MOTION BLUR PRIOR

Filip Šroubek, Jan Kotera

The Czech Academy of Sciences, Institute of Information Theory and Automation,
Prague, Czech Republic

ABSTRACT

Priors play an important role of regularizers in image deblurring algorithms. Image priors are frequently studied and many forms were proposed in the literature. Blur priors are considered less important and the most common forms are simple uniform distributions with domain constraints. We propose a more informative blur prior based on the notion of atomic norm which favors blurs composed of line segments and is suitable for motion blur. The prior is formulated as a linear program that can be inserted into any optimization task. Evaluation is conducted on blind deblurring of moving objects.

Index Terms— deblurring, deconvolution, motion blur, atomic norm, convolutional sparse coding

1. INTRODUCTION

Blur is common image degradation that limits resolving power of acquisition devices. It is caused by various physical phenomena such as blur of atmospheric turbulence, out-of-focus blur, lens aberrations, or blur produced by camera/object motion. Blur characteristics imply that the problem of deblurring, i.e. estimating the original sharp image, is ill-posed. The situation is even worse in blind deblurring when the blur shape is not known and has to be estimated together with the image.

Blind deblurring is an active field of research and the literature presents a wide range of methods. The most frequent are methods based on alternating maximization of a posteriori probability (MAP) [1], then we have methods marginalizing the posterior [2], and lastly we witness a surge of learning-based methods [3]. The majority of deblurring papers consider motion blur, which probably stems from the fact that it is the most common type of blur in everyday photography; see examples in Fig. 1.

Priors are essential for alleviating the ill-posed nature of deblurring. A lot of work was devoted to image priors but much less to blur priors. Image priors in MAP approaches are based on the central idea that image features, such as gradients, are more sparse for sharp images than for blurred ones [4, 5, 6]. Likewise, the success of learning-based methods is in accurate modeling of the sharp image (patch) distribution. On the other hand, blur priors are typically flat and enforce only non-negativity and constant energy as was advocated in [7, 8]. This reasoning is driven by the notion that the blur size is by several orders of magnitudes smaller than the image size and therefore inferring the blur from the posterior is driven primarily by the likelihood function. However, if the image estimation is inaccurate, which is the case in the initial stages of blind deblurring, a more informative blur prior is likely to help in avoiding local maxima and/or speeding up the convergence.

Fig. 1. Real examples of motion blur caused by camera motion (top) and by object motion (bottom). Blurred images (left) and estimated sharp images (right) with zoomed blur kernels shown in insets.

In this work, we focus on designing priors for motion blurs that look as curves. A straightforward method for enforcing curves is to use parametric models, yet such models are either too restrictive or not known at all. Inference over parameters directly in the deblurring problem is often intractable and curves must be fitted ex-post [9, 10].

We propose a non-parametric prior favoring blurs composed of curve segments which is particularly useful for motion blur. The proposed prior is differentiable, convex, and can be easily plugged into any optimization problem. We demonstrate its performance on a blind deblurring problem of fast moving objects, where the goal is to estimate the object motion blur from a single video frame of the object traveling over the static background.

This work was supported by Czech Science Foundation grant GA18-05360S and by the Praemium Academiae awarded by the Czech Academy of Sciences.

Fig. 2. Proposed atoms: 36 line segments of length 5 used to build the atomic set A.

2. PROBLEM FORMULATION

A general formation model for a blurred image g is formulated as

$$g = \mathcal{D}(h, u) + n\,, \tag{1}$$

where $\mathcal{D}$ is a blur operator whose mathematical form is assumed to be known, h are unknown parameters of the blur, u is the unknown sharp image, and n is additive noise. For camera motion in Fig. 1(top), the blur is standard convolution $\mathcal{D}(h, u) = h * u$ and h is called convolution kernel or Point Spread Function (PSF). When the camera motion is more complex, convolution becomes space-variant and h is also a function of the position in the image. For object motion in Fig. 1(bottom), the blur operator is convolution with occlusion $\mathcal{D}(h, [u, m]) = h * u + (1 - h * m)b$, where in this case u is a foreground object in motion, m is the object silhouette in the form of a binary mask, b is the static background and h is a convolution kernel corresponding to the 2D object trajectory. If the object undergoes rotation and/or the trajectory has a strong depth component, h is more complex and the degradation cannot be expressed as convolution anymore [11]. However, irrespective of the chosen problem, we will refer to h simply as blur.

