CAMBRIDGE

Cambridge University Press
978-1-107-04436-4 - Motion Deblurring: Algorithms and Systems
Edited by A. N. Rajagopalan and Rama Chellappa
Excerpt
More information

1 Mathematical models and practical solvers for uniform motion deblurring

Jiaya Jia

Recovering an un-blurred image from a single motion-blurred picture has long been a fundamental research problem. If one assumes that the blur kernel – or point spread function (PSF) – is shift invariant, the problem reduces to that of image deconvolution. Image deconvolution can be further categorized as non-blind and blind.

In non-blind deconvolution, the motion blur kernel is assumed to be known or computed elsewhere; the task is to estimate the un-blurred latent image. The general problems to address in non-blind deconvolution include reducing possible unpleasant ringing artifacts that appear near strong edges, suppressing noise, and saving computation. Traditional methods such as Wiener deconvolution (Wiener 1949) and the Richardson–Lucy (RL) method (Richardson 1972, Lucy 1974) were proposed decades ago and find many variants thanks to their simplicity and efficiency. Recent developments involve new models with sparse regularization and the proposal of effective linear and non-linear optimization to improve result quality and further reduce running time.

Blind deconvolution is a much more challenging problem, since both the blur kernel and the latent image are unknown. One can regard non-blind deconvolution as an inevitable step in blind deconvolution during the course of PSF estimation or after the PSF has been computed. Both blind and non-blind deconvolution are practically very useful; they are studied and employed in a variety of disciplines including, but not limited to, image processing, computer vision, medical and astronomic imaging, and digital communication.

This chapter discusses shift-invariant single image motion deblurring methods, which assume that the image is uniformly blurred with only one PSF, which may not be known a priori. This set of problems has a long history in theoretical and empirical research and has notably advanced in the last 5–10 years with a few remarkably effective models and solvers.

1.1 Non-blind deconvolution

Ideally, a blur observation is modeled as a linearly filtered version of the latent un-blurred signal. This process can be expressed as

$$\mathbf{b} = \mathbf{l} \otimes \mathbf{f}, \tag{1.1}$$

© in this web service Cambridge University Press www.cambridge.org

CAMBRIDGE

Cambridge University Press
978-1-107-04436-4 - Motion Deblurring: Algorithms and Systems
Edited by A. N. Rajagopalan and Rama Chellappa
Excerpt
More information

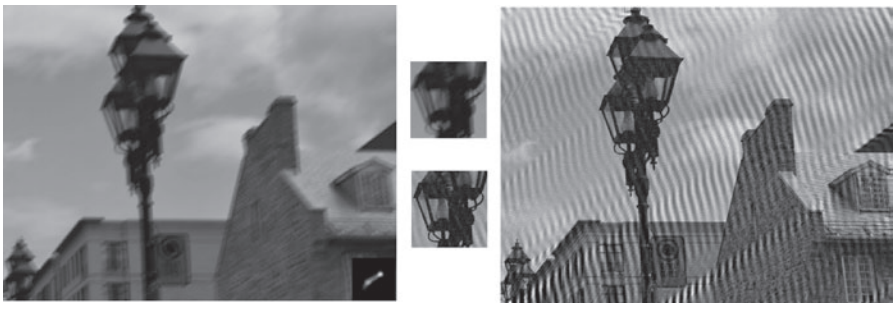

(a) (b) (c)

Figure 1.1 Inverse filtering problem. Visual artifacts caused by inverse filter. (a) Blurred image and PSF; (b) close-ups; (c) output of inverse filtering.

where **b**, **l** and **f** are the blurred image, latent un-blurred image, and PSF (or blur kernel), respectively. In the frequency domain,

$$\mathcal{F}(\mathbf{b}) = \mathcal{F}(\mathbf{l}) \cdot \mathcal{F}(\mathbf{f}), \tag{1.2}$$

where $\mathcal{F}$ is the Fourier transform.

