

Direct Sparse Deblurring

Yifei Lou

Andrea Bertozzi

Stefano Soatto

Abstract

We propose a deblurring algorithm that explicitly takes into account the sparse characteristics of natural images and does not entail solving a numerically ill-conditioned backward-diffusion. The key observation is that the sparse coefficients that encode a given image with respect to an over-complete basis are the same that encode a blurred version of the image with respect to a modified basis. Following an “analysis-by-synthesis” approach, an explicit generative model is used to compute a sparse representation of the blurred image, and the coefficients of which are used to combine elements of the original basis to yield a restored image. We compare our algorithm against the state of the art in variational methods as well as wavelet-based algorithms.

1. Introduction

Deblurring refers to the task of “undoing” the effects of convolving¹ the data with a known kernel. A common instance occurs when an image is taken with a finite-aperture system that is not well focused, so the measured image is a blurred version of the “ideal image,” convolved with the point-spread function of the lens. Ideally, one would like to recover, or “restore,” the image as would be captured by a well-focused lens. Unfortunately, deblurring is well-known to be an ill-posed inverse problem, so small perturbations in the data (for instance noise in the measured “blurred” image) lead to large errors in the reconstruction. These artifacts are usually kept at bay by means of regularization, following the classical work of Tikhonov [17]. Several choices of such *generic regularizers* have been proposed, mostly made for mathematical convenience, some based on empirical observations on the quality of the reconstruction. Recently, the research has converged towards regularizers that enforce the statistics of natural images, which are well-known to possess highly kurtotic behavior [10] due to the presence of large homogeneous regions bounded by sharp discontinuities at visibility boundaries such as occlusions and cast shadows. Some classic regularizers, such

¹Although usually deblurring refers to non-shift-invariant kernels, we use deblurring and deconvolution interchangeably in this manuscript.

as the Total Variation, implicitly favor these kind of solutions, and remain among the most competitive deblurring algorithms to this day. Nevertheless, deblurring in this context involves solving an inverse diffusion partial differential equation (PDE). In [6], Favaro et al. have approached the problem of deblurring using a “direct” method: Rather than deconvolving the measured image and the noise that goes with it, they convolve the “model image,” which is noiseless by definition, with the known kernel. This yields a simple diffusion PDE whose (space-varying) stopping time encodes the value of the kernel. They do not, however, exploit the statistics of natural images in their solution.

A related literature stream encodes the image as a discrete array of positive numbers, approximated by linear combinations of local overcomplete bases, where the natural statistics are captured by the fact that the vector of coefficients is *sparse*, so at any location only few bases contribute to the approximation [11]. In this case, there is no need for an explicit regularizer, due to the finite dimensionality of the representation, but there is still a trade-off between fidelity of the approximation and complexity of the model. While the measured image is undoubtedly a discrete object, with quantization of both the domain and the range, the object of inference, or the “ideal image,” is best represented in the continuum, with the final discretization left only for the numerical implementation of the optimization scheme. We therefore take the continuum approach, and explicitly write cost functionals that have the ideal image as an infinite-dimensional unknown.

In this manuscript, we devise a deblurring algorithm that (a) explicitly takes into account the “sparse” natural statistics of the image as a regularizer, and (b) does not suffer from the numerical conditioning issues associated with solving an inverse diffusion PDE.

1.1. The basic idea

The idea of direct sparse deblurring is simple and can be illustrated in three steps.

First, let us pretend that the image is square-integrable and sparse in some basis defined on the entire real plane. This is a common assumption underlying most image compression algorithms, in particular those based on over-complete bases, or “dictionaries” $\{d_k\} \in \mathcal{L}^2(\mathbb{R}^2 \rightarrow \mathbb{R})$

that can be used to approximate the original image f to an arbitrary degree ² by a sparse linear combination $u \doteq \sum_{k=1}^K d_k \alpha_k \doteq D\alpha$ where $\|\alpha\|_0$ is small [5].

