

Blind Image Deblurring Using Dark Channel Prior

Jinshan Pan^{1,2,3} Deqing Sun^{3,4} Hanspeter Pfister³ Ming-Hsuan Yang²
¹Dalian University of Technology ²UC Merced ³Harvard University ⁴NVIDIA




Figure 1. Deblurring result on a challenging low-light image. The blur process makes the dark channel of the blurred image less sparse (c). Enforcing sparsity on the dark channel of the recovered image favors clean images over blurred ones.

Abstract

We present a simple and effective blind image deblurring method based on the dark channel prior. Our work is inspired by the interesting observation that the dark channel of blurred images is less sparse. While most image patches in the clean image contain some dark pixels, these pixels are not dark when averaged with neighboring high-intensity pixels during the blur process. This change in the sparsity of the dark channel is an inherent property of the blur process, which we both prove mathematically and validate using training data. Therefore, enforcing the sparsity of the dark channel helps blind deblurring on various scenarios, including natural, face, text, and low-illumination images. However, sparsity of the dark channel introduces a non-convex non-linear optimization problem. We introduce a linear approximation of the min operator to compute the dark channel. Our look-up-table-based method converges fast in practice and can be directly extended to non-uniform deblurring. Extensive experiments show that our method achieves state-of-the-art results on deblurring natural images and compares favorably methods that are well-engineered for specific scenarios.

1. Introduction

Blind image deblurring aims to recover a blur kernel and a sharp latent image from a blurred image. This is a classical image and signal processing problem [22], which has been an active research effort in the vision and graphics community within the last decade. This problem becomes increasingly important as more photos are taken using hand-held

cameras, particularly with smart phones. Camera shake is often inevitable and the resulting image blur is usually undesirable. As captured moments are ephemeral and difficult to reproduce, it is of great interest to remove blur for a higher-quality image.

When the blur is uniform and spatially invariant, we can model the blur process with the convolution operation

$$B = I \otimes k + n, \quad (1)$$

where B , I , k , and n denote the blur image, latent image, blur kernel, and noise, respectively, and \otimes is the convolution operator. As only B is available, we need to recover both I and k simultaneously. This problem is highly ill-posed because many different pairs of I and k give rise to the same B , e.g., blurred images and delta blur kernels.

To make blind deblurring well posed, existing methods make assumptions on blur kernels, latent images, or both [2, 7, 17, 20, 26, 27, 36]. For example, numerous methods [2, 7, 19, 20] assume sparsity of image gradients, which has been widely used in low-level vision tasks including denoising, stereo, and optical flow. Levin *et al.* [19] show that deblurring methods based on this prior tend to favor blurry images over original clear images, especially for algorithms formulated within the maximum a posteriori (MAP) framework. To remedy this problem, a heuristic edge selection step [5, 34] is often necessary to achieve state-of-the-art results in the MAP framework. New natural image priors have also been introduced that favor clean images over blurred ones, e.g., normalized sparsity prior [17], L_0 -regularized prior [36], and internal patch recurrence [24]. However,

these natural image models do not generalize well to specific images, such as face [25], text [3, 4, 26], and low-illumination [12] images.

We present a deblurring algorithm that achieves competitive results on both natural and specific images. Our work is motivated by an interesting observation on the blur process: dark channels (smallest values in a local neighborhood) of blurred images are less dark. Intuitively, when a dark pixel is averaged with neighboring high-intensity pixels during the blur process, its intensity increases. We show theoretically and empirically that this generic property of the blur process holds for many images. This inspires us to propose an L_0 -regularization term to minimize the dark channel of the recovered image. This new term favors clean images over blurred images in the restoration process.

Optimizing the new L_0 -regularized dark channel term is challenging. The L_0 norm is highly non-convex and the optimization involves a non-linear minimum operation. We propose an approximate linear operator based on look-up tables for the \min operator, and solve the linearized L_0 minimization problem by half-quadratic splitting methods. The proposed algorithm converges quickly in practice and can be naturally extended to non-uniform deblurring tasks.

The contributions of this work are as follows: (1) we theoretically prove that the blur (convolution) operation increases the values of the dark channel pixels; (2) we empirically confirm our analysis using a dataset of 3,200 clean and blurred image pairs; (3) we introduce an L_0 -regularization term to enforce sparsity on the dark channel of latent images and develop an efficient optimization scheme; (4) our algorithm achieves state-of-the-art performance on widely-used natural image deblurring benchmarks [16, 19, 29], and competitive results on specific deblurring tasks, including text, face, and low-illumination images, which are not well handled by most recent deblurring methods for natural images. Further, our method also works on non-uniform deblurring.

2. Related Work

In recent years, we have witnessed significant advances in single image deblurring [14, 16] mainly due to the use of statistical priors on natural images and selection of salient edges for kernel estimation [5, 7, 17, 20, 27, 34, 36].

Fergus *et al.* [7] use a mixture of Gaussians to learn an image gradient prior via variational Bayesian inference. Levin *et al.* [19] show that the variational Bayesian inference method [7] is able to avoid trivial solutions while naive MAP based methods may not. However, the variational Bayesian approach is computationally expensive, and efficient methods require approximation [20].

Efficient methods based on MAP formulations have been developed with different likelihood functions and image priors [1, 17, 21, 27, 31, 36, 37]. In particular, heuristic edge selection methods for kernel estimation [5, 15, 34] have been

proposed and demonstrated effective for the MAP estimation framework [16]. However, the assumption that strong edges exist in the latent images may not always hold.