The classical solution to the ill-posed problem of estimating the sharp image (object) u and blur h from the single observed image g is formulated as a maximum a posteriori (MAP) estimator, which is equivalent to the minimization problem

$$\min_{h,u} \tfrac{\gamma}{2}\|\mathcal{D}(h, u) - g\|^2 + \phi(u) + \psi(h)\,. \tag{2}$$

The first term is given by the formation model (1) and it is called a data term (negative log of the likelihood function). In this particular case we assume n to be white Gaussian noise of distribution $N(0, 1/\gamma)$ and the data term is equivalent to least squares. The remaining two terms are regularizations (negative log of prior distributions) for the image and blur that define our prior knowledge and make the problem better posed. They must be carefully chosen as the final solution is strongly influenced by their form.

Image priors are based on the notion of image feature sparsity. It has been shown that simple features such as derivatives have heavy-tailed distribution (Laplacian or Gaussian mixture) in natural images independent of the image content. Gradient-based forms $\phi(u) := \sum_i |\nabla u_i|^p$ for $0 < p \le 1$ or $\sum_i \lambda_i |\nabla u_i|^2$ with learnable parameters λ are among the most frequently used image priors in the literature [12]. In the case of object motion with occlusion, additional prior terms enforcing relation between u and m were proposed in [10].

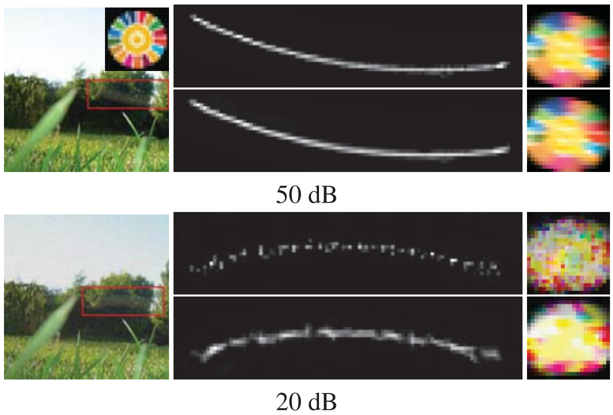

Fig. 3. Examples on synthetic data for 50 dB and 20 dB: (left) input images with a blurred moving object marked by red frames, the moving object appearance is enlarged in the inset; (center) estimated blurs using the standard blur prior on top and the proposed curve prior on bottom; (right) corresponding estimates of the moving object.

Blur priors are considered to be less informative compared to the likelihood function and the image prior. Recently, a low-rank prior was proposed in [13], yet it is nonconvex and difficult to optimize with. The most popular form for motion blur is the classical ℓ_1 norm, i.e. $\psi(h) = \sum_i |h_i|$. Frequently the non-negativity and constant energy constraints are assumed and then we obtain a flat prior with h lying in the simplex $S := \{h \,|\, h_i \ge 0, \sum_i h_i = 1\}$:

$$\psi(h) := \chi_S(h) = \begin{cases} 0, & h \in S; \\ +\infty, & h \notin S. \end{cases} \tag{3}$$

3. BLUR PRIOR FOR CURVES

This work focuses on priors that enforce curve-like blurs that are typical for motion. We propose to use an atomic norm [14], which is a versatile mathematical form for defining a broad family of convex functionals. Let $A := \{a^{(1)}, \ldots, a^{(N)}\}$ be a set of atoms $a^{(n)}$ that constitute simple building blocks of general images and they are the extreme points of the convex hull $\mathrm{conv}(A)$. The atomic norm induced by A is defined as

$$\|x\|_A := \inf\{t > 0 \,|\, \tfrac{1}{t}x \in \mathrm{conv}(A)\}\,. \tag{4}$$

If the atomic set A is centrally symmetric about the origin, i.e. $a \in A$ if and only if $-a \in A$, then $\|\cdot\|_A$ is a norm. For example, A is a set of unit-norm one-sparse elements $\{\pm e_i\}$ (the standard basis with both signs) then $\|\cdot\|_A$ corresponds to the ℓ_1 norm. However in our case, we do not require properties