If $\mathcal{F}(\mathbf{f})$ does not contain zero or very small values and the blurred image is noise-free, the latent image **l** can be obtained simply by inverting the convolution process using inverse filtering, the simplest method that solves for **l**. This process is expressed as

$$\mathcal{F}(\mathbf{l}) = \mathcal{F}(\mathbf{b})/\mathcal{F}(\mathbf{f}). \tag{1.3}$$

This strategy practically may produce severe visual artifacts, such as ringing, for the following reasons. First, the inversion of **f** may not exist, especially for low-pass filters. Second, motion PSFs caused by object or camera motion are typically band-limited and their spectra have zero or near-zero values at high frequency. Third, image formation causes problems including image noise, quantization error, color saturation, and non-linear camera response curve. They make blur violate the ideal convolution model and lead to a more flexible form

$$\mathbf{b} = \mathbf{l} \otimes \mathbf{f} + \mathbf{n}, \tag{1.4}$$

where **n** denotes error in the blurred image, which we call image noise in general. One deconvolution result by direct inverse filter is shown in Figure 1.1.

Development of more advanced non-blind deconvolution methods dates back to the 1970s. Early representative approaches include Wiener Deconvolution (Wiener 1949), Least Square Filtering (Miller 1970, Hunt 1973, Tikhonov & Arsenin 1977), Richardson–Lucy method (Richardson 1972, Lucy 1974), and recursive Kalman Filtering (Woods & Ingle 1981). Readers are referred to Andrews & Hunt (1977) for a review of these early approaches.

Simply put, many algorithms minimize an energy consisting of two terms, i.e. the *data* term E_{data} (corresponding to *likelihood* in probability) and *regularization* (also known as *prior*) E_{prior}. E_{data} measures the difference between the convolved image and

© in this web service Cambridge University Press www.cambridge.org

CAMBRIDGE

Cambridge University Press
978-1-107-04436-4 - Motion Deblurring: Algorithms and Systems
Edited by A. N. Rajagopalan and Rama Chellappa
Excerpt
More information

the blur observation, and is written as

$$E_{\text{data}} = \Phi(\mathbf{l} \otimes \mathbf{f} - \mathbf{b}), \tag{1.5}$$

where Φ is a distance function. A common definition is $\Phi(\cdot) = \|\cdot\|^2$ (Wiener 1949), representing the L_2-norm of all elements. It is also called the Gaussian likelihood. E_{prior} is denoted as a function $\Psi(\mathbf{l})$, which has different specifications in existing approaches. Given E_{data} and E_{prior}, the latent image $\mathbf{l}$ can be estimated by minimizing the energy incorporating these two terms, expressed as

$$\min_{\mathbf{l}} \|\mathbf{l} \otimes \mathbf{f} - \mathbf{b}\|^2 + \lambda \Psi(\mathbf{l}), \tag{1.6}$$

where λ is a weight. In what follows, we discuss a few representative non-blind deconvolution methods with respect to model design and solver construction. Their respective strengths, disadvantages, and relations are also presented.

1.1.1 Regularized approaches

A number of early methods incorporated square regularization constraints. Two representative forms are $\Psi(\mathbf{l}) = \|\mathbf{l}\|^2$ and $\Psi(\mathbf{l}) = \|\nabla\mathbf{l}\|^2$, where ∇ is the gradient operator. They enforce smoothness on image values and image gradients, and are called Tikhonov and Gaussian regularizers, respectively. Substituting them into Eq. (1.6) yields

$$\min_{\mathbf{l}} \|\mathbf{l} \otimes \mathbf{f} - \mathbf{b}\|^2 + \lambda \|\mathbf{l}\|^2 \tag{1.7}$$

and

$$\min_{\mathbf{l}} \|\mathbf{l} \otimes \mathbf{f} - \mathbf{b}\|^2 + \lambda \|\nabla\mathbf{l}\|^2, \tag{1.8}$$

for overall energy minimization. Weight λ is typically a small value.

The main advantage of these constrained least square methods is in the simplicity of formation, which results in a solver similar to an inverse filter. Taking the Tikhonov method as an example, there exists a closed form solution $\mathbf{l}^*$ for Eq. (1.7) by setting its first order derivative to zero with respect to $\mathbf{l}$. Rearranging Eq. (1.7) in a matrix form and denoting by E the total energy yield

$$\begin{aligned} E &= \|F\nu(\mathbf{l}) - \nu(\mathbf{b})\|^2 + \lambda\|\nu(\mathbf{l})\|^2 \\ &= \nu(\mathbf{l})^T F^T F \nu(\mathbf{l}) - 2\nu(\mathbf{b})^T F \nu(\mathbf{l}) + \nu(\mathbf{b})^T \nu(\mathbf{b}) + \lambda \nu(\mathbf{l})^T \nu(\mathbf{l}), \end{aligned}$$

where F is a sparse convolution matrix generated from $\mathbf{f}$, and ν is the operator that transforms the image into its vector form. The partial derivative is

$$\frac{\partial E}{\partial \nu(\mathbf{l})} = 2F^T F \nu(\mathbf{l}) - 2F^T \nu(\mathbf{b}) + 2\lambda\nu(\mathbf{l}). \tag{1.9}$$