To better reconstruct sharp edges for kernel estimation, recent exemplar-based methods [9, 25, 29] exploit information contained in both a blurred input and example images from an external dataset. However, querying a large external dataset is computationally expensive.

Numerous recent methods exploit domain-specific statistical properties for deblurring, such as text [3, 4, 26], face [25], and low-illumination images [12]. While these domain-specific methods generate better results than generic deblurring algorithms, each application requires specific operations or significant engineering effort. In this work, we propose a generic algorithm based on how the blur process affects the dark channel.

The dark channel prior was introduced by He *et al.* for single image dehazing [10] based on the assumption that the dark channel in the haze-free outdoor image is zero. In this work, we make the less restrictive assumption that the dark channel of the original image is sparse instead of zero, and we show that the proposed method is able to deblur a large variety of images. To enforce the sparsity of the dark channel, we develop a novel optimization scheme for the resulting non-linear non-convex problem.

3. Convolution and Dark Channel

To motivate our work, we first describe the dark channel and then its role in image deblurring. For an image I , the dark channel [10] is defined by

$$D(I)(x) = \min_{y \in \mathcal{N}(x)} \left(\min_{c \in \{r, g, b\}} I^c(y) \right), \quad (2)$$

where x and y denote pixel locations; $\mathcal{N}(x)$ is an image patch centered at x ; and I^c is the c -th color channel. If I is a gray-scale image, we have $\min_{c \in \{r, g, b\}} I^c(y) = I(y)$. The dark channel prior is mainly used to describe the minimum values in an image patch. He *et al.* [10] observe that the dark channel of outdoor, haze-free images is almost zero. We find that most, although not all, elements of the dark channel are zero for natural images (see Figure 2(a) and (c)). However, most elements in the dark channel of blurred images are nonzero, as shown in Figure 2(b) and (d).

To explain why the dark channel of blurred images are less sparse, we derive some properties of the blur (convolution) operation. For discrete signals (images), convolution is defined as the sum of the product of the two signals after one is reversed and shifted

$$B(x) = \sum_{z \in \Omega_k} I(x + \lfloor \frac{s}{2} \rfloor - z) k(z), \quad (3)$$

where Ω_k and s denote the domain and size of blur kernel k , $k(z) \geq 0$, $\sum_{z \in \Omega_k} k(z) = 1$, and $\lfloor \cdot \rfloor$ denotes the rounding




Figure 2. Blurred images have less sparse dark channels than clear images. The blur process (convolution) outputs a weighted average of pixels in a neighborhood and tends to increase the value of the minimum pixel. Top: images; bottom: corresponding dark channels computed with an image patch size of 35×35 .

operator. We note that (3) can be regarded as the sum of a locally weighted linear combination of I .

Why do blurred images have fewer dark pixels? Intuitively, the weighted sum of pixel values in a local neighborhood is larger than the minimum pixel value in the neighborhood, *i.e.*, convolution increases the values of the dark pixels. Mathematically, we have the following proposition.

Proposition 1: Let $\mathcal{N}(x)$ denote a patch centered at pixel x with size the same as the blur kernel. We have:

$$B(x) \geq \min_{y \in \mathcal{N}(x)} I(y). \quad (4)$$

Proof. Based on the definition of convolution (3), we have

$$\begin{aligned} B(x) &= \sum_{z \in \Omega_k} I(x + \left\lfloor \frac{s}{2} \right\rfloor - z) k(z) \geq \sum_{z \in \Omega_k} \min_{y \in \mathcal{N}(x)} I(y) k(z) \\ &= \min_{y \in \mathcal{N}(x)} I(y) \sum_{z \in \Omega_k} k(z) = \min_{y \in \mathcal{N}(x)} I(y). \quad \square \end{aligned}$$

Note that when x is the dark pixel in its neighborhood, *i.e.*, $I(x) = \min_{y \in \mathcal{N}(x)} I(y)$, $B(x) \geq I(x)$. This means that the intensity values of dark pixels in I tend to become larger after the convolution, as shown in Figure 2.

Proposition 1 enables us to derive two properties to describe the changes caused to blurred images by convolution:

Property 1: Let $D(B)$ and $D(I)$ denote the dark channel of the blurred and clear images, we have:

$$D(B)(x) \geq D(I)(x). \quad (5)$$

Please see the supplementary material for the detailed proof.

Property 2: Let Ω denote the domain of an image I . If there exist some pixels $x \in \Omega$ such that $I(x) = 0$, we have:

$$\|D(B)(x)\|_0 > \|D(I)(x)\|_0, \quad (6)$$




Figure 3. Intensity histograms for dark channels of both clear and blurred images in a dataset of 3,200 natural images. Blurred images have far fewer zero dark channel pixels than clear ones, confirming our analysis in the text. The dark channel of each image has been computed with an image patch size of 35×35 .

where the L_0 norm $\|\cdot\|_0$ counts the nonzero elements of $D(I)$. Property 2 directly follows from Property 1.

We further validate our analysis using a dataset of 3,200 natural images.¹ As shown in Figure 3, the dark channels of clear images have significantly more zero elements than those of blurred images. This property also holds for other image types, such as text and saturated images (please see Section 7 and the supplemental material for the statistics). Thus, the sparsity of dark channels is a natural metric to distinguish clear images from blurred images. This observation motivates us to introduce a new regularization term to enforce sparsity of dark channels in latent images.

