MRF-based Blind Image Deconvolution

Nikos Komodakis[1] and Nikos Paragios[2]

[1] Ecole des Ponts ParisTech, nikos.komodakis@enpc.fr
[2] Ecole Centrale de Paris, nikos.paragios@ecp.fr

Abstract. This paper proposes an optimization-based blind image deconvolution method. The proposed method relies on imposing a discrete MRF prior on the deconvolved image. The use of such a prior leads to a very efficient and powerful deconvolution algorithm that carefully combines advanced optimization techniques. We demonstrate the extreme effectiveness of our method[3] by applying it on a wide variety of very challenging cases that involve the inference of large and complicated blur kernels.

1 Introduction

Blind image deconvolution is a fundamental but very challenging problem, which has a long history in the image and signal processing literature [1]. Perhaps its most well known use is for removing the blur from consumer photographs (*e.g.*, due to camera shake), but it also has important applications in areas such as computational photography and astronomical imaging. The input to this problem consists of a degraded image $\mathbf{I}$ that equals the convolution of a true image $\mathbf{x}$ with a kernel $\mathbf{k}$ plus some noise $\mathbf{n}$, or

$$\mathbf{I} = \mathbf{x} \otimes \mathbf{k} + \mathbf{n}, \tag{1}$$

where $\otimes$ denotes the convolution operator. Given only the image $\mathbf{I}$ as input, the goal of blind image deconvolution is to inverse the above process and to recover both $\mathbf{x}$ and $\mathbf{k}$, which are assumed to be the unknowns in this case.

Over the past years, the problem of blind image deconvolution has attracted a significant amount of attention from the computer vision and image processing community, and thus a variety of algorithms [2–7] have been proposed that try to contribute to the state of the art in various ways. Obviously, one of the main difficulties of blind deconvolution relates to the fact that there can be exponentially many images $\mathbf{x}$ and kernels $\mathbf{k}$ that satisfy equation (1), which, in other words, means that inverting the above equation is a severely ill-posed problem [8]. As a result, a number of different strategies have been employed in order to deal with this issue. For instance, a variety of regularization based methods [2–4, 9–11] have been presented recently, which try to make use of statistical priors of natural images in order to narrow down the set of possible solutions

[3] To encourage others to experiment with our algorithm, a Matlab implementation will become publicly available from http://imagine.enpc.fr/~komodakn/deblurring

to eq. (1). A characteristic example of such a prior relates to the well known property of natural images that its gradient follows a heavy tailed distribution, which naturally leads to imposing a sparsity constraint on the gradient of image $\mathbf{x}$. Methods of this type have recently been able to show impressive deconvolution results. On the other hand, many of these approaches rely on sophisticated energy minimization methods [2], [3], and as a result they can often carry a high computational cost. Another approach to handle the ill-posedness of eq. (1) is to employ a marginalization method (instead of a MAP estimation algorithm) for integrating out the unknown kernel $\mathbf{k}$. Such a method has recently been proposed by Levin *et al.* in [12], [13].

Besides all the single-image based approaches mentioned above, another strand of deconvolution methods relies instead on making use of multiple images (possibly of different modalities) and/or specific hardware during image capturing [14–17]. A characteristic example of this trend is the so-called epsilon-photography [18] techniques (which have also been applied to the image deconvolution problem), in which case multiple images of the scene are captured by an epsilon variation on the settings of the camera. All these approaches make it much easier to infer the blur kernel $\mathbf{k}$ and thus have been shown to yield greatly improved results. On the other hand, due to the additional input and/or specific hardware requirements of these methods, their applicability is more limited compared to approaches that rely on single images. Also, it is important to note that over the last years a number of works have appeared that study the quite challenging problem of image deconvolution with spatially-varying blur (caused *e.g.*, due to camera rotation) [19–25]. Furthermore, besides blind image deconvolution, there has also been a considerable amount of work on the problem of non-blind deconvolution [26]. This is, of course, an easier task given that the blur kernel is known in advance. Yet, very impressive results have recently been demonstrated on this problem for the quite challenging case of blur kernels with very large spatial support [27].

Our apporach: The core difficulty of blind image deconvolution stems from the fact that the kernel $\mathbf{k}$ is unknown. This leads to the number of measurements being much lower than the number of unknowns, thus contributing to the ill-posedness of the problem and also severely complicating the deconvolution process. To deal with this problem, one therefore needs to impose additional constraints on the structure of the deconvolved image. To illustrate what kind of constraints are going to be imposed by our method and to also motivate our approach, we will use the example presented in Fig. 1. We there show two images, $\mathbf{x}$ and $\bar{\mathbf{x}}$, where $\bar{\mathbf{x}}$ is essentially a quantized version of image $\mathbf{x}$ (the total number of distinct colors in $\bar{\mathbf{x}}$ is not more than 15). Yet, notice that the convolution of these two images with the same kernel $\mathbf{k}$ generates the blurred images $\mathbf{I}$ and $\bar{\mathbf{I}}$ respectively, which are also shown in Fig. 1. As can be seen, despite the fact that $\bar{\mathbf{x}}$ is a vastly simpler image than $\mathbf{x}$, the corresponding blurred images $\mathbf{I}$ and $\bar{\mathbf{I}}$ are visually quite similar. Our approach will try to exploit this fact in order to implicitly reduce the ill-posedness of blind deconvolution. To this end, it proceeds by trying to concurrently estimate the unknown kernel $\mathbf{k}$ and a deconvolved im-