By setting the above equation to zero, the optimal solution $\mathbf{l}^*$ is

$$\nu(\mathbf{l}^*) = \frac{F^T}{F^T F + \lambda\Lambda}\nu(\mathbf{b}), \tag{1.10}$$

where Λ is an identity matrix, the same size as $F^T F$.

© in this web service Cambridge University Press www.cambridge.org

CAMBRIDGE

Cambridge University Press
978-1-107-04436-4 - Motion Deblurring: Algorithms and Systems
Edited by A. N. Rajagopalan and Rama Chellappa
Excerpt
More information

Regularization bias

If there is neither kernel error nor image noise and the kernel matrix F is invertible, the ground truth latent image $\hat{\mathbf{l}}$ is simply the reversion of convolution, expressed as

$$\nu(\hat{\mathbf{l}}) = F^{-1}\nu(\mathbf{b}) = \frac{F^T\nu(\mathbf{b})}{F^TF}. \tag{1.11}$$

The difference between Eqs. (1.10) and (1.11) makes it possible to analyze how the regularization term introduces bias in deconvolution in an ideal noise-free situation. It serves as guidance for future deconvolution model design.

We denote the error map of the recovered image as

$$\delta\mathbf{l} = \mathbf{l}^* - \hat{\mathbf{l}}, \tag{1.12}$$

where $\delta\mathbf{l}$ is the error introduced in deconvolution. Equations (1.11) and (1.12) together lead to

$$\nu(\delta\mathbf{l}) = \nu(\mathbf{l}^*) - \nu(\hat{\mathbf{l}}) = -\frac{\lambda\nu(\mathbf{b})}{(F^TF+\lambda\Lambda)F} = -\frac{\lambda}{F^TF+\lambda\Lambda}\nu(\hat{\mathbf{l}}). \tag{1.13}$$

Because $\lambda/(F^TF+\lambda\Lambda)$ can be regarded as a weight fixed by blur, this equation indicates that $\delta\mathbf{l}$ generally appears as a high frequency map dependent on image structures in $\nu(\hat{\mathbf{l}})$ as shown in Figure 1.2(d). Intuitively, a large λ makes the result lose detail.

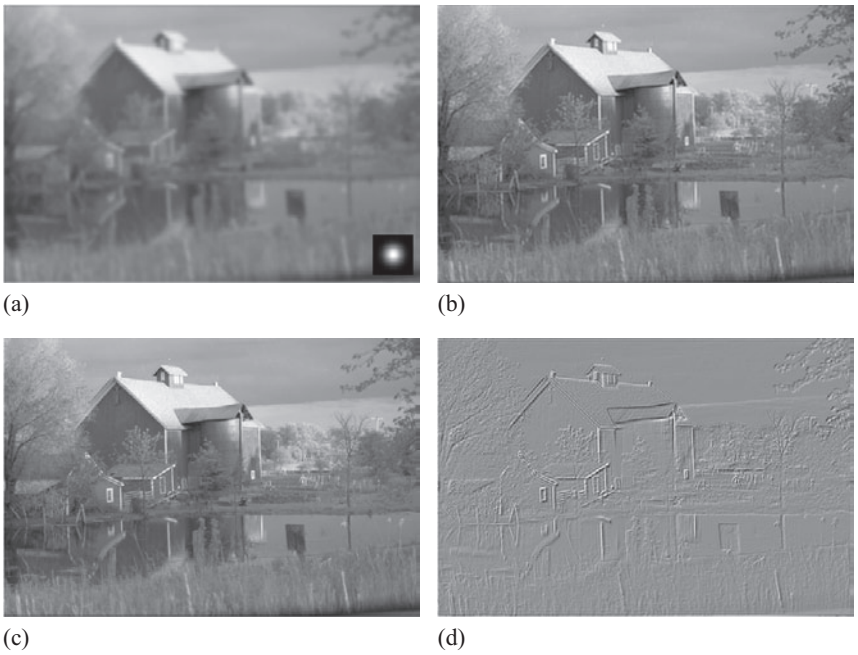

(a) (b) (c) (d)

Figure 1.2 Error introduced with Tikhonov regularization. (a) Blurred image with the ground truth PSF; (b) deblurred image with Tikhonov regularization; (c) ground truth latent image; (d) map of $\delta\mathbf{l}$ computed using Eq. (1.13).