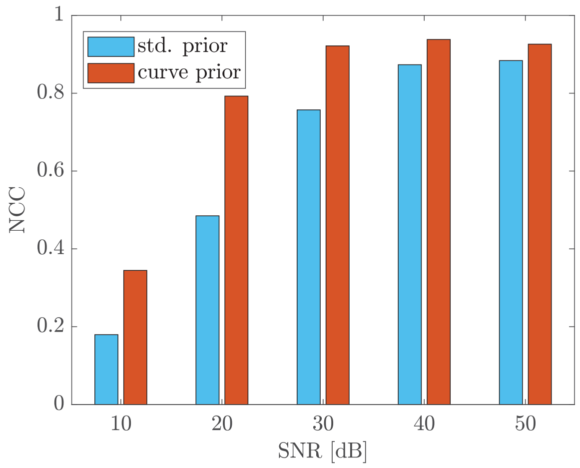

Fig. 4. Performance with respect to noise using the standard blur prior (blue) and the proposed curve prior (red). Higher values of NCC are better.

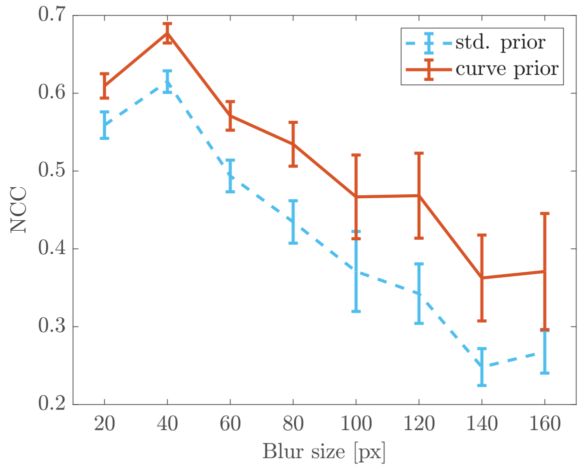

Fig. 5. Mean performance on the FMO dataset with respect to the trajectory length using the standard blur prior (blue dashed) and the proposed curve prior (red solid). Vertical bars show the standard error of the mean.

of the norm. More important is that (4) is a convex function for any set A and it can be written as a linear program

$$\begin{aligned} \|x\|_A = \min_c \sum_i c_i \\ \text{s.t.} \quad x = \sum_i c_i a^{(i)},\ c_i \geq 0\,. \end{aligned} \tag{5}$$

To construct a prior that favors curves, we design atoms $\{a^{(k)}\}$ as short line segments with different orientation and curvature; see an example of 5-pixel segments in Fig. 2. The atomic set A contains all shifted versions of such atoms to make the prior translation invariant and the linear combination $x = \sum_i c_i a^{(i)}$ can be then expressed in a more compact form with convolution as $x = \sum_k a^{(k)} * c^{(k)}$, where $c^{(k)}$ are c_i's corresponding to shifted versions of $a^{(k)}$ and arranged in an image. Note that (5) is then similar to convolutional sparse coding [15].

Using the atomic set A defined above, we propose to design the blur prior as

$$\psi(h) = \begin{cases} \alpha\|h\|_A, & h \in S; \\ +\infty, & h \notin S\,, \end{cases} \tag{6}$$

where α is the prior weight.

