,j)} \phi(I_1(i,j) - I_2(i+u_{ij}, j+v_{ij}))$$

ϕ = potential function

Brightness constancy

$$p_{\mathrm{BC}}(\mathbf{I}_2 | \mathbf{u}, \mathbf{v}, \mathbf{I}_1) \propto \prod_{(i,j)} \phi(I_1(i,j) - I_2(i + u_{ij}, j + v_{ij}))$$

Histogram from training data (images + ground truth flow)

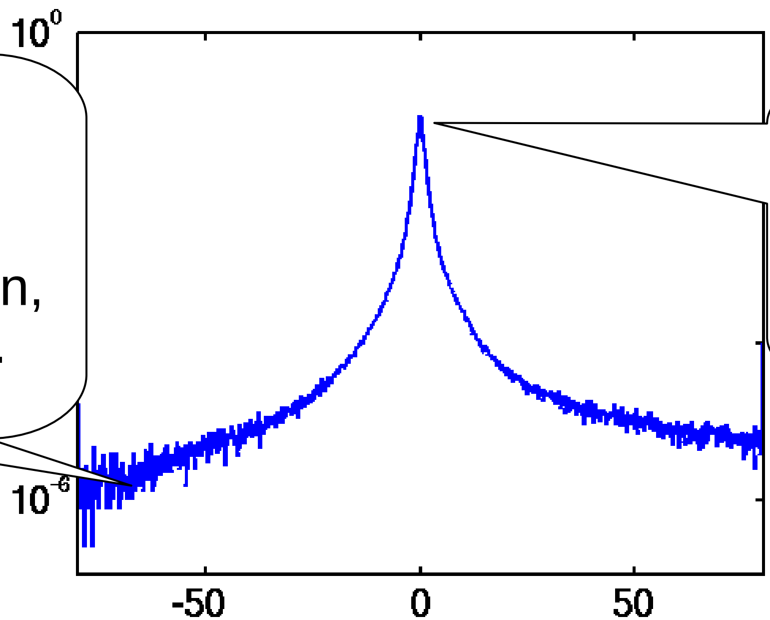

$$I_1(i,j) - I_2(i + u_{ij}, j + v_{ij})$$

Heavy tails

violations due to reflection, occlusion, transparency etc.

Sharp peak

brightness constancy usually holds

Brightness constancy

Idea: Fit the histogram (Gaussian Scale Mixture)