4. Model and Optimization

From our analysis and observations, we use the $\|D(I)\|_0$ norm to measure sparsity of dark channels. We add this constraint to a standard formulation for image deblurring as

$$\min_{I, k} \|I \otimes k - B\|_2^2 + \gamma \|k\|_2^2 + \mu \|\nabla I\|_0 + \lambda \|D(I)\|_0, \quad (7)$$

where the first term imposes that the convolution output of the recovered image and the blur kernel should be similar to the observation; the second term is used to regularize the solution of the blur kernel; the third term on image gradients retains large gradients and removes tiny details [26, 36]; γ , μ , and λ are weight parameters. We use coordinate descent to alternatively solve for the latent image I :

$$\min_I \|I \otimes k - B\|_2^2 + \mu \|\nabla I\|_0 + \lambda \|D(I)\|_0, \quad (8)$$

and the blur kernel k :

$$\min_k \|I \otimes k - B\|_2^2 + \gamma \|k\|_2^2. \quad (9)$$

¹The images are from both BSDS [23] and the Internet. The datasets are available on the authors' websites.

4.1. Estimating the Latent Image I

Minimizing (8) is computationally intractable because of the L_0 -regularized term and the non-linear function $D(\cdot)$. To tackle the L_0 -regularized term, we use the half-quadratic splitting L_0 minimization approach [35]. Similar to [26], we introduce the auxiliary variables u with respect to $D(I)$ and $g = (g_h, g_v)$ corresponding to image gradients in the horizontal and vertical directions. The objective function (8) can be rewritten as:

$$\begin{aligned} \min_{I, u, g} & \|I \otimes k - B\|_2^2 + \alpha \|\nabla I - g\|_2^2 \\ & + \beta \|D(I) - u\|_2^2 + \mu \|g\|_0 + \lambda \|u\|_0, \end{aligned} \quad (10)$$

where α and β are penalty parameters. When α and β are close to infinity, the solution of (10) approaches that of (8) [32]. We can solve (10) by alternatively minimizing I , u , and g while fixing the other variables. Note that given I , the subproblems of solving for the auxiliary variables u and g do not involve the nonlinear function $D(\cdot)$.

Now we will explain how to deal with the nonlinear min operator when solving for I :

$$\min_I \|I \otimes k - B\|_2^2 + \alpha \|\nabla I - g\|_2^2 + \beta \|D(I) - u\|_2^2. \quad (11)$$

Our observation is that the non-linear operation $D(I)$ is equivalent to a linear operator \mathbf{M} applied to the vectorized image \mathbf{I} .² Let $y = \operatorname{argmin}_{z \in \mathcal{N}(x)} I(z)$. \mathbf{M} satisfies:

$$\mathbf{M}(x, z) = \begin{cases} 1, & z = y, \\ 0, & \text{otherwise.} \end{cases} \quad (12)$$

Multiplying the x -th row of \mathbf{M} with \mathbf{I} gives the value of the pixel y , *i.e.*, $\mathbf{I}(y)$ or equivalently $\mathbf{D}(\mathbf{I})(x)$ (see the top row in Figure 4). Given the previous estimated intermediate latent image, we can construct the desired matrix \mathbf{M} according to (12), as shown in Figure 4.

For the true clear image, $\mathbf{M}\mathbf{I} = \mathbf{D}(\mathbf{I})$ strictly holds. Without the clear image, we compute an approximation of \mathbf{M} using the intermediate result at each iteration. As the intermediate result becomes closer to the clear image, \mathbf{M} approaches to the desired \mathbf{D} . Empirically, we find that the approximation scheme converges well, as shown in Figure 15.

Given the selection matrix \mathbf{M} , we solve for I by:

$$\min_I \|\mathbf{T}_k \mathbf{I} - \mathbf{B}\|_2^2 + \alpha \|\nabla \mathbf{I} - \mathbf{g}\|_2^2 + \beta \|\mathbf{M}\mathbf{I} - \mathbf{u}\|_2^2, \quad (13)$$

where \mathbf{T}_k is a Toeplitz (convolution) matrix of k , \mathbf{B} , \mathbf{g} , and \mathbf{u} denote vector forms of B , g , and u , respectively. The matrix-vector production with respect to the Toeplitz matrix can be achieved using the Fast Fourier Transform (FFT) [32]. The solution of (13) can be obtained according to [26, 27, 36].

²For consistency, we use $\mathbf{D}(\mathbf{I})$ to denote the vector form of $D(I)$.

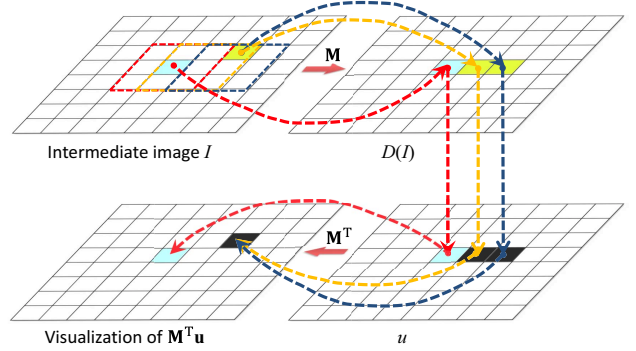


Figure 4. Top: computing the dark channel $D(I)$ of an image I by the non-linear min operator is equivalent to multiplying a linear selection matrix \mathbf{M} with the vectorized image \mathbf{I} . The three squares in the intermediate image denote adjacent image patches for computing the dark channel, where the minimum intensity value in each patch is marked with different colors. Bottom: the transpose \mathbf{M}^T enforces identified dark pixels to be consistent with u .