4. OPTIMIZATION

The proposed solution to blind restoration of motion-blurred images is obtained by substituting (6) and (5) for the blur prior in the minimization problem (2):

$$\begin{aligned} \min_{h,c,u} \tfrac{\gamma}{2}\|\mathcal{D}(h,u) - g\|^2 + \phi(u) + \chi_S(h) + \alpha \sum_i c_i \\ \text{s.t.} \quad h = \sum_i c_i a^{(i)},\ c_i \geq 0 \end{aligned} \tag{7}$$

If $\mathcal{D}$ is a linear operator such as convolution, the problem is non-smooth convex (quadratic) w.r.t. h and c. Depending on the choice of ϕ it is also convex w.r.t. u. The literature presents many techniques for solving such problems. We use the popular alternating direction method of multipliers (ADMM) [16] and here show the solution only for h. The solution for u can be derived similarly. To simplify notation, we assume the standard convolution model and write it in a matrix form $\mathcal{D}(h,u) = u * h \equiv \mathcal{U}h$, where $\mathcal{U}$ is a matrix performing convolution with u, rewrite the constraint in a matrix form $h = \sum_i c_i a^{(i)} \equiv \mathcal{A}c$, and define $\rho(c) := \alpha \sum_i P(c_i)$, where $P(x) = x$ for $x \geq 0$ and $P(x) = \infty$ for $x < 0$. To tackle (7) w.r.t h and c, we split variables h and c, and convert the constrained minimization to the unconstrained minimization of the scaled augmented Lagrangian:

$$\begin{aligned} L(h,c,v,w) = \tfrac{\gamma}{2}\|\mathcal{U}h - g\|^2 + \chi_S(v) + \rho(w) + \\ \tfrac{\beta}{2}\left(\|h - v - \lambda_v\|^2 + \|c - w - \lambda_w\|^2 + \|h - \mathcal{A}c - \lambda_h\|^2\right), \end{aligned} \tag{8}$$

where λ_v, λ_w and λ_h are scaled Lagrange multipliers of constraints $v = h$, $w = c$ and $h = \mathcal{A}c$, respectively.

We minimize L in an alternating manner and get two update equations

$$h = \left(\mathcal{U}^T\mathcal{U} + \tfrac{2\beta}{\gamma}I\right)^{-1}\left(\mathcal{U}^T g + \tfrac{\beta}{\gamma}(v + \lambda_v + \mathcal{A}c + \lambda_h)\right), \tag{9}$$

$$c = \left(\mathcal{A}^T\mathcal{A} + I\right)^{-1}\left(\mathcal{A}^T(h - \lambda_h) + w + \lambda_w\right), \tag{10}$$

