



Chapter 4

Variational Image Restoration

Variational Methods for Computer Vision
Winter 2013/14

Inverse Problems and
Image Restoration

Image Denoising

Image Deblurring

Inverse Problems and
Bayesian Inference

Motion Blur and
Defocus Blur

Video Super
Resolution

Prof. Daniel Cremers
Chair for Computer Vision and Pattern Recognition
Department of Computer Science
Technische Universität München



- 1 Inverse Problems and Image Restoration
- 2 Image Denoising
- 3 Image Deblurring
- 4 Inverse Problems and Bayesian Inference
- 5 Motion Blur and Defocus Blur
- 6 Video Super Resolution

Inverse Problems and
Image Restoration

Image Denoising

Image Deblurring

Inverse Problems and
Bayesian Inference

Motion Blur and
Defocus Blur

Video Super
Resolution

Inverse Problems, Ill-Posedness and Regularization



In mathematics, the conversion of measurement data into information about the observed object or the observed physical system is referred to as an **inverse problem**.

Following **Hadamard (1902)**, a mathematical problem is called **well-posed** iff:

- 1 A solution exists.
- 2 The solution is unique.
- 3 The solution's behavior changes continuously with the initial conditions.

Inverse problems are often **ill-posed**. Since the measurement data is often not sufficient to uniquely characterize the observed object or system, one introduces **prior knowledge** to disambiguate which solutions are a priori more likely. In the context of variational methods this prior knowledge gives rise to the **regularity term**.

Image Restoration: Denoising

Image restoration is a classical inverse problem: Given an observed image $f : \Omega \rightarrow \mathbb{R}$ and a (typically stochastic) model of an image degradation process, we want to restore the original image $u : \Omega \rightarrow \mathbb{R}$.

Image denoising is an example of image restoration where we assume that the true image u is corrupted by (additive) noise:

$$f = u + \eta, \quad \eta \sim \mathcal{N}(0, \sigma).$$

Assuming that u is smooth, one can compute a denoised image by minimizing (Rudin, Osher, Fatemi '92):

$$\min_u \frac{1}{2} \int |u - f|^2 dx + \int |\nabla u| dx.$$

This gives rise to the **Euler-Lagrange equation**

$$u - f - \operatorname{div} \left(\frac{\nabla u}{|\nabla u|} \right) = 0.$$

Other noise models and regularizers are conceivable.



Image Restoration: Denoising



original



noisy



denoised

(Goldlücke, Strekalovskiy, Cremers, SIAM J. Imaging Sci. '12)



Image Restoration: Deblurring

A prototypical **blur model** is given by

$$f = A * u + \eta \quad \eta \sim \mathcal{N}(0, \sigma),$$

with a blur kernel A .

In a variational setting, this process can be inverted by minimizing the **TV deblurring functional**:

$$\min_u \frac{1}{2} \int |A * u - f|^2 dx + \int |\nabla u| dx.$$

For symmetric kernels A , the **Euler-Lagrange equation** is given by:

$$A * (A * u - f) - \operatorname{div} \left(\frac{\nabla u}{|\nabla u|} \right) = 0,$$

and the **gradient descent equation**


$$\frac{\partial u}{\partial t} = -A * (A * u - f) + \operatorname{div} \left(\frac{\nabla u}{|\nabla u|} \right).$$



Image Restoration: Deblurring



Original



blurred and noisy



deblurred

(Source: Goldluecke, Cremers, ICCV 2011)

Inverse Problems and
Image Restoration

Image Denoising

Image Deblurring

Inverse Problems and
Bayesian Inference

Motion Blur and
Defocus Blur

Video Super
Resolution

Inverse Problems and Bayesian Inference

The framework of **Bayesian inference** allows to systematically derive functionals for different image formation models.

Let u be the unknown true image and f the observed one, then we can write the joint probability for u and f as:

$$\mathcal{P}(u, f) = \mathcal{P}(u|f) \mathcal{P}(f) = \mathcal{P}(f|u) \mathcal{P}(u).$$

Rewriting this expression we obtain the **Bayesian formula** (Thomas Bayes 1887):

$$\mathcal{P}(u|f) = \frac{\mathcal{P}(f|u) \mathcal{P}(u)}{\mathcal{P}(f)}.$$

Maximum A posteriori (MAP) estimation aims at computing the most likely solution \hat{u} given f by maximizing the **posterior probability** $\mathcal{P}(u|f)$

$$\hat{u} = \arg \max_u \mathcal{P}(u|f) = \arg \max_u \mathcal{P}(f|u) \mathcal{P}(u).$$

$\mathcal{P}(f|u)$ is called the **likelihood** and $\mathcal{P}(u)$ the **prior**.