Given I , we compute u and g separately by:

$$\begin{aligned} \min_u & \beta \|D(I) - u\|_2^2 + \lambda \|u\|_0, \\ \min_g & \alpha \|\nabla I - g\|_2^2 + \mu \|g\|_0. \end{aligned} \quad (14)$$

We note that (14) is an element-wise minimization problem. Thus, the solution of u is:

$$u = \begin{cases} D(I), & |D(I)|^2 \geq \frac{\lambda}{\beta}, \\ 0, & \text{otherwise,} \end{cases} \quad (15)$$

and similarly for the solution of g . The algorithmic details of (10) are presented in the supplemental material.

4.2. Estimating Blur Kernel k

Given I , the kernel estimation in (9) is a least squares problem. We note that kernel estimation methods based on gradients have been shown to be more accurate [5, 20, 36] (see analysis in the supplemental material). Thus, we estimate the blur kernel k by:

$$\min_k \|\nabla I \otimes k - \nabla B\|_2^2 + \gamma \|k\|_2^2. \quad (16)$$

Similar to existing approaches [5, 26, 36], we obtain the solution of (16) by FFTs. After obtaining k , we set the negative elements of k to 0, and normalize k so that k satisfies our definition of the blur kernel. Similar to state-of-the-art methods, the proposed kernel estimation process is carried out in a coarse-to-fine manner using an image pyramid [5]. Algorithm 1 shows the main steps for the kernel estimation algorithm on one pyramid level.

5. Extension to Non-Uniform Deblurring

Our method can be directly extended to handle non-uniform deblurring where the blurred images are acquired from moving cameras (*e.g.*, rotational and translational

Algorithm 1 Blur kernel estimation algorithm

Input: Blurred image B .
 initialize k with results from the coarser level.
while $i \leq \text{max_iter}$ **do**
 solve for I using (10).
 solve for k using (16).
end while
Output: Blur kernel k and intermediate latent image I .




Figure 5. Quantitative evaluations on two benchmark datasets. Our method performs competitively against the state-of-the-art.

movements) [8, 11, 28, 30, 33]. Based on the geometric model of camera motion [30, 33], the non-uniform blur model can be expressed as:

$$\mathbf{B} = \sum_t k_t \mathbf{H}_t \mathbf{I} + \mathbf{n}, \quad (17)$$

where \mathbf{I} and \mathbf{n} denote vector forms of I , n in (1); t is the index of camera pose samples; \mathbf{H}_t is a matrix derived from the homography matrix in [33]; k_t is the weight corresponding to the t -th camera pose, which satisfies $k_t \geq 0$ and $\sum_t k_t = 1$. Similar to [33], (17) can be expressed as:

$$\mathbf{B} = \mathbf{K}\mathbf{I} + \mathbf{n} = \mathbf{A}\mathbf{k} + \mathbf{n}, \quad (18)$$

where \mathbf{k} is a vector and its element is composed of the weight k_t . Based on (18), the non-uniform deblurring process is achieved by alternatively minimizing:

$$\min_{\mathbf{I}} \|\mathbf{K}\mathbf{I} - \mathbf{B}\|_2^2 + \lambda \|\mathbf{D}(\mathbf{I})\|_0 + \mu \|\nabla \mathbf{I}\|_0 \quad (19)$$

and

$$\min_{\mathbf{k}} \|\mathbf{A}\mathbf{k} - \mathbf{B}\|_2^2 + \gamma \|\mathbf{k}\|_2^2. \quad (20)$$

We employ the fast forward approximation [11] to estimate the latent image \mathbf{I} and the weight \mathbf{k} . The algorithmic details are presented in the supplementary material. Our MATLAB code is publicly available on the authors' websites.

6. Experimental Results

We examine our method on two natural image deblurring datasets [16, 29] and compare it to state-of-the-art natural image deblurring methods. Then, we evaluate our method using text [26], face [25], and low-illumination [12] images and further compare it to methods specially designed for these tasks. Finally, we report results on images undergoing non-uniform blurs. Due to the comprehensive experiments performed, we only show a small portion of the results in the main paper. Please see the supplementary document for more and larger result images.

Parameter setting: In all experiments, we set $\lambda = \mu = 0.004$, $\gamma = 2$, and the neighborhood size to compute the dark channel in (2) to be 35 (please see the supplemental material for analysis). We empirically set $\text{max_iter} = 5$ as a trade-off between accuracy and speed. As our focus is on the kernel estimation, We follow the practice [7, 19, 34] to use a non-blind deblurring method to recover the final latent image with our estimated kernel. We use the non-blind method [26] unless otherwise mentioned. Our MATLAB code is publicly available on the authors' websites.