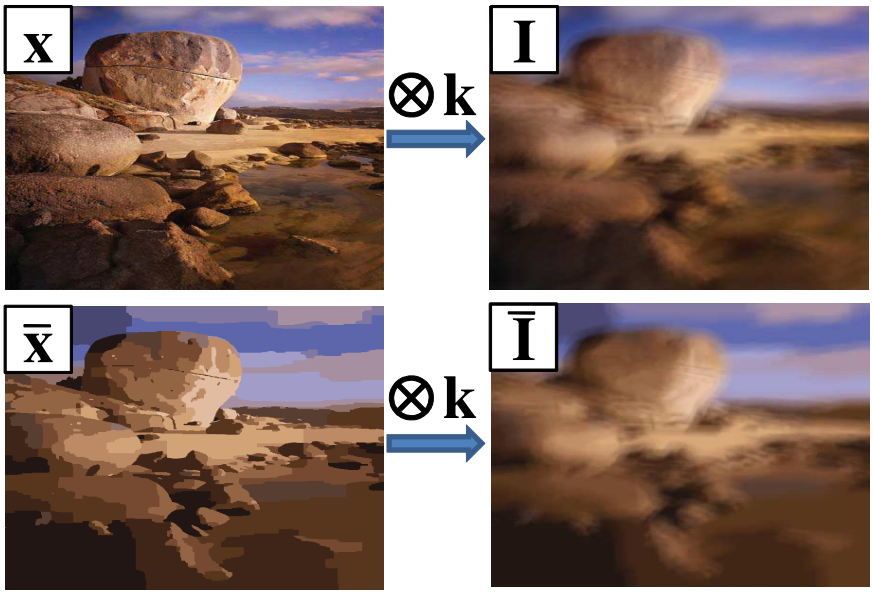

Fig. 1: The blurred images generated from image $\mathbf{x}$ and its quantized version $\bar{\mathbf{x}}$ are visually quite similar.

age. However, instead of aiming to directly compute a deconvolved image that coincides with the complex natural image $\mathbf{x}$, our approach focuses on estimating the much simpler deconvolved image $\bar{\mathbf{x}}$. The latter image is supposed to provide a coarse approximation to image $\mathbf{x}$ but at the same time it has a much simpler structure as it is assumed to satisfy the following constraints: (i) it contains only a limited number of colors, (ii) and is a piecewise constant image.

Intuitively, the role of $\bar{\mathbf{x}}$ is two-fold: on the one hand, by making use of a deconvolved image $\bar{\mathbf{x}}$ that has such a simple structure, what we are effectively doing is to significantly reduce the number of unknowns in our problem, which in turn helps us to reduce the ill-posedness of deconvolution and to accurately estimate the unknown kernel $\mathbf{k}$, which is a key issue in image deconvolution. Essentially, $\bar{\mathbf{x}}$ is able to provide a segmentation of the deconvolved image $\mathbf{x}$, thus revealing its structure, which is known to be crucial for the successful estimation of $\mathbf{k}$. At the same time, as will become clear shortly, the quantized deconvolved image $\bar{\mathbf{x}}$ can also helps us to estimate the true deconvolved image $\mathbf{x}$.

We conclude this introduction by summarizing the main contributions of this work, which are the following: (i) We propose a novel discrete MRF-based prior for image deconvolution, whose role is to impose the above mentioned simple structure on image $\bar{\mathbf{x}}$. (ii) Furthermore, based on this, we develop a very efficient and easily implementable algorithmic scheme for blind deconvolution that carefully combines advanced optimization techniques such as fast inference for discrete MRFs and the alternating direction method of multipliers . (iii) Last,

despite the computational efficiency of our proposed scheme, we can use it to easily handle very challenging deblurring problems that involve even large and complicated blur kernels.

2 Blind Image Deconvolution Scheme

In this section, we proceed to describe our blind deconvolution algorithm in detail. At a high level, our method uses an iterative optimization scheme that alternates between the following two main steps:

- **Step 1:** estimating a quantized image $\bar{\mathbf{x}}$ that obeys a discrete MRF prior given the blur kernel $\mathbf{k}$ (section 2.1). During this step, we are also able to compute an estimation of the deconvolved and non-quantized image $\mathbf{x}$.
- **Step 2:** estimating the kernel $\mathbf{k}$ given a quantized deconvolved image $\bar{\mathbf{x}}$ (section 2.2),

Based on the above, the following variables are maintained during the execution of our algorithm: $\mathbf{k}$, $\bar{\mathbf{x}}$, and the auxiliary variable $\mathbf{x}$. Furthermore, in order to accelerate the convergence of our iterative scheme and to cope with the case of large blur kernels $\mathbf{k}$, we apply our algorithm in a multiresolution fashion (section 2.3).

It is important to note at this point that the estimation of image $\bar{\mathbf{x}}$ plays an absolutely crucial role in the whole process. On the one hand, it is exactly due to $\bar{\mathbf{x}}$ and its constrained image structure that we are able to compute a blur kernel $\mathbf{k}$ that converges to the correct kernel during step 2. If instead we choose to use the non-quantized image $\mathbf{x}$ during that step, the whole process breaks down and leads to computing the trivial no-blur solution (*i.e.*, to the identity kernel). On the other hand, $\bar{\mathbf{x}}$ also helps us to correctly estimate the non-quantized image $\mathbf{x}$ during step 1 of our method. In the following we describe the main steps of our deconvolution scheme in more detail.

2.1 Estimating quantized image $\bar{\mathbf{x}}$ given the kernel $\mathbf{k}$

Given a fixed kernel $\mathbf{k}$, to estimate the quantized image $\bar{\mathbf{x}}$ we solve the following optimization problem

$$\min_{\bar{\mathbf{x}} \in \mathcal{L}_{\mathbf{n}}, \mathbf{x}} \|\mathbf{k} \otimes \mathbf{x} - \mathbf{I}\|^2 + \lambda \|\nabla \mathbf{x}\|^2 + \mu \|\bar{\mathbf{x}} - \mathbf{x}\|^2 + \sum_{pq \in \mathcal{E}} w_{pq} [\bar{x}_p \neq \bar{x}_q] \qquad (2)$$

In the above formula, the symbol $\otimes$ denotes the convolution operator, $\mathcal{L}_n$ denotes the set of quantized images that contain at most n different colors/intensities, $\mathcal{E}$ denotes the set of pairs of neighboring[4] pixels of image $\bar{\mathbf{x}}$, and $[\cdot]$ equals 1 if the expression inside the brackets is true and zero otherwise. At a high level, the above functional has the following interpretation: On the one hand, the first

[4] Here we assume 4-connectivity, but in general any higher-connectivity can be used as well.

two terms express the fact that $\mathbf{x}$ must be a natural image corresponding to the deconvolution of image $\mathbf{I}$ with respect to blur kernel $\mathbf{k}$. On the other hand, the last two terms (along with the constraint $\bar{\mathbf{x}} \in \mathcal{L}_n$) express the fact that $\bar{\mathbf{x}}$ should correspond to a piecewise-constant quantization of the aforementioned deconvolved image $\mathbf{x}$ (the latter is achieved through constraining the structure of image $\bar{\mathbf{x}}$ by imposing a discrete MRF prior on it).