MAP Estimation in the Discrete Setting

Let us assume n independent pixels. For each the measured intensity f_i is given by the true intensity u_i plus additive Gaussian noise. This corresponds to the likelihood

$$\mathcal{P}(f_i|u_i) \propto \exp\left(-\frac{(u_i - f_i)^2}{2\sigma^2}\right).$$

Since all measurements are mutually independent, we obtain for the entire vector $f = (f_1, \dots, f_n)$ of pixel intensities:

$$\mathcal{P}(f|u) = \prod_{i=1}^n \mathcal{P}(f_i|u) = \prod_{i=1}^n \mathcal{P}(f_i|u_i) \propto \prod_{i=1}^n \exp\left(-\frac{(u_i - f_i)^2}{2\sigma^2}\right).$$

We now expand the prior:

$$\mathcal{P}(u) = \mathcal{P}(u_1 \dots u_n) = \mathcal{P}(u_1|u_2 \dots u_n)\mathcal{P}(u_2 \dots u_n) \propto \prod_{i=1}^{n-1} \mathcal{P}(u_i|u_{i+1}),$$

where we assumed a **Markov property**, namely that the probability of u_i is sufficiently characterized by its neighbor.



MAP Estimation in the Discrete Setting

Assuming a simple **smoothness prior**, we have:

$$\mathcal{P}(u) \propto \prod_{i=1}^{n-1} \mathcal{P}(u_i | u_{i+1}) \propto \prod_{i=1}^{n-1} \exp(-\lambda |u_i - u_{i+1}|).$$

With these assumptions, the **posterior distribution** is given by:

$$\mathcal{P}(u|f) \propto \prod_{i=1}^n \exp\left(-\frac{|f_i - u_i|^2}{2\sigma^2}\right) \prod_{i=1}^{n-1} \exp(-\lambda |u_i - u_{i+1}|)$$

Rather than maximizing this probability distribution, one can equivalently **minimize its negative logarithm** (because the logarithm is strictly monotonous).

It is given by the **energy**

$$E(u) = -\log \mathcal{P}(u|f) = \sum_{i=1}^n \frac{|f_i - u_i|^2}{2\sigma^2} + \lambda \sum_{i=1}^{n-1} |u_i - u_{i+1}| + \text{const.}$$



MAP Estimation in the Continuous Setting

Similarly one can define **Bayesian MAP optimization in the continuous setting**, where the likelihood is given by:

$$\mathcal{P}(f|u) \propto \exp\left(-\int \frac{|f(x) - u(x)|^2}{2\sigma^2} dx\right),$$

and the prior is given by

$$\mathcal{P}(u) \propto \exp\left(-\lambda \int |\nabla u(x)| dx\right).$$

Thus the data term in variational methods corresponds to the likelihood, whereas the regularizer corresponds to the prior:

$$E(u) = -\log \mathcal{P}(u|f) = \int \frac{|f(x) - u(x)|^2}{2\sigma^2} dx + \lambda \int |\nabla u(x)| dx + \text{const.}$$

A systematic derivation of **probability distributions on infinite-dimensional spaces** requires a more formal derivation (introduction of measures etc). This is beyond our scope.



Image Restoration: Motion Blur

Assume the camera lens opens instantly and remains open during the time interval $[0, T]$ in which the camera moves with constant velocity V in x -direction. The observed brightness is

$$g(x, y) = \frac{1}{T} \int_0^T f(x - Vt, y) dt.$$

Inserting $x' \equiv Vt$, we get a convolution

$$g(x, y) = \frac{1}{VT} \int_0^{VT} f(x-x', y) dx' = \int_{-\infty}^{\infty} f(x-x', y-y') h(x', y') dx' dy',$$

with the anisotropic blur kernel:

$$h(x', y') = \frac{1}{VT} \cdot \delta(y') \cdot \chi_{[0, VT]}(x'),$$

and


$$\chi_{[a,b]}(x') = \begin{cases} 1, & \text{if } x' \in [a, b] \\ 0, & \text{else} \end{cases} \quad (\text{box filter})$$



Example: Motion Blur



Original



Motion-blurred

(Author: D. Cremers)



Inverse Problems and
Image Restoration

Image Denoising

Image Deblurring

Inverse Problems and
Bayesian Inference

Motion Blur and
Defocus Blur

Video Super
Resolution

Image Restoration: Defocus Blur

Defocus blur arises with real (in particular thick) lenses because structures are increasingly blurred, the further they are from the focal plane.

Depending on the focal setting and the depth of the scene, we will therefore observe a **space-varying blur** which allows us to infer the local depth (**shape from focus / defocus**).



Scene captured with three different focal settings.

(Source: Favaro, Soatto, PAMI 2005)



Image Restoration: Super Resolution

Super resolution from video exploits the redundancy available in multiple images. We assume that each image f_i is a blurred and downsampled version of a high-resolution scene.

We can try to recover a high-resolution image u with a variational approach of the form:

$$\min_u \sum_{i=1}^n \int |A(u \circ w_i) - f_i| dx + \lambda \int |\nabla u| dx.$$

The deformation field $w_i : \Omega \rightarrow \Omega$ models the warping from the original scene into image i , and A is a linear operator modeling the blurring and downsampling. Again, the variational approach aims at inverting an image formation process of the form:

$$f_i = A(u \circ w_i) + \eta,$$

which states that the observed image is obtained from the “true” image by nonrigid deformation, blurring and downsampling plus additive Poisson-distributed noise η .



Image Restoration: Super Resolution



One of several input images



Superresolution estimate

(Schoenemann, Cremers, IEEE T. on Image Processing 2012)

Inverse Problems and
Image Restoration

Image Denoising

Image Deblurring

Inverse Problems and
Bayesian Inference

Motion Blur and
Defocus Blur

Video Super
Resolution