Natural images: We use the image dataset by Köhler *et al.* [16], which contains 4 images and 12 blur kernels. The PSNR value is computed by comparing each restored image with 199 clear images captured along the camera motion trajectory. As shown in Figure 5(a), our method has the highest average PSNR among all the methods evaluated. Figure 6 shows results on a challenging example with heavy blur. Although state-of-the-art methods [5, 34] are able to deal with large blur in most places, their deblurred images contain moderate ringing artifacts. In contrast, our result has fewer artifacts and clearer details.

Next, we evaluate our method on the dataset by Sun *et al.* [29], which contains 80 images and 8 blur kernels. For fair comparisons, we use the provided codes of state-of-the-art methods [5, 17, 20, 24, 26, 29, 34, 36] to estimate blur kernels and use the non-blind deblurring method [38] to generate the final deblurring results. We use the error ratio [19] as the quality metric. As Figure 5(b) shows, our method consistently outperforms state-of-the-art methods.

We further test our method using a real natural image (Figure 7). We use the same non-blind deconvolution method [26] with blur kernels estimated by each method. While several state-of-the-art methods [17, 26, 36] produce strong ringing artifacts and blur effects, our method generate clearer images. The deblurred image by our method without the dark channel prior contains considerable artifacts, suggesting the effectiveness of the dark channel prior.

Text images: Table 1 summarizes the PSNR results on the text image dataset [26], which contains 15 clear text images and 8 blur kernels. The average PSNR by our method is at least 1.7dB higher than those by other natural image deblurring methods [5, 17, 20, 34, 36] and less than 0.9dB lower than that by the specially-designed method [26]. Visually, the recovered image by our method compares favorably to that by [26] (Figure 8).

Low-illumination images: Blurred images captured in low-illumination scenes are particularly challenging for most deblurring methods, because they often have saturated pixels that interfere with the kernel estimation process [6, 12]. For example, the kernel estimate by [36] looks like a delta kernel due to the influence of saturated regions as shown in Figure 9(b); and the deblurred image has significant residual blur. Compared with the clean image, the




Figure 6. Visual comparisons using one challenging image from the dataset [16]. The deblurred images from other methods are from the reported results in [16]. The recovered image by the proposed algorithm with the dark channel prior is visually more pleasing.

Table 1. Quantitative evaluations on the text image dataset [26]. Our method outperforms several recent deblurring methods for natural images and is comparable to the method designed for text images [26].

	Cho and Lee [5]	Xu and Jia [34]	Krishnan <i>et al.</i> [17]	Levin <i>et al.</i> [20]	Xu <i>et al.</i> [36]	Pan <i>et al.</i> [26]	Ours
Average PSNRs	23.80	26.21	20.86	24.90	26.21	28.80	27.94

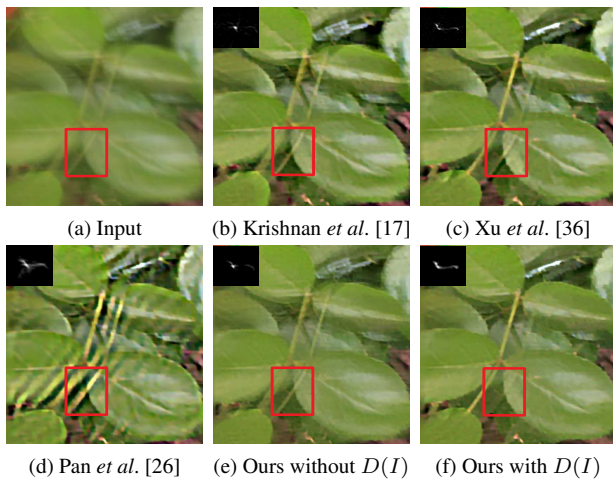


Figure 7. Comparisons on a real natural image. The parts in red boxes in (b)-(e) still contain significant residual blur. (Best viewed on high-resolution display with zoom-in.)




Figure 8. On text images, our generic method generates results comparable to methods tailored to text. (Best viewed on high-resolution display with zoom-in.)

blurred image with saturated regions also has a less sparse dark channel. As a result, directly applying our method produces results comparable to [12], which has been specifically designed for low-light conditions.

Face images: Blurred face images are also challenging for methods designed for natural images, because they contain fewer edges or textures [25] for kernel estimation. As shown in Figure 10, our method compares favorably against [25], which explicitly explores facial structures using an exemplar dataset.

Non-uniform deblurring: As our method can naturally be




Figure 9. Results on a saturated image. The deblurring results are all generated by the non-blind deconvolution method [12]. Residual blur and ringing artifacts exist in the red boxes in (b)-(c). (Best viewed on high-resolution display with zoom-in.)



Figure 10. Comparisons on blurred face images. Our method compares favorably with [25], which uses a face dataset to explore face structures for deblurring face images.

extended to deal with non-uniform blur, we also report results on an image degraded by spatially-variant motion blur in Figure 11 (please see the supplemental material for more examples and large images). Compared with the state-of-the-art non-uniform deblurring method [36], our method generates images with fewer artifacts and clearer textures.