More specifically, the terms in (2) have the following meaning:

- The data term $\|\mathbf{k} \otimes \mathbf{x} - \mathbf{I}\|^2$ expresses the fact that the convolution of $\mathbf{x}$ with $\mathbf{k}$ should be close to the observed image $\mathbf{I}$. This corresponds to a typical negative log-likelihood term under the assumption of an additive Gaussian noise on $\mathbf{I}$.
- The term $\lambda\|\nabla\mathbf{x}\|^2$ corresponds to a standard gradient-based regularizer for natural images.
- The term $\mu\|\bar{\mathbf{x}} - \mathbf{x}\|^2$ along with the constraint $\bar{\mathbf{x}} \in \mathcal{L}_n$ expresses the fact that $\bar{\mathbf{x}}$ should form a quantization of image $\mathbf{x}$ with at most n colors.
- Last, the term $\sum_{pq\in\mathcal{E}} w_{pq}[\bar{x}_p \neq \bar{x}_q]$ corresponds to a Markov Random Field (MRF) prior imposed on $\bar{\mathbf{x}}$ in the form of a weighted Potts model, which is used for ensuring the piecewise constancy assumption on $\bar{\mathbf{x}}$. Regarding the weights $\{w_{pq}\}$, they can be adjusted based on the gradient of the continuous image $\mathbf{x}$ (*i.e.*, a standard contrast sensitive weight term can be used) so that the color changes in $\bar{\mathbf{x}}$ can be better aligned with the edges of image $\mathbf{x}$.

To solve optimization problem (2), we apply alternating minimization between $\bar{\mathbf{x}}$ and $\mathbf{x}$ as follows.

(a) Solving for $\bar{\mathbf{x}}$: For fixed $\mathbf{x}$, the minimization of (2) over $\bar{\mathbf{x}}$ reduces to

$$\min_{\bar{\mathbf{x}}\in\mathcal{L}_n} \; \mu\|\bar{\mathbf{x}} - \mathbf{x}\|^2 + \sum_{pq\in\mathcal{E}} w_{pq}[\bar{x}_p \neq \bar{x}_q] \;\; . \tag{3}$$

Since image $\bar{\mathbf{x}}$ should approximate image $\mathbf{x}$, to determine the set of n possible colors/intensities for $\bar{\mathbf{x}}$ we apply k-means clustering (with $k = n$) to the colors/intensities of the current image $\mathbf{x}$, and denote the resulting set of cluster centers by $L_{\mathbf{x}}$. Optimization problem (3) is then equivalent to minimizing the energy of a discrete MRF with unary potentials $V_p(\bar{x}_p) = \mu\|\bar{x}_p - x_p\|^2$, pairwise potentials $V_{pq}(\bar{x}_p, \bar{x}_q) = w_{pq}[\bar{x}_p \neq \bar{x}_q]$, and label set $L_{\mathbf{x}}$. To optimize this MRF, we use the FastPD[5] algorithm [28]. Due to the very high efficiency of that method, optimization of (3) proceeds extremely fast, while the estimated solution $\bar{\mathbf{x}}$ is guaranteed to be almost optimal. Intuitively, the update of $\bar{\mathbf{x}}$ via this MRF optimization step can be roughly viewed as some sort of smooth quantization of the continuous image $\mathbf{x}$.

(b) Solving for $\mathbf{x}$: For fixed $\bar{\mathbf{x}}$, minimization of (2) over $\mathbf{x}$ reduces to the following continuous optimization problem

$$\min_{\mathbf{x}} \|\mathbf{k} \otimes \mathbf{x} - \mathbf{I}\|^2 + \lambda\|\nabla\mathbf{x}\|^2 + \mu\|\mathbf{x} - \bar{\mathbf{x}}\|^2. \tag{4}$$

[5] Publicly available implementation from `http://www.csd.uoc.gr/~komod/FastPD`

The above functional can be rewritten as

$$\min_{\mathbf{x}} \|\mathbf{M_k x} - \mathbf{I}\|^2 + \lambda(\|\mathbf{M}_{\mathbf{g}_h}\mathbf{x}\|^2 + \|\mathbf{M}_{\mathbf{g}_v}\mathbf{x}\|^2) + \mu\|\mathbf{x} - \bar{\mathbf{x}}\|^2 ,$$

where $\mathbf{M_k}$, $\mathbf{M}_{\mathbf{g}_h}$, $\mathbf{M}_{\mathbf{g}_v}$ denote respectively the convolution matrices for the filter $\mathbf{k}$, the horizontal derivative filter $\mathbf{g}_h$ and the vertical derivative filter $\mathbf{g}_v$ (note that in the above formula $\mathbf{x}$, $\bar{\mathbf{x}}$, and $\mathbf{I}$ are assumed to represent vectorized images).