Now, convolving an image with a shift-invariant kernel h yields a *blurred image* $f = h * u = h * D\alpha$; this shows that *the coefficients α that encode the sharp image u relative to the basis $\{d_k\}$ are the same as the coefficients that encode the blurred image f relative to the blurred basis $\{b_k\} \doteq \{h * d_k\}$.*

But while we do not have access to the sharp image u , we do have access to the original basis elements $\{d_k\}$. Therefore, all we need to recover the encoding of the sharp image are the coefficients of the encoding of the blurred image relative to the blurred basis. The ensuing algorithm is as follows:

1. Take a dictionary $\{d_k\}$, either from a generic over-complete basis or learned from the image using any of a variety of sparse coding algorithms, for instance [1]. Convolve the basis with the kernel h to obtain a “blurred basis” $\{b_k\} \doteq \{h * d_k\}$.
2. Perform sparse coding of the blurred image f relative to the blurred basis $\{b_k\}$ to obtain $\hat{\alpha}$ with $\|\hat{\alpha}\|_0$ small such that $\|f - \sum_{k=1}^K b_k \hat{\alpha}_k\|$ is also small.
3. Reconstruct the original (deblurred) image *directly* via $\hat{u} = \sum_{k=1}^K d_k \hat{\alpha}_k$.

Note that this algorithm performs deblurring without solving a backward diffusion or other numerically ill-conditioned problem. Instead, it solves the inverse problem by “direct methods,” an approach sometimes referred to as “analysis by synthesis” [9], whereby an explicit generative model is used to match the statistics of the measured data, and the model itself provides both the necessary regularization and the solution to the desired inverse problem.

Now, of course the devil is in the details, as real images are not defined on the entire real plane. If we wish to keep the complexity of the coding step manageable we will have to break down the image into patches, which raises the issue of boundary effects and scale. We now derive the algorithm for a partition of the image, explicitly taking these issues into account.

1.2. State of the art

The concept of sparsity has recently re-gained popularity in image processing and computer vision; similarly, there is a sizeable literature on the problems of deblurring, deconvolution and shape from defocus or motion blur. The two streams, however, have never crossed, at least to the best of our knowledge.

² $\forall \epsilon \exists K = K(\epsilon)$ such that if $M > K$, then $\|f - \sum_{k=1}^M d_k \alpha_k\| < \epsilon$.

The literature on “exemplars” in computer vision, although not explicitly about sparsity, goes in this direction, including the early work of non-local super-resolution [8, 4, 15]. The idea of sparse representation of local patches is widely used in the many applications of image processing, such as denoising [5], color denoising and inpainting [12] and super-resolution [18, 3]. Deconvolution, on the other hand, is mainly solved by regularization. For example, the Wiener filter [2] uses the H^1 semi-norm of the solution, which favors smooth reconstructions. Total Variation (TV) [16], as already mentioned, favors piece-wise constant solutions, whereas wavelet-based deconvolution does not explicitly enforce a regularizer, but exercises regularization through complexity bounds [14, 7].

There is, however, relatively little work on the use of sparse priors for deblurring, which is the goal in this paper. Although super-resolution is a close relative to deblurring (the point-spread function corresponds to block-averaging of neighboring pixels), the latter has not been addressed directly in a sparse setting.

Since we use the K-SVD algorithm [5] as a building block, we will briefly review it here to make the manuscript self-contained. We will then extend it to the continuum and apply it to deblurring in Sect. 2. More details on the implementation are presented in Sect. 2.1, and Sect. 3 contains numerical experiments.

1.3. Sparse representations for denoising

If we consider discrete image patches, i.e. positive-valued matrices of size $\sqrt{n} \times \sqrt{n}$ pixels, ordered lexicographically as column vectors $x \in \mathbb{R}^n$, then the sparsity assumption corresponds to assuming the existence of a matrix $D \in \mathbb{R}^{n \times K}$, the “dictionary,” such that every image patch x can be represented as a linear combination of its columns with a vector of coefficients with small \mathcal{L}^0 norm. If we measure y , a version of x corrupted by additive Gaussian noise that is spatially white (independent and identically distributed) with standard deviation σ , then the maximum a-posteriori (MAP) estimator of the “denoised” patch x is given by $D\hat{\alpha}$, where

$$\hat{\alpha} = \arg \min_{\alpha} \|\alpha\|_0 \quad \text{s.t.} \quad \|D\alpha - y\|_2^2 \leq T, \quad (1)$$

where T is dictated by σ . If one wishes to encode a larger image X of size $\sqrt{N} \times \sqrt{N}$ ($N \gg n$), with a combination of columns of the low-dimensional dictionary D , a natural approach is to use a block-coordinator relaxation.