solved by Conjugate Gradients, two update equations with proximal operators

$$v = \operatorname{prox}_{\chi_S}(h - \lambda_v)\,, \tag{11}$$

$$w = \operatorname{prox}_{\rho}(c - \lambda_w)\,, \tag{12}$$

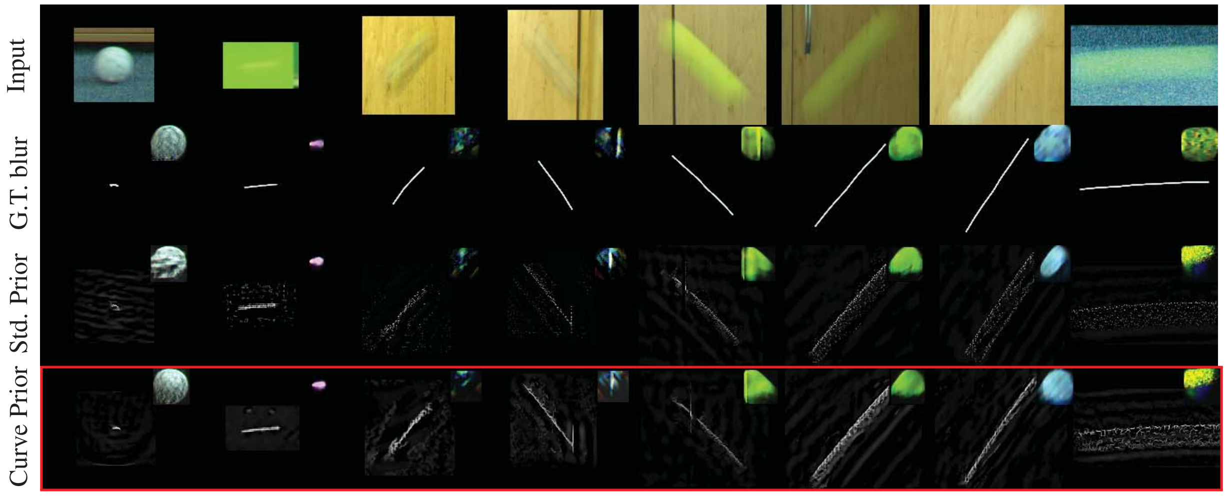

Fig. 6. FMO dataset examples of different trajectory length and estimated blurs using the standard prior (3rd row) and the proposed curve prior (4th row in the red frame). Estimated blurs are brightened to better visualize small values. In insets are sharp objects estimated using the corresponding blur.

solved by projection to the simplex [17] and by rectified soft-thresholding, respectively, and three simple update equations for Lagrange multipliers

$$\lambda_v = \lambda_v - h + v,\ \lambda_w = \lambda_w - c + w,\ \lambda_h = \lambda_h - h + \mathcal{A}c\,. \quad (13)$$

5. EXPERIMENTS

We demonstrate the performance of the proposed blur prior on the problem of fast moving object (FMO), in which the blur is a curve corresponding to the object trajectory. Refer to [11] for more details about this problem and how to solve it with ADMM. Two sets of experiments were conducted. First is on synthetic data and evaluates performance with respect to the noise level. Second is on the FMO dataset [18] and shows performance on real data with respect to the blur size.

We compare the proposed curve blur prior in (6) with the standard blur prior in (3) that only enforces non-negativity and constant energy. Normalized cross correlation (NCC) was chosen as the quality measure for evaluating the performance. If h is the ground-truth blur and $\hat{h}$ is the estimated blur, then NCC $= (\sum_i h_i \hat{h}_i)/(\|h\|_2 \|\hat{h}\|_2)$. For all experiments, the blur prior weight α was set to 1. For the synthetic data, the data term weight was set according to the noise level (SNR) to maximize the mean NCC: $\gamma = 10$ (50 dB), $\gamma = 5$ (40 dB), $\gamma = 2.5$ (30 dB), $\gamma = 1$ (20 dB) and $\gamma = 0.1$ (10 dB). For the FMO dataset, γ was set to 1.

The synthetic data simulate a simplified 2D world: a color disk of the diameter 20 flying along a random trajectory of the length 100 over the static background. Two examples for noise levels 50 dB and 20 dB, and the estimated blurs and flying object for the standard prior and the proposed curve prior are in Fig. 3. The curve prior helps the restoration method to converge to a better solution especially for lower SNRs, when the likelihood term is less informative. Notice that the curve blur prior has a positive impact also on the estimation of the moving object appearance. The mean performance over 20 random trajectories for different noise levels is summarized in Fig. 4.

The FMO dataset comprises 30-fps videos of various fast moving objects and includes ground-truth trajectories estimated from high-speed (240 fps) footage. We took 470 frames from the dataset and a video frame preceding each chosen frame was used as the background. Results were grouped into 8 categories according to the trajectory length from small blurs < 20 pixels up to blurs exceeding 160 pixels. Mean performance over the dataset is summarized in Fig. 5. Fig. 6 shows one example per category with the ground-truth trajectories and estimated blurs using the standard prior and the proposed curve prior. The curve prior outperforms the standard prior for all blur sizes with a gain slightly increasing towards larger blurs. For both priors, the performance decreases with an increasing blur size.

6. CONCLUSIONS

We proposed a blur prior tailored for motion blurs and solvable by simple linear programming. The prior is constructed using the atomic norm which is a highly flexible mathematical mapping and we believe it can be used as a general methodology for designing priors for other types of blurs in future. Performance of the proposed prior was tested on the blind restoration problem of fast moving objects and it outperformed standard uniform priors that are currently used as motion blur priors.