This, in turn, leads to having to solve the following linear system of equations

$$\Big(\mathbf{M}_{\mathbf{k}}^T\mathbf{M_k} + \lambda(\mathbf{M}_{\mathbf{g}_h}^T\mathbf{M}_{\mathbf{g}_h} + \mathbf{M}_{\mathbf{g}_v}^T\mathbf{M}_{\mathbf{g}_v}) + \mu\Big)\mathbf{x} = \mathbf{M}_{\mathbf{k}}^T\mathbf{I} + \mu\bar{\mathbf{x}} ,$$

which can be done very efficiently by using the Fast Fourier Transform and working directly in the frequency domain as follows

$$\mathbf{x} = \mathcal{F}^{-1}\Big(\frac{\overline{\mathcal{F}(\mathbf{k})} \circ \mathcal{F}(\mathbf{I}) + \mu\mathcal{F}(\bar{\mathbf{x}})}{|\mathcal{F}(\mathbf{k})|^2 + \lambda(|\mathcal{F}(\mathbf{g}_h)|^2 + |\mathcal{F}(\mathbf{g}_v)|^2) + \mu}\Big) . \tag{5}$$

In the above formula $\mathcal{F}^{-1}$ and $\mathcal{F}$ denote the inverse and forward DFT operator respectively. It should be noted that solving the above linear system using formula (5) is very fast but can sometimes introduce ringing artifacts (due to boundary effects). Alternatively, one can solve the above system in the spatial domain through a preconditioned conjugate gradient algorithm in order to avoid such artifacts (at the expense of an increased computational cost).

To summarize, one round of updates of $(\bar{\mathbf{x}}, \mathbf{x})$ for minimizing (2) reduces to the following operation

$$\text{UPDATE OF } (\bar{\mathbf{x}}, \mathbf{x}) := \begin{cases} \text{optimize MRF energy (3)} , \\ \text{apply formula (5)} \quad . \end{cases} \tag{6}$$

2.2 Estimating the kernel k given the quantized image $\bar{\mathbf{x}}$

To update the blur kernel $\mathbf{k}$ given the current image $\bar{\mathbf{x}}$, we solve the following convex optimization problem

$$\min_{\mathbf{k}} \|\bar{\mathbf{x}} \otimes \mathbf{k} - \mathbf{I}\|^2 + \tau\|\mathbf{k}\|_1 \quad . \tag{7}$$

The first term in (7) expresses the fact that the convolution of the quantized image $\bar{\mathbf{x}}$ with $\mathbf{k}$ should be close to the observed image $\mathbf{I}$. The second term in (7) is an ℓ_1-norm penalty term, which corresponds to a Laplacian prior used for imposing sparsity on the unknown convolution kernel $\mathbf{k}$. The sparsity of $\mathbf{k}$ is a reasonable assumption, since, in practice, the most common and challenging cases for image deconvolution are due to camera motion, where the kernel essentially coincides with the camera path.

Before proceeding, we want to emphasize at this point the importance of using the quantized image $\bar{\mathbf{x}}$ during this step. The fact that image $\bar{\mathbf{x}}$ is required to have a simple structure (through imposing the aforementioned discrete MRF prior

on it) regularizes the whole process and plays a crucial role for estimating the correct $\mathbf{k}$. Essentially, as will also become clear from the experimental results, $\bar{\mathbf{x}}$ manages to reveal the structure of the unknown deconvolved image by providing a segmentation of it, which is known to play a crucial role in image deconvolution. For instance, if we were to use image $\mathbf{x}$ (instead of $\bar{\mathbf{x}}$) in (7), the resulting overall iterative scheme would typically converge to the trivial no-blur solution [12].

If we let $\mathbf{M}_{\bar{\mathbf{x}}}$ denote the convolution matrix for $\bar{\mathbf{x}}$ (*i.e.*, $\mathbf{M}_{\bar{\mathbf{x}}}\mathbf{k} = \bar{\mathbf{x}} \otimes \mathbf{k}$), then problem (7) can be rewritten as follows

$$\min_{\mathbf{k}} \ \|\mathbf{M}_{\bar{\mathbf{x}}}\mathbf{k} - \mathbf{I}\|^2 + \tau\|\mathbf{k}\|_1 \ . \tag{8}$$

Although (8) can be solved via an interior point method, this leads to an algorithm with very high computational cost given that problem (8) must be solved multiple times during deconvolution. Instead, we apply the well known Alternating Direction Method of Multipliers (ADMM) [29–31], which allows us to decouple problem (8) to a series of highly efficient optimization tasks through the introduction of an auxiliary replicating variable $\mathbf{k}'$ as described next.

Alternating Direction Method of Multipliers (ADMM). By using such a variable $\mathbf{k}'$, problem (8) can be rewritten as

$$\min_{\mathbf{k}',\mathbf{k}} \ g_1(\mathbf{k}') + g_2(\mathbf{k}) \tag{9}$$

$$\text{s.t. } \mathbf{k}' = \mathbf{k} \ , \tag{10}$$

where $g_1(\mathbf{k}') = \|\mathbf{M}_{\bar{\mathbf{x}}}\mathbf{k}' - \mathbf{I}\|^2$, $g_2(\mathbf{k}) = \tau\|\mathbf{k}\|_1$. ADMM is essentially a variant of the Augmented Lagrangian method. The latter provably solves problems like (9) via an iterative scheme that is equivalent to

$$(\mathbf{k}'_{i+1}, \mathbf{k}_{i+1}) \in \arg\min_{\mathbf{k}',\mathbf{k}} g_1(\mathbf{k}') + g_2(\mathbf{k}) + \frac{b}{2}\|\mathbf{k}' - \mathbf{k} - \mathbf{a}_i\|^2 \tag{11}$$

$$\mathbf{a}_{i+1} = \mathbf{a}_i - (\mathbf{k}'_{i+1} - \mathbf{k}_{i+1}) \ , \tag{12}$$

where $\mathbf{a}$ corresponds to a transformed vector of Lagrange multipliers, and b is a fixed penalty parameter.