7. Analysis and Discussions

It is surprising that the dark channel prior enables us to design a method that outperforms state-of-the-art methods on natural images but also obtains competitive results on

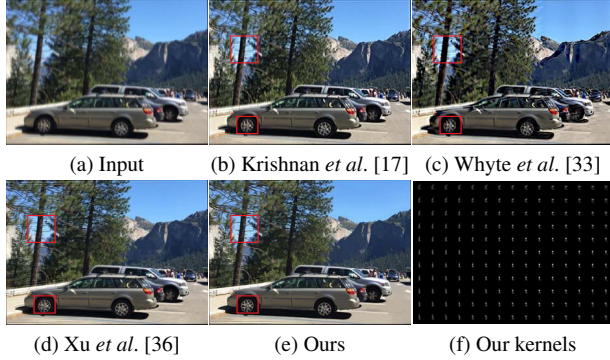


Figure 11. The dark channel prior directly applies to images with non-uniform blur. The parts in red boxes in (b)-(d) still contain ringing artifacts and residual blurs. (Best viewed on high-resolution display with zoom-in.)

specific scenarios without using domain knowledge. In this section, we further analyze the proposed method, compare it with related methods, and discuss its limitations.

Effectiveness of the dark channel prior: Our method without the dark channel prior reduces to the deblurring method of Xu *et al.* [36]. To ensure fair comparison, we disable the dark channel prior in our implementation. As shown in Figure 12(f) and (g), using the dark channel prior generates intermediate results with more sharp edges, which favors clear images and facilitates kernel estimation. Also, the dark channel of the intermediate results becomes sparser with more iterations (Figure 12(h)). We quantitatively evaluate our method with and without the dark channel prior using two benchmark datasets [16, 19]. The results in Figure 13 show that the dark channel prior consistently improves deblurring. In particular, our method with the dark channel prior has 100% success rate on the dataset by Levin *et al.* [19]. All these results concretely demonstrate the effectiveness of the dark channel prior.

Favored minimum of the energy function: The dark channel prior is effective because it has lower energy for clear images than for blurred ones. Two notable methods [17, 24] also have energy functions with similar properties. However, they are mainly designed for natural images and are less effective for specific scenarios (*e.g.*, text and low-illumination images). For example, the normalized sparsity prior [17] gives lower energy to clear natural images than blurred images, but does not always favor clear text images (Figure 14(b)). In contrast, the dark channel prior favors clear text images (Figure 14(a)). In [24], internal patch recurrence is exploited for image deblurring. The method performs well when images have repeated patterns among patches, but may fail otherwise. Our analysis and observation suggest that the dark channel prior can broadly apply to scenarios where blur makes the dark channel less sparse.

He *et al.* [10] first introduce the dark channel prior for image dehazing. They assume that all elements of the dark channel are zero, which mainly holds for outdoor haze-free

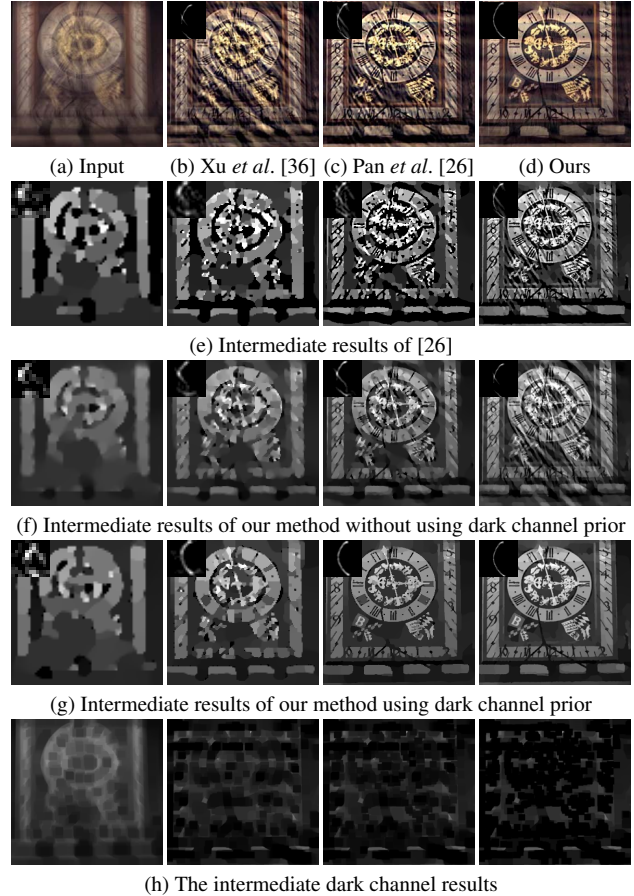


Figure 12. Deblurred images by several methods are shown in (a)-(d), and the intermediate results over iterations (from left to right) are shown in (e)-(h). With the dark channel prior, our method recovers intermediate results containing more sharp edges for kernel estimation. The dark channels of the intermediate results become darker, which favor clear images and facilitate kernel estimation.

images. In contrast, our analysis shows that, generally, the blur operation makes the dark channel of clean images less sparse. Therefore, we assume that the dark channel of clear images is sparse. Empirically, this assumption holds not only for natural images, but also for specific scenarios, including text (Figure 14(a)) and saturated images (Figure 1). Note that the dark channel prior and domain knowledge are more likely to be complementary than contradictory. Future work could study the relationship between these complementary priors.

Relation with L_0 -regularized deblurring methods: Two previous methods [26, 36] have used L_0 -regularized priors for deblurring. The method [36] assumes L_0 sparsity on image gradients, which performs well on natural images but is less effective for text images (Figure 8(b)). The method [26] assumes L_0 sparsity on both the intensity and gradients for deblurring text images. The L_0 -regularized intensity term plays a key role in text image deblurring, because the intensity values (histograms) of text images are close to two-




Figure 13. Quantitative results of our method with and without the dark channel prior on two benchmark datasets. The dark channel prior consistently improves the results. In particular, our method with the dark channel prior has 100 % success at error ratio 2 on the dataset by Levin *et al.* [19].