$$\hat{X} = \arg \min_{X, \alpha_{ij}, D} \|X - Y\|_2^2 + \lambda \sum_{i,j} \|\alpha_{ij}\|_0 + \mu \sum_{i,j} \|D\alpha_{ij} - R_{ij}X\|_2^2. \quad (2)$$

The first term measures the fidelity between the measured image Y and its denoised (and unknown) version X . The

second term enforces sparsity of each patch; the $n \times N$ matrix R_{ij} extracts the (i, j) th block from the image. A simple denoising algorithm based on sparse coding goes as follows,

1. Initialization: Set $X = Y, D =$ an overcomplete discrete cosine transform (DCT) dictionary.
2. Repeat until it converges:
 - *Sparse Coding*: fix X and D , compute the representation vectors α_{ij} for each patch $R_{ij}X$

$$\hat{\alpha}_{ij} = \arg \min_{\alpha} \|\alpha\|_0 \quad (3)$$

s.t. $\|D\alpha - R_{ij}X\|_2^2 \leq T$.

- *Dictionary Update*: fix X and $\{\alpha_{ij}\}$, compute D via K-SVD [1] one column at a time.

3. Set:

$$X = \frac{Y + \mu \sum_{ij} R_{ij}^T D \alpha_{ij}}{Id + \mu \sum_{ij} R_{ij}^T R_{ij}}, \quad (4)$$

which is a simple averaging of shifted patches.

One could also fix the dictionary and only perform sparse coding in the iteration. Alternatively, the dictionary can be learned from a large number of patches in natural images via K-SVD [1] so that it is tailored to the data.

2. Direct Sparse Deblurring

In this section we formalize the problem of direct sparse deblurring. We find the formalization to be clearer when written in the continuum, so one knows on what domain each function is calculated. The previous claim, that the blurred image can be sparsely represented in the blurred basis by the same coefficients that the “ideal” image would have on the original basis, will become clear.

Let $u : \Omega \subset \mathbb{R}^2 \rightarrow \mathbb{R}^+$; $x \mapsto u(x)$ be the “ideal image”, corresponding to X in the discrete model. The procedure of extracting a small patch from an epsilon ball centered at x ,

$$\{u(y) : y \in \mathcal{B}_\epsilon(x)\} = \{u(x+y) : y \in \mathcal{B}_\epsilon(0)\} \quad (5)$$

can be expressed by an indicator function χ_ϵ acting on the image u .

Let $d_k : \mathcal{B}_\epsilon(0) \subset \mathbb{R}^2 \rightarrow \mathbb{R}$, $k = 1, \dots, K$ be a given overcomplete basis of $\mathcal{L}_{loc}(\mathcal{B}_\epsilon(0) \rightarrow \mathbb{R})$, and $\alpha_k(x) \in \mathbb{R}^K$ be the k th sparse coefficient of the patch centered at x . Then the sparse representation of one image patch is given by

$$\begin{aligned} u_x(y) &= \chi_\epsilon(y)u(x+y) \\ &= \sum_{k=1}^K d_k(y)\alpha_k(x) \doteq \mathbf{d}(y)\boldsymbol{\alpha}(x), \end{aligned} \quad (6)$$




Figure 1. Diagram of the continuum formulation.

where $\mathbf{d}(y) = [d_1, d_2, \dots, d_K](y)$, $y \in \mathcal{B}_\epsilon(0)$ and $\boldsymbol{\alpha}(x) = [\alpha_1, \alpha_2, \dots, \alpha_K](x)$, $x \in \Omega$. Fig. 1 illustrates the dictionary elements and how they match to local patches.