An efficient solution to (11) is hard due to the non-separable quadratic term. To overcome this, ADMM applies to (11) a single step of alternating minimization between $\mathbf{k}'$ and $\mathbf{k}$, which allows a nice decoupling of the problem into

$$\mathbf{k}'_{i+1} \in \arg\min_{\mathbf{k}'} g_1(\mathbf{k}') + \frac{b}{2}\|\mathbf{k}' - \mathbf{k}_i - \mathbf{a}_i\|^2 \tag{13}$$

$$\mathbf{k}_{i+1} \in \arg\min_{\mathbf{k}} g_2(\mathbf{k}) + \frac{b}{2}\|\mathbf{k}'_{i+1} - \mathbf{k} - \mathbf{a}_i\|^2 \tag{14}$$

$$\mathbf{a}_{i+1} = \mathbf{a}_i - (\mathbf{k}'_{i+1} - \mathbf{k}_{i+1}) \ . \tag{15}$$

Notice now that, due to equality $g_2(\mathbf{k}) = \tau\|\mathbf{k}\|_1$, a solution to (14) is easily recognized to be $\mathbf{k}_{i+1} = \text{soft}_{\tau/b}(\mathbf{k}'_{i+1} - \mathbf{a}_i)$, where $\text{soft}_t(\mathbf{x})_i \equiv \max(|x_i| - t, 0) \cdot \text{sign}(x_i)$ denotes an efficient soft thresholding/shrinkage operator.

Furthermore, due to equality $g_1(\mathbf{k}') = \|\mathbf{M}_{\bar{\mathbf{x}}}\mathbf{k}' - \mathbf{I}\|^2$, problem (13) reduces to the linear system

$$(\mathbf{M}_{\bar{\mathbf{x}}}^T\mathbf{M}_{\bar{\mathbf{x}}} + \frac{b}{2})\mathbf{k}'_{i+1} = \mathbf{M}_{\bar{\mathbf{x}}}^T\mathbf{I} + \frac{b}{2}(\mathbf{k_i} + \mathbf{a}_i) \ , \tag{16}$$

which can be solved numerically very efficiently via a conjugate gradient (CG) method (the most expensive part of each CG iteration is the calculation $(\mathbf{M}_{\bar{\mathbf{x}}}^T\mathbf{M}_{\bar{\mathbf{x}}} + \frac{b}{2})\mathbf{k}'_{i+1}$, which requires just 3 FFTs, *i.e.*, 2 forward and 1 inverse).

To summarize, a round of updates for minimizing (8) thus consists of the following operations

$$\begin{aligned} \mathbf{k'}_{i+1} &= \text{solution of linear system (16)} \\ \mathbf{k}_{i+1} &= \text{soft}_{\tau/b}(\mathbf{k'}_{i+1} - \mathbf{a}_i) \\ \mathbf{a}_{i+1} &= \mathbf{a}_i - (\mathbf{k'}_{i+1} - \mathbf{k}_{i+1}) \ . \end{aligned} \tag{17}$$

Input: images $\{\mathbf{I}_i\}$, where $\mathbf{I}_i$ = image $\mathbf{I}$ resized by factor s_i

$(\mathbf{x}, \mathbf{k}) = (\mathbf{I}_0, \text{delta kernel})$ `// use trivial initialization`
for $i \in \{0, 1, \ldots, \#\text{scales} - 1\}$ **do** `// outer loop`
  for $j \in \{0, 1, \ldots, \#\text{inner iterations} - 1\}$ `// inner loop`
    `// Update` $(\bar{\mathbf{x}}, \mathbf{x})$
    Apply $m (m \geq 1)$ rounds of updates (6)
    `// Update` $\mathbf{k}$
    Apply $m (m \geq 1)$ rounds of updates (17)
  end for
  Upsample images $\bar{\mathbf{x}}$, $\mathbf{x}$, $\mathbf{k}$
end for

Fig. 2: Pseudocode of the blind deconvolution algorithm

2.3 Multiresolution deconvolution scheme

The iterative deconvolution algorithm described in sections 2.1 and 2.2 is used within the context of a multiresolution scheme. This essentially means that this algorithm is applied successively for deconvolving a series of resized images $\{\mathbf{I}_i\}$, where $\mathbf{I}_i$ represents image $\mathbf{I}$ resized by a factor $s_i \leq 1$ (we start from a low resolution image $\mathbf{I}_0$ and proceed with images of finer resolution until we reach the last image in the series that coincides with the input image $\mathbf{I}$). As explained in sections 2.1 and 2.2, to solve the problem at the i-th level we alternate a few times between updates (6) of $(\bar{\mathbf{x}}, \mathbf{x})$ and updates (17) of $\mathbf{k}$. The final images $\bar{\mathbf{x}}$, $\mathbf{x}$, and $\mathbf{k}$ estimated by the algorithm at level i are then upsampled and used as initialization for the problem at the next finer level $i+1$. This process leads to a gradual refinement of the variables $\bar{\mathbf{x}}$, $\mathbf{x}$, $\mathbf{k}$ (see Fig. 4) and allows the algorithm to deal even with large blur kernels $\mathbf{k}$. It also leads to a much faster convergence of the deconvolution process, which thus requires only a very small number of iterations per level. At the first level, $\mathbf{x}$ is always initialized as the coarse image $\mathbf{I}_0$, and $\mathbf{k}$ is initialized to the delta kernel. Also, the resize image factor s_0 at the

first level is always chosen small enough so that the size of the corresponding kernel $\mathbf{k}$ is small as well (typically 3×3). We show the overall structure of the resulting algorithm in Fig. 2.

3 Experimental results

We next proceed to present some experimental results of our deconvolution method. Regarding the choice of the parameters of our algorithm, the following values have been used in all of the experiments: $\lambda = \mu = 0.4 \cdot 10^{-3}$ and $\tau = 10^{-3}$. Also, with regard to the number of colors n of the quantized image $\bar{\mathbf{x}}$, we found that a value of $n = 10$ was typically enough. We should note at this point that our method has been quite robust with respect to the choice of all of these parameters.

We display in Fig. 3 deconvolution results computed by our method for some quite challenging cases, where very large and complicated blur kernels need to be inferred. Notice, for instance, that the recovered images $\mathbf{x}$ are very sharp and contain all the fine details of the objects appearing in them.