© in this web service Cambridge University Press www.cambridge.org

CAMBRIDGE

Cambridge University Press
978-1-107-04436-4 - Motion Deblurring: Algorithms and Systems
Edited by A. N. Rajagopalan and Rama Chellappa
Excerpt
More information

(a) (b) (c)

Figure 1.3 Influence of image noise in deconvolution using Tikhonov regularization. (a) Input images with additive noise; (b) deconvolution results; artifacts are primarily the amplified noise; (c) difference maps between (b) and the ground truth un-blurred noise-free images.

If we consider inevitable image noise and PSF error, the Tikhonov regularizer actually enhances the stability of deconvolution, discussed below.

Noise amplification

Now consider the case that image noise $\delta\mathbf{b}$ is present, which is common in natural images. With a derivation similar to Eq. (1.9), which takes derivatives and sets them to zero, the expression obtained is

$$\nu(\delta\mathbf{l}) = \frac{F^T\nu(\delta\mathbf{b})}{F^TF+\lambda\Lambda} + \frac{-\lambda\nu(\hat{\mathbf{l}})}{F^TF+\lambda\Lambda}. \tag{1.14}$$

We have explained the second term given the same expression in Eq. (1.13). It produces a map that in general contains high frequency structures.

In the first term, setting $\kappa = F^T/(F^TF+\lambda\Lambda)$ to represent a coefficient matrix, the expression simplifies to $\kappa\nu(\delta\mathbf{b})$. It actually functions as adding noise with a ratio κ, which yields noisy results. Summing up the effects of the two terms in Eq. (1.14), it can be concluded that the deconvolution results contain noise while lacking an amount of structural detail compared to the ground truth image. Two examples are shown in Figure 1.3.

Relation to Wiener deconvolution

Wiener filter is a method which has been widely used in non-blind deconvolution (Wiener 1949). Its specialty is in the use of image and noise power spectra to suppress noise, expressed as

$$\mathcal{F}(\mathbf{l}) = \frac{1}{\mathcal{F}(\mathbf{f})}\left(\frac{|\mathcal{F}(\mathbf{f})|^2}{|\mathcal{F}(\mathbf{f})|^2+\frac{1}{\mathrm{SNR}(\mathbf{f})}}\right)\cdot\mathcal{F}(\mathbf{b}), \tag{1.15}$$

where $\mathrm{SNR}(\mathbf{f})$ is the signal-to-noise ratio and $|\mathcal{F}(\mathbf{f})|^2$ denotes autocorrelation.

© in this web service Cambridge University Press www.cambridge.org

CAMBRIDGE

Cambridge University Press
978-1-107-04436-4 - Motion Deblurring: Algorithms and Systems
Edited by A. N. Rajagopalan and Rama Chellappa
Excerpt
More information

It can be proven that the Tikhonov regularized method is equivalent to the Wiener filter with a proper λ. First, Eq. (1.10) can be rewritten as

$$(F^T F + \lambda\Lambda)\nu(\mathbf{l}) = F^T \nu(\mathbf{b}). \tag{1.16}$$

Because it holds that

$$\begin{aligned} F\nu(\mathbf{l}) &= \nu(\mathbf{f} \otimes \mathbf{l}), \\ F^T\nu(\mathbf{l}) &= \nu(\mathbf{f} \oplus \mathbf{l}) = \nu(\mathbf{f}' \otimes \mathbf{l}), \\ \mathcal{F}(\mathbf{f}) \cdot \mathcal{F}(\mathbf{f}') &= |\mathcal{F}(\mathbf{f})|^2, \end{aligned}$$

where $\mathbf{f}'$ is the flipped version of $\mathbf{f}$, $\oplus$ denotes correlation, and $\cdot$ is an element-wise multiplication operator. Equation (1.16) finds the solution in image domain as

$$\nu(\mathbf{f}' \otimes (\mathbf{f} \otimes \mathbf{l})) + \lambda\mathbf{l} = \nu(\mathbf{f}' \otimes \mathbf{b}). \tag{1.17}$$