Figure 14. Statistics of different priors on the text image deblurring dataset [26]. The normalized sparsity prior [17] (*i.e.*, L_1/L_2) sometimes favors blurred text images.

tone. However, the intensity histograms of natural images are more complex than those of text images, and this prior is not applicable to natural image deblurring problems (Figure 7(d)). The intermediate results in Figure 12(e) also show that although this L_0 -regularized intensity term helps preserve significant contrast compared to (f), it fails to recover useful structures for kernel estimation.

Convergence property: As our energy function is non-linear and highly non-convex, a natural question is whether our optimization method converges (to a good local minimum). We quantitatively evaluate convergence properties of our method on the benchmark dataset by Levin *et al.* [19]. Figure 15(a) and (b) suggest that the proposed method converges after less than 50 iterations, in terms of the average kernel similarity values [13] and the energies computed from (7). Note that the kernel estimation methods based on image intensity (*i.e.*, (9)) and gradients (*i.e.*, (16)) have similar convergence properties. More discussions are included in the supplemental material.

Computational complexity: Compared to the L_0 -regularized methods [26, 36], our method additionally requires computing the dark channel and look-up table. The complexity of this step is $O(N)$ and independent of patch size [18], where N is the number of pixels. This is the main bottleneck. Other steps can be accelerated by FFTs. Our method takes about 17 seconds for a 255×255 image on a computer with an Intel Core i7-4790 processor and 28 G-




Figure 15. Fast convergence property of our method, which empirically validates our approximation of the non-linear operator.

B RAM (see supplemental material for the running time of other methods and more discussions).

Limitations: Despite its robust performance on a variety of challenging datasets, our method has limitations. When a clear image has no dark pixels, the dark channel prior is less likely to help kernel estimation. In this situation, Property 2 does not hold and $\|D(B)(x)\|_0 = \|D(I)(x)\|_0$. The solution of u given by (15) is likely to be $D(I)$ as the value of $\frac{\lambda}{\beta}$ will be much smaller than that of $D(I)$. Thus, the constraint $\|D(I)\|_0$ would have no effect on the intermediate latent image estimation. As a result, our method with and without the dark channel have almost the same result (see the supplemental material). In addition, our method assumes that only the blur process changes the sparseness of the dark channel. Significant noise may affect the dark pixels of an image, which accordingly interferes with the kernel estimation (see the supplemental material for examples and more discussions). Future work will consider joint deblurring and denoising using the dark channel prior.

8. Concluding Remarks

Based on an analysis of the convolution operation and its effect on the dark channel of blurred images, we have introduced a simple and effective blind image deblurring algorithm. The proposed dark channel prior captures the changes to blurred images caused by the blur process, and favors clear images over blurred ones in the deblurring process. To restore images regularized by the dark channel prior, we develop an effective optimization algorithm based on a half-quadratic splitting strategy and look-up tables. The proposed algorithm does not require heuristic edge selection steps or any complex processing techniques in kernel estimation, *e.g.*, shock filtering and bilateral filtering. Furthermore, the proposed algorithm is easily extended to handle non-uniform blur. Our algorithm achieves state-of-the-art results on deblurring natural images, and performs favorably against specialized methods for faces, texts, and low-illumination conditions.

Acknowledgements: This work has been supported in part by NSF CAREER (No. 1149783), NSF IIS (No. 1152576), NSF OIA (No. 1125087), NSFC (No. 61572099), and a gift from Adobe. J. Pan has been supported by a scholarship from the China Scholarship Council.