We want to use this local sparsity to enforce a global reconstruction prior in the sense that $u(x)$ is the minimizer of the sparse representation error for all the local patches.

$$\hat{u}(x) = \min J(u), \quad (7)$$

where $J(u)$ is defined to be

$$\begin{aligned} &\int_{x \in \Omega} \int_{y \in \mathcal{B}_\epsilon(0)} \|\chi_\epsilon(y)u(x+y) - \mathbf{d}(y)\boldsymbol{\alpha}(x)\|^2 dy dx \\ &= \iint_{\Omega \times \Omega} \chi_\epsilon(y) \|u(x+y) - \mathbf{d}(y)\boldsymbol{\alpha}(x)\|^2 dy dx \\ &= \iint_{\Omega \times \Omega} \chi_\epsilon(z-x) \|u(z) - \mathbf{d}(z-x)\boldsymbol{\alpha}(x)\|^2 dz dx. \end{aligned} \quad (8)$$

To solve for $u(x)$, we compute its Euler-Lagrange equation

$$\begin{aligned} &\partial_u J(u)(z) \\ &= \int \chi_\epsilon^2(z-x) (\chi_\epsilon(z-x)u(z) - \mathbf{d}(z-x)\boldsymbol{\alpha}(x)) dy \\ &= \left[\int \chi_\epsilon(z-x) dx \right] u(z) - \int \chi_\epsilon(z-x) \mathbf{d}(z-x) \boldsymbol{\alpha}(x) dx. \end{aligned} \quad (9)$$

There is a closed-form solution for $u(x)$ w.r.t $\boldsymbol{\alpha}(x)$ that minimizes the objective function $J(u)$. It is obtained by setting the Euler-Lagrange equation to zero:

$$\hat{u}(x) = \frac{1}{\omega} \int \chi_\epsilon(x-y) \mathbf{d}(x-y) \boldsymbol{\alpha}(y) dy, \quad (10)$$

where $\omega = \int \chi_\epsilon(x-y) dy$ is the area of the ϵ -ball. Now, the measured image is, by assumption

$$\begin{aligned} f(x) &\doteq \int h(x-\bar{x})u(\bar{x})d\bar{x} + n(x) \\ &= \frac{1}{\omega} \int h(x-\bar{x}) \int \chi_\epsilon(\bar{x}-y) \mathbf{d}(\bar{x}-y) \boldsymbol{\alpha}(y) d\bar{x} dy + n(x) \\ &= \frac{1}{\omega} \iint [h(x-\bar{x})\chi_\epsilon(\bar{x}-y) \mathbf{d}(\bar{x}-y) d\bar{x}] \boldsymbol{\alpha}(y) dy + n(x), \end{aligned} \quad (11)$$

where h is a space-invariant blurring kernel and $n(x)$ is the additive noise, whose variance is σ^2 . The blurred basis is easily defined as

$$b_k(z) \doteq \int h(z - \bar{x})\chi_\epsilon(\bar{x})d_k(\bar{x})d\bar{x}. \quad (12)$$

The characteristic function χ_ϵ implies that the boundary condition for the convolution is zero-padding. Denote with r the support of the blurred basis, in particular $r = \epsilon + \text{supp}(h)$. Therefore the measured image is a sparse representation under this blurred basis:

$$f(x) = \frac{1}{\omega} \int \chi_r(x - y)\mathbf{b}(x - y)\boldsymbol{\alpha}(y)dy + n(x). \quad (13)$$

We solve for the sparse coefficients in the following,

$$\hat{\boldsymbol{\alpha}}(x) = \arg \min \int \|\boldsymbol{\alpha}(x)\|_0 dx, \quad (14)$$

$$\text{s.t.} \int \|f(x) - \omega^{-1} \int \chi_r(x - y)\mathbf{b}(x - y)\boldsymbol{\alpha}dy\|^2 dx \leq T.$$

This optimization problem (a) is finite-dimensional (the only unknown is $\boldsymbol{\alpha}$) and (b) *does not involve de-blurring*. All that is required is to find the finite-dimensional sparse set of coefficient that best approximates the given image. Note that this is accomplished by *blurring the base*. In other words, one is only required to solve a *direct* problem, rather than the inverse problem of deblurring. Once the coefficients $\hat{\boldsymbol{\alpha}}$ are obtained, we can compute the “deblurred” image \hat{u} via eq. (10). Note that the deblurred image is a sparse combination of the (original, non-blurred) basis, and therefore – by construction – one should expect the reconstruction to exhibit the same spatial frequencies of the original (unblurred) data from which the overcomplete basis has been learned.