Taking the Fourier transform on both sides of Eq. (1.17) yields

$$\mathcal{F}(\mathbf{f}') \cdot \mathcal{F}(\mathbf{f}) \cdot \mathcal{F}(\mathbf{l}) + \lambda\mathcal{F}(\mathbf{l}) = \mathcal{F}(\mathbf{f}') \cdot \mathcal{F}(\mathbf{b}),$$

which can be further expressed as

$$\mathbf{l} = \mathcal{F}^{-1}\left(\frac{\mathcal{F}(\mathbf{f}') \cdot \mathcal{F}(\mathbf{b})}{\mathcal{F}(\mathbf{f}') \cdot \mathcal{F}(\mathbf{f}) + \lambda\Lambda}\right). \tag{1.18}$$

Equation (1.18) is the same as Eq. (1.15) when $\lambda\Lambda = 1/(\mathrm{SNR}(\mathbf{f}))$. This equivalence implies that Wiener deconvolution has similar noise amplification and structure information loss properties to Tikhonov regularized deconvolution.

1.1.2 Iterative approaches

Iterative computation was also used in several methods. The van Cittert (van Cittert 1931) solver can be applied to iteratively estimate the deconvolved image as

$$\mathbf{l}^{t+1} = \mathbf{l}^t + \beta(\mathbf{b} - \mathbf{l}^t \otimes \mathbf{f}), \tag{1.19}$$

where β is a parameter, adjustable automatically or manually, controlling the convergence speed while t and $t+1$ index iterations. Equation (1.19) converges ideally to a result close to that produced by the inverse filter expressed in Eq. (1.3), which does not incorporate any prior or regularization.

The widely employed Richardson–Lucy (RL) deconvolution (Richardson 1972, Lucy 1974) can be expressed as

$$\mathbf{l}^{t+1} = \mathbf{l}^t\left(\mathbf{f}' \otimes \left(\frac{\mathbf{b}}{\mathbf{l}^t \otimes \mathbf{f}}\right)\right), \tag{1.20}$$

where $\mathbf{f}'$ is the flipped version of $\mathbf{f}$, used in correlation instead of convolution. How Eq. (1.20) is constructed is explained in quite a number of papers and tutorials available online, and is thus omitted here. Different from direct inversion (Eq. (1.3)), RL deconvolution is iterative and can be stopped halfway, which *empirically* alleviates, in part, noise amplification. Performing it for many iterations or making it converge,

© in this web service Cambridge University Press www.cambridge.org

CAMBRIDGE

Cambridge University Press
978-1-107-04436-4 - Motion Deblurring: Algorithms and Systems
Edited by A. N. Rajagopalan and Rama Chellappa
Excerpt
More information

contrarily, could yield less satisfactory results. The following derivation shows that the RL method is equivalent to the Poisson maximum likelihood, without imposing any image or kernel prior.

When assuming independent and identically distributed (i.i.d.) Gaussian noise $\mathbf{n} = \mathbf{b} - \mathbf{l} \otimes \mathbf{f}$, the maximum likelihood estimation of $\mathbf{l}$ is generally expressed as

$$p(\mathbf{b}|\mathbf{l}) \propto \prod_i \exp\left(-\frac{(\mathbf{b}_i - (\mathbf{l} \otimes \mathbf{f})_i)^2}{2\sigma^2}\right), \tag{1.21}$$

where $p(\mathbf{b}|\mathbf{l})$ is the conditional probability (also known as likelihood), i indexes pixels, and σ^2 is the Gaussian variance. Similarly, assuming that noise $\mathbf{n} = \mathbf{b} - \mathbf{l} \otimes \mathbf{f}$ follows a Poisson distribution, yields

$$p(\mathbf{b}|\mathbf{bl}) \propto \prod_i \frac{(\mathbf{l} \otimes \mathbf{f})_i^{\mathbf{b}_i} \exp(-(\mathbf{l} \otimes \mathbf{f})_i)}{\mathbf{b}_i!}, \tag{1.22}$$

where i indexes pixels. Its logarithmic energy is

$$\log\left(p(\mathbf{b}|\mathbf{l})\right) \propto \sum_i \left(\mathbf{b}_i \log\left(\mathbf{l} \otimes \mathbf{f}\right)_i - (\mathbf{l} \otimes \mathbf{f})_i\right), \tag{1.23}$$

where the constant $\mathbf{b}_i!$ term is omitted. Taking partial derivatives w.r.t. each pixel on the log energy and setting them to zero yields

$$\mathbf{f}' \otimes \left(\frac{\mathbf{b}}{\mathbf{l} \otimes \mathbf{f}} - 1\right) = 0. \tag{1.24}$$