References

- [1] J.-F. Cai, H. Ji, C. Liu, and Z. Shen. Framelet based blind motion deblurring from a single image. *IEEE TIP*, 21(2):562–572, 2012. 2
- [2] T. Chan and C. Wong. Total variation blind deconvolution. *IEEE TIP*, 7(3):370–375, 1998. 1
- [3] X. Chen, X. He, J. Yang, and Q. Wu. An effective document image deblurring algorithm. In *CVPR*, pages 369–376, 2011. 2
- [4] H. Cho, J. Wang, and S. Lee. Text image deblurring using text-specific properties. In *ECCV*, pages 524–537, 2012. 2
- [5] S. Cho and S. Lee. Fast motion deblurring. In *SIGGRAPH Asia*, volume 28, page 145, 2009. 1, 2, 4, 5, 6
- [6] S. Cho, J. Wang, and S. Lee. Handling outliers in non-blind image deconvolution. In *ICCV*, pages 495–502, 2011. 5
- [7] R. Fergus, B. Singh, A. Hertzmann, S. T. Roweis, and W. T. Freeman. Removing camera shake from a single photograph. *ACM SIGGRAPH*, 25(3):787–794, 2006. 1, 2, 5
- [8] A. Gupta, N. Joshi, C. L. Zitnick, M. F. Cohen, and B. Curless. Single image deblurring using motion density functions. In *ECCV*, pages 171–184, 2010. 5
- [9] Y. HaCohen, E. Shechtman, and D. Lischinski. Deblurring by example using dense correspondence. In *ICCV*, pages 2384–2391, 2013. 2
- [10] K. He, J. Sun, and X. Tang. Single image haze removal using dark channel prior. In *CVPR*, pages 1956–1963, 2009. 2, 7
- [11] M. Hirsch, C. J. Schuler, S. Harmeling, and B. Schölkopf. Fast removal of non-uniform camera shake. In *ICCV*, pages 463–470, 2011. 5
- [12] Z. Hu, S. Cho, J. Wang, and M.-H. Yang. Deblurring low-light images with light streaks. In *CVPR*, pages 3382–3389, 2014. 2, 5, 6
- [13] Z. Hu and M.-H. Yang. Good regions to deblur. In *ECCV*, pages 59–72, 2012. 8
- [14] J. Jia. *Mathematical models and practical solvers for uniform motion deblurring*. Cambridge University Press, 2014. 2
- [15] N. Joshi, R. Szeliski, and D. J. Kriegman. PSF estimation using sharp edge prediction. In *CVPR*, 2008. 2
- [16] R. Köhler, M. Hirsch, B. J. Mohler, B. Schölkopf, and S. Harmeling. Recording and playback of camera shake: Benchmarking blind deconvolution with a real-world database. In *ECCV*, pages 27–40, 2012. 2, 5, 6, 7, 8
- [17] D. Krishnan, T. Tay, and R. Fergus. Blind deconvolution using a normalized sparsity measure. In *CVPR*, pages 2657–2664, 2011. 1, 2, 5, 6, 7, 8
- [18] D. Lemire. Streaming maximum-minimum filter using no more than three comparisons per element. *Nordic Journal of Computing*, 13(4):328–339, 2006. 8
- [19] A. Levin, Y. Weiss, F. Durand, and W. T. Freeman. Understanding and evaluating blind deconvolution algorithms. In *CVPR*, pages 1964–1971, 2009. 1, 2, 5, 7, 8
- [20] A. Levin, Y. Weiss, F. Durand, and W. T. Freeman. Efficient marginal likelihood optimization in blind deconvolution. In *CVPR*, pages 2657–2664, 2011. 1, 2, 4, 5, 6
- [21] Y. Lou, A. L. Bertozzi, and S. Soatto. Direct sparse deblurring. *Journal of Mathematical Imaging and Vision*, 39(1):1–12, 2011. 2
- [22] L. B. Lucy. An iterative technique for the rectification of observed distributions. *Astronomy Journal*, 79(6):745–754, 1974. 1
- [23] D. Martin, C. Fowlkes, D. Tal, and J. Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *ICCV*, pages 416–423, 2001. 3
- [24] T. Michaeli and M. Irani. Blind deblurring using internal patch recurrence. In *ECCV*, pages 783–798, 2014. 1, 5, 7
- [25] J. Pan, Z. Hu, Z. Su, and M.-H. Yang. Deblurring face images with exemplars. In *ECCV*, pages 47–62, 2014. 2, 5, 6
- [26] J. Pan, Z. Hu, Z. Su, and M.-H. Yang. Deblurring text images via L_0 -regularized intensity and gradient prior. In *CVPR*, pages 2901–2908, 2014. 1, 2, 3, 4, 5, 6, 7, 8
- [27] Q. Shan, J. Jia, and A. Agarwala. High-quality motion deblurring from a single image. *ACM SIGGRAPH*, 27(3):73, 2008. 1, 2, 4
- [28] Q. Shan, W. Xiong, and J. Jia. Rotational motion deblurring of a rigid object from a single image. In *ICCV*, pages 1–8, 2007. 5
- [29] L. Sun, S. Cho, J. Wang, and J. Hays. Edge-based blur kernel estimation using patch priors. In *ICCP*, 2013. 2, 5
- [30] Y.-W. Tai, P. Tan, and M. S. Brown. Richardson-lucy deblurring for scenes under a projective motion path. *IEEE TPAMI*, 33(8):1603–1618, 2011. 5
- [31] H. Takeda, S. Farsiu, and P. Milanfar. Deblurring using regularized locally adaptive kernel regression. *IEEE TIP*, 17(4):550–563, 2008. 2
- [32] Y. Wang, J. Yang, W. Yin, and Y. Zhang. A new alternating minimization algorithm for total variation image reconstruction. *SIAM Journal on Imaging Sciences*, 1(3):248–272, 2008. 4
- [33] O. Whyte, J. Sivic, A. Zisserman, and J. Ponce. Non-uniform deblurring for shaken images. *IJCV*, 98(2):168–186, 2012. 5, 7
- [34] L. Xu and J. Jia. Two-phase kernel estimation for robust motion deblurring. In *ECCV*, pages 157–170, 2010. 1, 2, 5, 6
- [35] L. Xu, C. Lu, Y. Xu, and J. Jia. Image smoothing via L_0 gradient minimization. In *SIGGRAPH Asia*, volume 30, page 174, 2011. 4
- [36] L. Xu, S. Zheng, and J. Jia. Unnatural L_0 sparse representation for natural image deblurring. In *CVPR*, pages 1107–1114, 2013. 1, 2, 3, 4, 5, 6, 7, 8
- [37] H. Zhang, J. Yang, Y. Zhang, and T. S. Huang. Sparse representation based blind image deblurring. In *ICME*, pages 1–6, 2011. 2
- [38] D. Zoran and Y. Weiss. From learning models of natural image patches to whole image restoration. In *ICCV*, pages 479–486, 2011. 5