To better illustrate how our algorithm works, we also provide more detailed results in Fig. 4. That figure shows (for 2 cases) how the kernel $\mathbf{k}$, and the images $\bar{\mathbf{x}}$, $\mathbf{x}$ vary during the course of the algorithm. Notice, for instance, how the estimated quantized image $\bar{\mathbf{x}}$ manages to gradually conform to the underlying structure of the true deblurred image, thus essentially providing a rough segmentation of the latter and also helping $\mathbf{x}$ to converge to the correct solution. This is achieved in a robust and completely automatic manner that is driven entirely by our iterative optimization scheme. The role of the automatically quantized image $\bar{\mathbf{x}}$ is crucial in this regard. On the one hand, it prevents the estimated blur kernel $\mathbf{k}$ from converging to an incorrect result such as the trivial solution corresponding to the identity kernel. On the other hand, it also helps $\mathbf{x}$ to avoid drifting to the wrong image structure during our deconvolution scheme. The robustness of our algorithm is also evidenced from the fact that it does not require from the user to manually select an image patch that will then be used as input for the blind deconvolution process, as done in [2]. Instead, our method simply uses as input the whole blurred image.

Fig. 4 also includes energy plots that show how the objective function (7) (which measures the dissimilarity between the convolution $\bar{\mathbf{x}} \otimes \mathbf{k}$ and the input image $\mathbf{I}$) varies during the course of our algorithm. Notice how quickly this function decreases, thus also verifying the effectiveness of our deconvolution procedure. We should also mention that no special initialization for $\mathbf{k}$ is required by our method, which, as already explained, always uses the delta kernel as initial kernel.

Further results and comparisons. Fig. 5 shows more results of our method on 3 other challenging examples from [2], [3] and [27]. As can be seen, they all require estimating very large blur kernels. In fact, as has been shown in [27], examples such as the last one in Fig. 5 (that involves a 50×50 kernel) can prove to be quite challenging even for some non-blind deconvolution algorithms, *i.e.*, when the blur kernel is known in advance. Note that, as mentioned earlier in

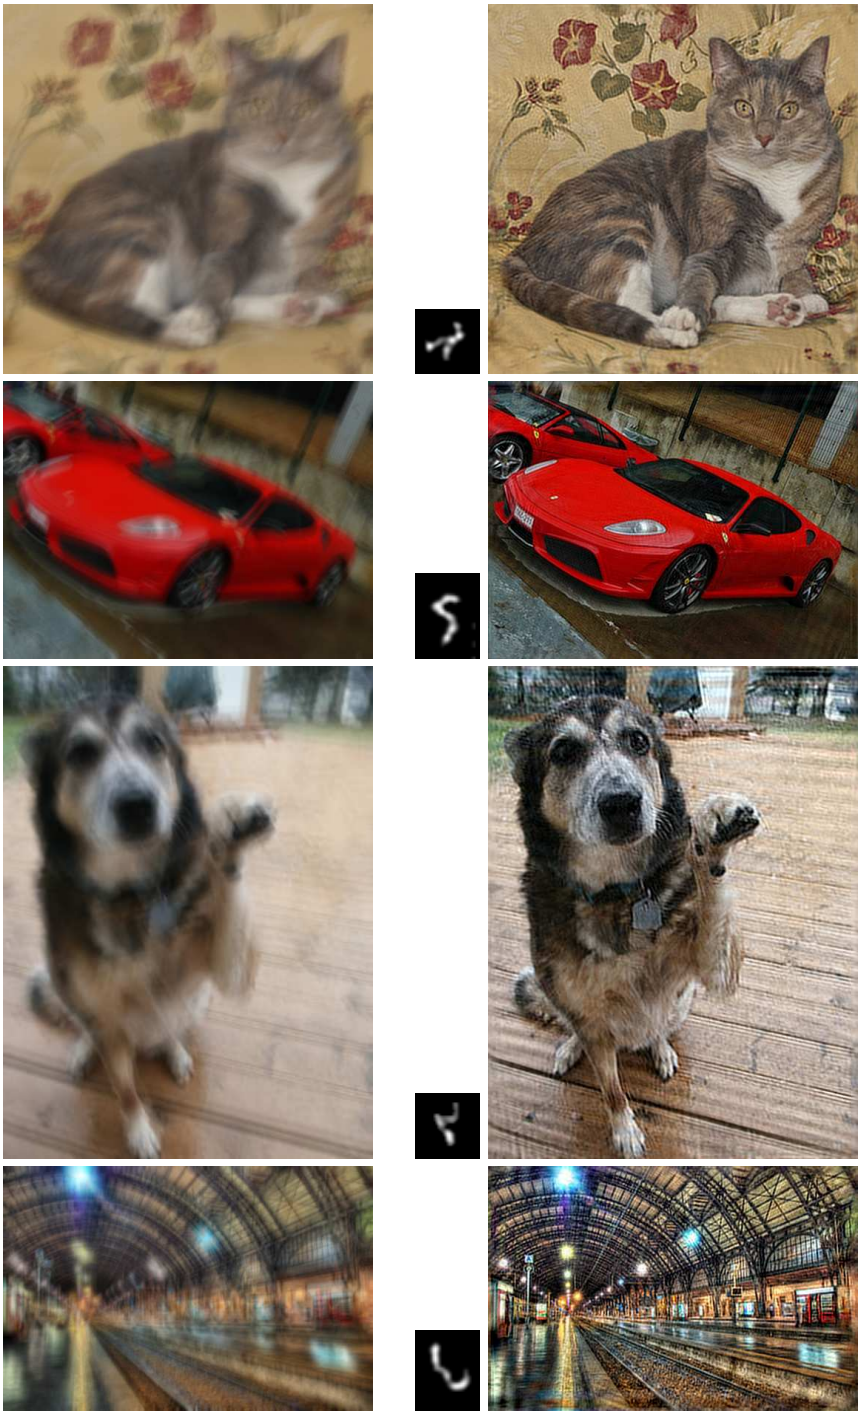

Fig. 3: Left: input images. **Right:** kernels **k** and deconvolved images **x** as computed by our method. The size (in pixels) of kernel **k** was, from top to bottom: 40×40, 29×38, 35×35, 30×37.