Since $\mathbf{f}$ is a PSF, its elements sum to 1, making $\mathbf{f}' \otimes 1 = 1$. Thus Eq. (1.24) can be approximated by Richardson–Lucy deconvolution, in iterations, as

$$\mathbf{l}^{t+1} = \mathbf{l}^t \left(\mathbf{f}' \otimes \left(\frac{\mathbf{b}}{\mathbf{l}^t \otimes \mathbf{f}}\right)\right). \tag{1.25}$$

The above derivation shows that the RL method is equivalent to the Poisson maximum likelihood estimator in theory. Because there is no prior on the latent image $\mathbf{l}$, the algorithm should be stopped halfway to reduce noise and other visual artifacts. There has been research to improve RL. For example, Yuan, Sun, Quan & Shum (2008), in the multi-scale refinement scheme, applied edge-preserving bilateral filtering to the RL result. This nonlocal regularizer makes the iterative method a bit more robust against noise.

1.1.3 Recent advancements

Effective non-blind deconvolution needs to deal with noise and suppress ringing artifacts introduced by incorrect blur kernel estimates and sometimes by compression or tone management in image formation. Understanding these issues has led to better means of *regularizing* the deconvolution process in recent years, giving prior E_{prior} (denoted as $\Psi(\mathbf{l})$) a number of new forms. A general principle is that the prior should not penalize excessively estimation outliers – in order not to wrongly deviate from

© in this web service Cambridge University Press www.cambridge.org

CAMBRIDGE

Cambridge University Press
978-1-107-04436-4 - Motion Deblurring: Algorithms and Systems
Edited by A. N. Rajagopalan and Rama Chellappa
Excerpt
More information

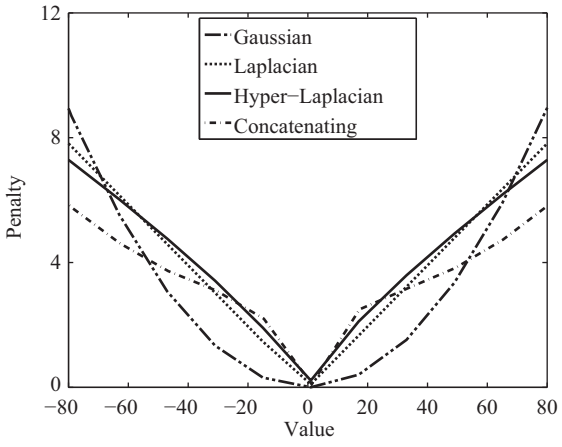

Figure 1.4 Illustration of different regularization terms. Different prior functions penalize values differently. The Gaussian prior increases energy most quickly for large absolute values.

final results. In what follows, without special mention, the overall objective function is still the one expressed in Eq. (1.6):

$$\min_{\mathbf{l}} \|\mathbf{l} \otimes \mathbf{f} - \mathbf{b}\|^2 + \lambda \Psi(\mathbf{l}). \tag{1.26}$$

Chan & Wong (1998) used a total variation regularizer, which is also known as a Laplacian prior, by setting

$$\Psi(\mathbf{l}) = \|\nabla \mathbf{l}\|^1, \tag{1.27}$$

where ∇ denotes the first-order derivative operator, i.e. $\nabla \mathbf{l} = (\partial_x \mathbf{l}, \partial_y \mathbf{l})$, a concatenation of the two gradient images. $\|\cdot\|^1$ is the L_1-norm operator for all image gradients. This prior, illustrated in Figure 1.4 by solid lines, has a stronger effect on reducing the influence of large errors than the Gaussian prior used in Eq. (1.8) (and shown as a dash–dot curve in Figure 1.4).