7. REFERENCES

[1] J. Pan, Z. Lin, Z. Su, and M. H. Yang, "Robust kernel estimation with outliers handling for image deblurring," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, June 2016, pp. 2800–2808.

[2] R. Molina, J. Mateos, and A. K. Katsaggelos, "Blind deconvolution using a variational approach to parameter, image, and blur estimation," *IEEE Transactions on Image Processing*, vol. 15, no. 12, pp. 3715–3727, Dec. 2006.

[3] O. Kupyn, V. Budzan, M. Mykhailych, D. Mishkin, and J. Matas, "DeblurGAN: Blind motion deblurring using conditional adversarial networks," in *Proc. IEEE/CVF Conf. Computer Vision and Pattern Recognition*, June 2018, pp. 8183–8192.

[4] Rob Fergus, Barun Singh, Aaron Hertzmann, Sam T. Roweis, and William T. Freeman, "Removing camera shake from a single photograph," in *SIGGRAPH '06: ACM SIGGRAPH 2006 Papers*, New York, NY, USA, 2006, pp. 787–794, ACM.

[5] A. Levin, Y. Weiss, F. Durand, and W. T. Freeman, "Understanding blind deconvolution algorithms," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 33, no. 12, pp. 2354–2367, 2011.

[6] David Wipf and Haichao Zhang, "Revisiting bayesian blind deconvolution," *Journal of Machine Learning Research (JMLR)*, November 2014.

[7] S Derin Babacan, Rafael Molina, Minh N Do, and Aggelos K Katsaggelos, "Bayesian blind deconvolution with general sparse image priors," in *Computer Vision–ECCV 2012*, pp. 341–355. Springer, 2012.

[8] D. Perrone and P. Favaro, "A clearer picture of total variation blind deconvolution," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 38, no. 6, pp. 1041–1055, June 2016.

[9] L. Fang, H. Liu, F. Wu, X. Sun, and H. Li, "Separable kernel for image deblurring," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, June 2014, pp. 2885–2892.

[10] J. Kotera, D. Rozumnyi, F. Šroubek, and J. Matas, "Intra-frame object tracking by deblatting," in *Proceedings of the IEEE International Conference on Computer Vision 2019 (ICCV 2019)*. 2019, IEEE.

[11] J. Kotera and F. Šroubek, "Motion estimation and deblurring of fast moving objects," in *Proceedings of the 25th IEEE International Conference on Image Processing 2018*. 2018, pp. 2860–2864, IEEE.

[12] Pablo Ruiz, Xu Zhou, Javier Mateos, Rafael Molina, and Aggelos K. Katsaggelos, "Variational Bayesian blind image deconvolution: A review," *Digital Signal Processing*, vol. 47, pp. 116–127, dec 2015.

[13] L. S. Yao, D. Ren, and Q. Yin, "Understanding kernel size in blind deconvolution," in *Proc. IEEE Winter Conf. Applications of Computer Vision (WACV)*, Jan. 2019, pp. 2068–2076.

[14] Venkat Chandrasekaran, Benjamin Recht, Pablo A Parrilo, and Alan S Willsky, "The convex geometry of linear inverse problems," *Foundations of Computational mathematics*, vol. 12, no. 6, pp. 805–849, 2012.

[15] T. Blumensath and M. Davies, "Sparse and shift-invariant representations of music," *and Language Processing IEEE Transactions on Audio, Speech*, vol. 14, no. 1, pp. 50–57, Jan. 2006.

[16] Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends® in Machine Learning*, vol. 3, no. 1, pp. 1–122, 2011.

[17] W. Wang and M. A. Carreira-Perpinan, "Projection onto the probability simplex: An efficient algorithm with a simple proof, and an application," *arXiv preprint arXiv:1309.1541*, 2013.

[18] D. Rozumnyi, J. Kotera, F. Šroubek, L. Novotný, and J. Matas, "The world of fast moving objects," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2017, pp. 5203–5211, IEEE.