Fig. 4: We show for the first two examples from Fig. 3 a few snapshots of $\mathbf{k}$, $\bar{\mathbf{x}}$ and $\mathbf{x}$ as estimated during the course of our deconvolution algorithm. We also show corresponding energy plots (where red dashed lines indicate a change in resolution).

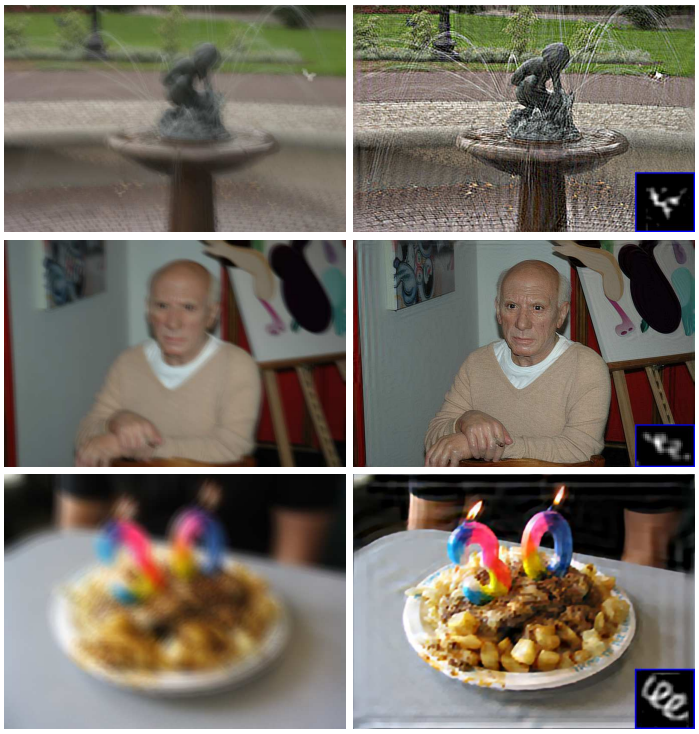

Fig. 5: Left: input images. **Right:** kernels $\mathbf{k}$ and images $\mathbf{x}$ as estimated by our method. The following sizes were used in our method for kernel $\mathbf{k}$, from top to bottom: 50×50, 30×20, 50×50.

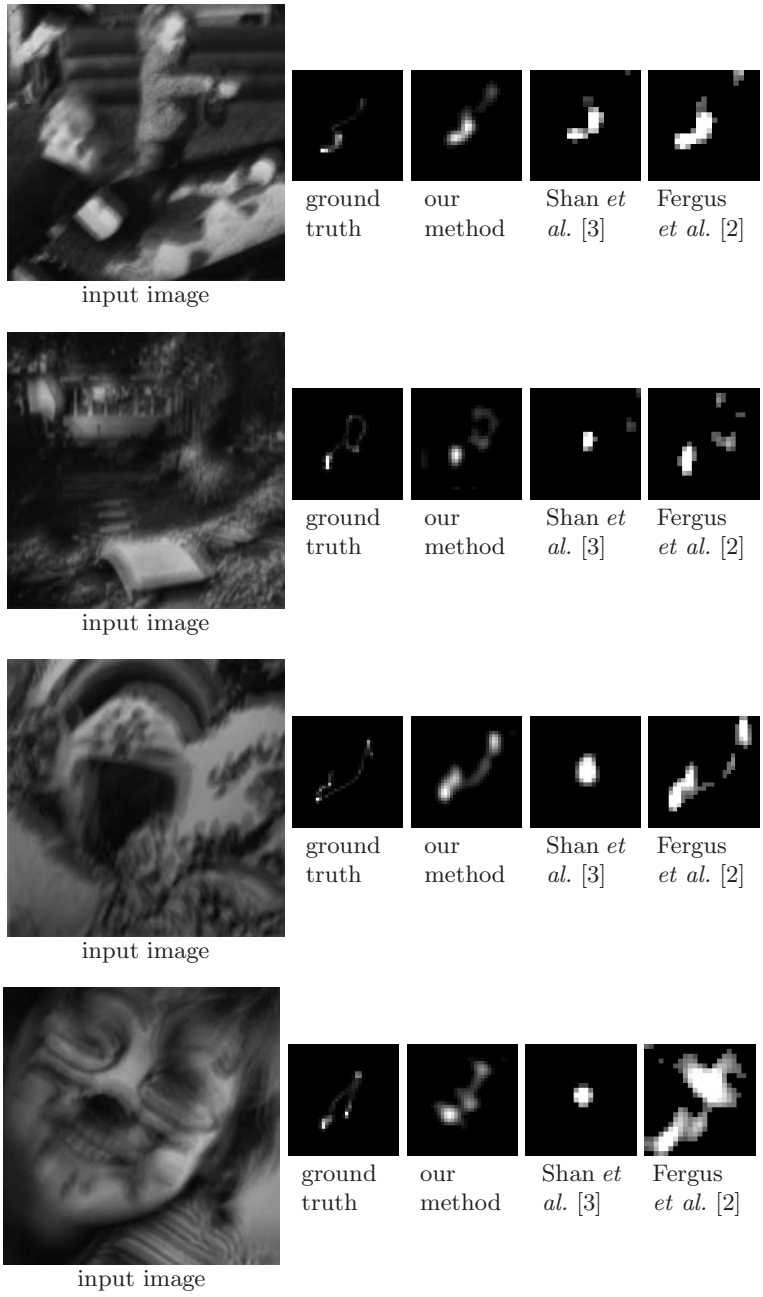

Fig. 6: Estimated kernels by our method, and the methods of [3], [2] for input images from the dataset in [12].

section 2.1, one could further improve our deconvolution results in this case by solving (4) in the spatial domain (in order to suppress any ringing artifacts due to boundary effects from the use of large kernels). Moreover, to avoid all ringing artifacts that appear in the case of large blur kernels one can apply an advanced non-blind deconvolution method such as [27].