There are other ways to define $\Psi(\mathbf{l})$. Shan, Jia & Agarwala (2008) constructed a natural prior for the latent image by concatenating two piecewise continuous convex functions, plotted as the concatenating curve in Figure 1.4. The expression is

$$\Psi(\mathbf{l}_i) = \begin{cases} a|\nabla \mathbf{l}_i| & |\nabla \mathbf{l}_i| \leqslant \xi \\ b(\nabla \mathbf{l}_i)^2 + c & |\nabla \mathbf{l}_i| > \xi \end{cases} \tag{1.28}$$

where i indexes pixels and $\nabla \mathbf{l}_i$ represents a partial derivative for $\mathbf{l}_i$ in either the x or y direction. ξ is the value on which the linear and quadratic functions are concatenated. a, b, and c are three parameters. $\Psi(\mathbf{l})$ can actually be used to approximate natural image statistics when a is large and b is very small.

© in this web service Cambridge University Press www.cambridge.org

CAMBRIDGE

Cambridge University Press
978-1-107-04436-4 - Motion Deblurring: Algorithms and Systems
Edited by A. N. Rajagopalan and Rama Chellappa
Excerpt
More information

Table 1.1 Comparison of a few non-blind deconvolution methods with respect to the employed likelihood and prior.

	Likelihood	Prior
Wiener (Wiener 1949)	Gaussian	Gaussian
L-S[a] (Tikhonov *et al*. 1977)	Gaussian	Gaussian
RL (Richardson 1972, Lucy 1974)	Poisson	–
Chan & Wong 1998	Gaussian	Laplacian
Wang *et al*. 2008	Gaussian	Laplacian
Shan *et al*. 2008	Gaussian	Concatenating
Yuan *et al*. 2008	Poisson	Nonlocal (bilateral)
Krishnan & Fergus 2009	Gaussian	Hyper-Laplacian
Yang *et al*. 2009	Laplacian	Laplacian
Xu & Jia 2010	Laplacian	Laplacian

[a]L-S is an abbreviation for least square.

To make the resulting structure less smooth, Levin, Fergus, Durand & Freeman (2007) suggested a hyper-Laplacian prior, written as

$$\Psi(\mathbf{l}) = \|\nabla \mathbf{l}\|^{\alpha}, \tag{1.29}$$

where $\alpha < 1$, representing a norm corresponding to a sparser distribution.

Additionally, the methods of Yang, Zhang & Yin (2009) and Xu & Jia (2010) suppress noise via a TV-L_1 (total variation) objective, which uses the Laplacian data term, i.e. $E_{\text{data}} = \|\mathbf{l} \otimes \mathbf{f} - \mathbf{b}\|^1$, so the objective function can then be expressed as

$$\min_{\mathbf{l}} \|\mathbf{l} \otimes \mathbf{f} - \mathbf{b}\|^1 + \|\nabla \mathbf{l}\|^1. \tag{1.30}$$

This function can suppress strong Gaussian and, particularly, impulse image noise with robust constraint on the data term.

The likelihood and prior forms of different recent methods are listed in Table 1.1. Gaussian likelihood and Laplacian prior are most frequently used thanks to their simplicity in expression and reasonable ability to resist noise and error.

Albeit not quadratic, objective functions incorporating the Laplacian prior in Eq. (1.27), the concatenating term in Eq. (1.28), the hyper-Laplacian prior in Eq. (1.29), and the robust data term in Eq. (1.30), as a TV-L_1 energy can be solved efficiently through half-quadratic splitting, which decomposes the original problem into a single-variable quadratic minimization process. Details are provided in Section 1.1.4.

1.1.4 Variable splitting solver

An effective scheme to solve sparsely constrained non-blind deconvolution is variable splitting, implemented by half-quadratic penalty methods (Geman & Reynolds 1992, Geman & Yang 1995). This scheme has been used in many recent methods (Shan *et al*. 2008, Wang, Yang, Yin & Zhang 2008, Krishnan & Fergus 2009, Xu & Jia 2010). In what follows, we discuss the half-quadratic penalty solver for minimizing

$$E = \|\mathbf{l} \otimes \mathbf{f} - \mathbf{b}\|^2 + \lambda \|\nabla \mathbf{l}\|^{\alpha}, \tag{1.31}$$

© in this web service Cambridge University Press www.cambridge.org

CAMBRIDGE

Cambridge University Press
978-1-107-04436-4 - Motion Deblurring: Algorithms and Systems
Edited by A. N. Rajagopalan and Rama Chellappa
Excerpt
More information

with a (hyper) Laplacian prior, where $0.5 \leqslant \alpha \leqslant 1$. Objective functions with the concatenating prior expressed in Eq. (1.28) and the TV-L_1 function in Eq. (1.30) can be solved similarly.