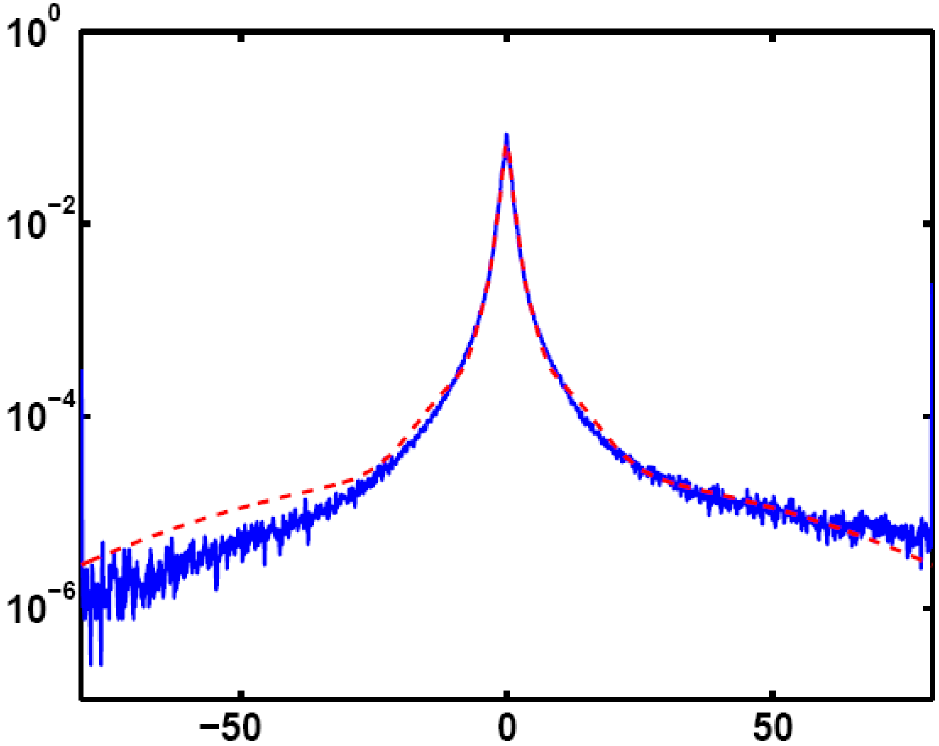

$$I_1(i,j) - I_2(i+u_{ij}, j+v_{ij})$$

Spatial term

- Smoothness expressed by canonical derivative

$$E_{\mathrm{Su}}(\mathbf{u}) = \sum_{(i,j)} \rho(u_{i,j+1} - u_{ij}) + \rho(u_{i+1,j} - u_{ij})$$

Penalty functions and parameters?

- Probabilistic interpretation: Pairwise MRF (PW)

$$p_{\mathrm{PW}}(\mathbf{u}) = \frac{1}{Z} \prod_{(i,j)} \phi(u_{i,j+1} - u_{ij}) \cdot \phi(u_{i+1,j} - u_{ij})$$

Spatial term

Pairwise MRF (PW) model

$$p_{\mathrm{PW}}(\mathbf{u}) = \frac{1}{Z} \prod_{(i,j)} \phi(u_{i,j+1} - u_{ij}) \cdot \phi(u_{i+1,j} - u_{ij})$$

Approximate ML learning by contrastive divergence (CD) algorithm (Hinton2000)

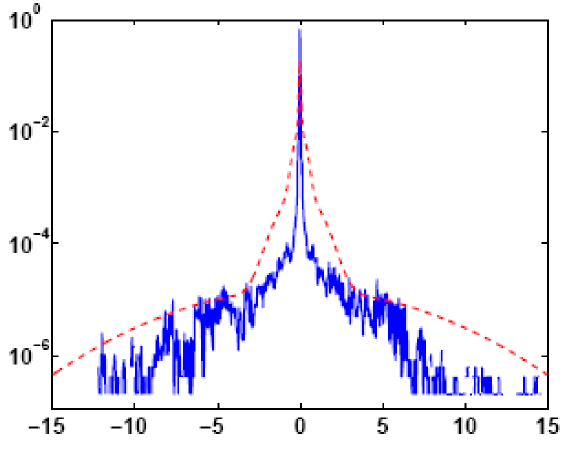

$\partial_x \mathbf{u}$

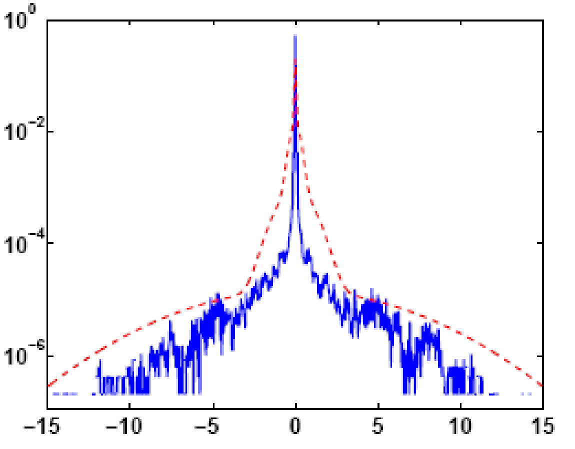

$\partial_y \mathbf{u}$

Conclusion

- Bayesian modeling encompasses regularization
- It also provides
  - Uncertainty modeling (2nd order statistics)
  - Sample to validate model
  - Optimal estimation (loss function)
  - Regularization parameter estimation
  - Learning model parameters