Finally, we compare our method with the algorithms of Fergus *et al.* [2] and Shan *et al.* [3] on the established benchmark dataset introduced by [12] (in this case the results shown for the algorithms [2] and [3] were taken from [12]). We show in Fig. 6 results for some challenging cases from that dataset (for the full set of results please check the supplemental material). Compared to the other methods, the kernels recovered by our algorithm resemble the ground truth more closely. Moreover, compared to those methods, our algorithm is significantly less computationally intensive. For instance, although the MATLAB implementation of our algorithm has not been fully optimized, its running times for the examples of Fig. 6 were only 6.8 seconds, 7.2 seconds, 7.7 seconds, and 7.5 seconds respectively. The kernel sizes (in pixels) that were used for these examples were 30×30, 34×34, 38×38, and 34×34, while the input images in Fig. 6 had size 255×255.

4 Conclusions

In this paper we presented a powerful MRF-based blind deconvolution algorithm. It relies crucially on estimating (along with the kernel $\mathbf{k}$) a deconvolved image $\bar{\mathbf{x}}$, which is quantized and piecewise smooth, and thus has a much simpler structure. We have shown experimentally that our method is able to produce state of the art deconvolution results and can successfully handle very challenging cases with large and complicated blur kernels. At the same time we have shown that it is very fast since it relies on a highly efficient optimization-based deconvolution scheme that we derived.

References

1. Kundur, D., Hatzinakos, D.: Blind image deconvolution. IEEE Signal Processing Magazine (1996)
2. Fergus, R., Singh, B., Hertzmann, A., Roweis, S.T., Freeman, W.T.: Removing camera shake from a single photograph. SIGGRAPH (2006)
3. Shan, Q., Jia, J., Agarwala, A.: High-quality motion deblurring from a single image. SIGGRAPH (2008)
4. Joshi, N., Zitnick, C.L., Szeliski, R., Kriegman, D.J.: Image deblurring and denoising using color priors. In: CVPR. (2009)
5. Cho, S., Lee, S.: Fast motion deblurring. In: SIGGRAPH ASIA. (2009)
6. Xu, L., Jia, J.: Two-phase kernel estimation for robust motion deblurring. In: ECCV. (2010) 157–170
7. Babacan, S.D., Molina, R., Katsaggelos, A.K.: Variational bayesian blind deconvolution using a total variation prior. IEEE Trans. on Image Processing **18** (2009) 12–26
8. Campisi, P., Egiazarian, K.: Blind Image Deconvolution: Theory and Applications. CRC Press (2007)

9. Jia, J.: Single image motion deblurring using transparency. In: CVPR. (2007)
10. Levin, A.: Blind motion deblurring using image statistics. In: NIPS. (2006)
11. Joshi, N., Szeliski, R., Kriegman, D.J.: Psf estimation using sharp edge prediction. In: CVPR. (2008)
12. Levin, A., Weiss, Y., Durand, F., Freeman, W.: Understanding and evaluating blind deconvolution algorithms. In: CVPR. (2009) 1964–1971
13. Levin, A., Weiss, Y., Durand, F., Freeman, W.T.: Efficient marginal likelihood optimization in blind deconvolution. In: CVPR. (2011)
14. Yuan, L., Sun, J., Quan, L., Shum, H.Y.: Image deblurring with blurred/noisy image pairs. In: SIGGRAPH. (2007)
15. Joshi, N., Kang, S.B., Zitnick, C.L., Szeliski, R.: Image deblurring using inertial measurement sensors. ACM Trans. Graph. **29** (2010) 30:1–30:9
16. Raskar, R., Agrawal, A., Tumblin, J.: Coded exposure photography: motion deblurring using fluttered shutter. SIGGRAPH (2006) 795–804
17. Levin, A., Fergus, R., Durand, F., Freeman, W.T.: Image and depth from a conventional camera with a coded aperture. In: SIGGRAPH. (2007)
18. Raskar, R., Tubmlin, J., Mohan, A., Agrawal, A., Li, Y.: Computational photography. In: EUROGRAPHICS. (2006)
19. Whyte, O., Sivic, J., Zisserman, A., Ponce, J.: Non-uniform deblurring for shaken images. In: CVPR. (2010)
20. Hirsch, M., Schuler, C.J., Harmeling, S., Schölkopf, B.: Fast removal of non-uniform camera shake. In: ICCV. (2011) 463–470
21. Harmeling, S., Hirsch, M., Schlkopf, B.: Space-variant single-image blind deconvolution for removing camera shake. In: NIPS. (2010)
22. Gupta, A., Joshi, N., Zitnick, L., Cohen, M., Curless, B.: Single image deblurring using motion density functions. In: ECCV. (2010)
23. Tai, Y., P.Tan, Brown, M.: Richardson-lucy deblurring for scenes under a projective motion path. PAMI (2011)
24. Tai, Y.W., Du, H., Brown, M.S., Lin, S.: Correction of spatially varying image and video motion blur using a hybrid camera. PAMI (2010)
25. Shan, Q., Xiong, W., Jia, J.: Rotational motion deblurring of a rigid object from a single image. In: ICCV. (2007)
26. Bar, L., Sochen, N.A., Kiryati, N.: Semi-blind image restoration via mumford-shah regularization. IEEE Transactions on Image Processing **15** (2006) 483–493
27. Yuan, L., Sun, J., Quan, L., Shum, H.Y.: Progressive inter-scale and intra-scale non-blind image deconvolution. In: SIGGRAPH. (2008)
28. Komodakis, N., Tziritas, G., Paragios, N.: Fast, approximately optimal solutions for single and dynamic MRFs. In: CVPR. (2007)
29. Eckstein, J., Bertsekas, D.P.: On the douglas-rachford splitting method and the proximal point algorithm for maximal monotone operators. Math. Program. (1992)
30. Figueiredo, M.A., Bioucas-Dias, J.M., Afonso, M.V.: Fast frame-based image deconvolution using variable splitting and constrained optimization. In: SSP. (2009)
31. Wang, Y., Yang, J., Yin, W., Zhang, Y.: A new alternating minimization algorithm for total variation image reconstruction. SIAM Journal on Imaging Sciences (2008)