The basic idea is to separate variables involved in convolution from those in other terms, so that they can be estimated quickly and reliably using Fourier transforms. This is realized by using a set of auxiliary variables $\psi = (\psi_x, \psi_y)$ for $\nabla \mathbf{l} = (\partial_x \mathbf{l}, \partial_y \mathbf{l})$, and adding the extra condition $\psi \approx \nabla \mathbf{l}$. Eq. (1.31) is accordingly updated to

$$E_{\mathbf{L}} = \|\mathbf{l} \otimes \mathbf{f} - \mathbf{b}\|^2 + \lambda \|\psi_x\|^\alpha + \lambda \|\psi_y\|^\alpha + \gamma \|\psi_x - \partial_x \mathbf{l}\|^2 + \gamma \|\psi_y - \partial_y \mathbf{l}\|^2, \tag{1.32}$$

given the isotropic implementation of the α-norm. γ is a weight. When its value is infinitely large, the desired conditions $\psi_x = \partial_x \mathbf{l}$ and $\psi_y = \partial_y \mathbf{l}$ can be satisfied. In this case, minimizing $E_{\mathbf{L}}$ converges to minimizing E.

Given this variable substitution, it is possible now to iterate between optimizing ψ and $\mathbf{l}$. This process is efficient and is able to converge to an optimal point, since, in each iteration, the global optimum of ψ is reached in a closed form, while a fast Fourier transform can be used to update $\mathbf{l}$.

Updating ψ

With an estimated $\mathbf{l}$ from the previous pass, Eq. (1.32) is simplified to

$$E'_\psi = \lambda \|\psi_x\|^\alpha + \lambda \|\psi_y\|^\alpha + \gamma \|\psi_x - \partial_x \mathbf{l}\|^2 + \gamma \|\psi_y - \partial_y \mathbf{l}\|^2. \tag{1.33}$$

With a few algebraic operations to decompose ψ into the set containing all elements $\psi_{i,x}$ and $\psi_{i,y}$ corresponding to all pixels i, E'_ψ can be written as a sum of sub-energy terms

$$E'_\psi = \sum_i \left(E'_{\psi_{i,x}} + E'_{\psi_{i,y}} \right), \tag{1.34}$$

where each $E'_{\psi_{i,\upsilon}}$, $\upsilon \in \{x, y\}$ only contains a single variable $\psi_{i,\upsilon} \in \psi_\upsilon$, given by

$$E'_{\psi_{i,\upsilon}} = \lambda |\psi_{i,\upsilon}|^\alpha + \gamma (\psi_{i,\upsilon} - \partial_\upsilon \mathbf{l}_i)^2, \tag{1.35}$$

where $\mathbf{l}_i$ is pixel i in $\mathbf{l}$. Each $E'_{\psi_{i,\upsilon}}$ contains only one variable $\psi_{i,\upsilon}$ so it can be optimized independently. For any α smaller than 1, minimizing Eq. (1.35) depends on two variables, i.e. joint weight γ/λ and image-dependent $\partial_\upsilon \mathbf{l}_i$. By sampling values from them, a 2D lookup table can be constructed offline, from which optimal results can be obtained efficiently. Possible errors caused by the discrepancy of actual values and nearest samples are controllable (Krishnan & Fergus 2009). For the special cases where $\alpha = 1/2$, $\alpha = 2/3$, and $\alpha = 1$, analytic solutions are available. We discuss the case $\alpha = 1$, where $\psi_{i,\upsilon}$ is expressed as

$$\psi_{i,\upsilon} = \text{sign}(\partial_\upsilon \mathbf{l}_i) \max \left\{ |\partial_\upsilon \mathbf{l}_i| - \frac{\lambda}{2\gamma}, 0 \right\}, \tag{1.36}$$

which is the 1D shrinkage formula (Wang *et al*. 2008). Its computation is efficient.

© in this web service Cambridge University Press www.cambridge.org