2.1. Boundary issues

In practice, solving for $\boldsymbol{\alpha}$ from (14) is not an easy task, since $\boldsymbol{\alpha}$ at different y contribute to one value. Instead we minimize an upper bound.

$$\begin{aligned} & \int \left\| \int \chi_\epsilon(x - y)f(x)dy - \int \chi_r(x - y)\mathbf{b}(x - y)\boldsymbol{\alpha}(y)dy \right\|^2 dx \\ & \leq \iint \chi_\epsilon(x - y)\|f(x) - \mathbf{b}(x - y)\boldsymbol{\alpha}(y)\|^2 dx dy. \end{aligned} \quad (15)$$

The above equation suggests coding the blurry patch centered at y in terms of the blurred basis $\{b_i\}$. Furthermore, the characteristic function indicates that the blurry patch has to be zero-padded in order to be consistent with the dimension of the blurred basis. However, these two terms can not match due to the boundary issue of convolution, as shown in Fig. 2. Instead we match them in the region where the




Figure 2. Left: the blurry patch with zero-padding. Right: the blur basis. The red square indicates the region for our refined dictionary.

convolution is computed without the zero-padded edges. In particular, we refine our blurred basis to be

$$\tilde{b}_i(z) = \begin{cases} b_i(z) & |z| < \epsilon - \text{supp}(h) \doteq \epsilon_0 \\ 0 & \text{otherwise.} \end{cases} \quad (16)$$

Experimentally we find that it is better to tailor the clear basis to have the same domain size as the blurred one. Therefore we have a two step algorithm for sparse deblurring

1. Solve the coefficients from the measured image

$$\hat{\boldsymbol{\alpha}}(x) = \arg \min \|\boldsymbol{\alpha}\|_0, \quad (17)$$

$$\text{s.t.} \iint \chi_{\epsilon_0}(x - y)\|f(x) - \mathbf{b}(x - y)\boldsymbol{\alpha}(y)\|^2 dx dy \leq T.$$

2. Stitch all the patches by averaging

$$\hat{u}(x) = \frac{\int \chi_{\epsilon_0}(x - y)\mathbf{d}(x - y)\hat{\boldsymbol{\alpha}}dy}{\int \chi_{\epsilon_0}(x - y)dy}. \quad (18)$$

2.2. Weighted Averaging

In the case where an image patch corresponds to a single entry in the dictionary, we could use weighted averaging to reduce the sparse coding artifacts. Consider, for example, the problem of deblurring text: We apply the same sparse coding step to get the coefficients $\boldsymbol{\alpha}(y)$, then use weighted averaging to combine the result

$$\hat{u}(x) = \frac{\int_{y \in \mathcal{B}_\epsilon(x)} \mathbf{d}(x - y)\boldsymbol{\alpha}(y)w(y)dy}{\int_{y \in \mathcal{B}_\epsilon(x)} w(y)dy}. \quad (19)$$

The weight is chosen to penalize the large \mathcal{L}^0 norm of the coefficients, for example,

$$w(y) = \exp \left\{ \frac{1 - |\boldsymbol{\alpha}(y)|_0}{s} \right\}, \quad (20)$$

where s is a control parameter.

3. Experiments

In this section we compare our algorithm to more traditional methods, such as ROF [16] and the wavelet-based approach called ForWaRD [14]. The optimal method parameters for both ROF and ForWaRD are chosen from a series of values with wide range. The parameters in our algorithm are determined by the data: the size of the dictionary is proportional to the width of the blurring kernel and the stopping criterion for the sparse coding stage is when the residual is below the variance of the noise.

We use root-mean-square (RMS) as a means of judging performance, $RMS(u, I) = \sqrt{\int_{x \in \Omega} (u(x) - I(x))^2 dx}$, where $I(x)$ is the original image and $u(x)$ is the recovery.

3.1. Binary Text images

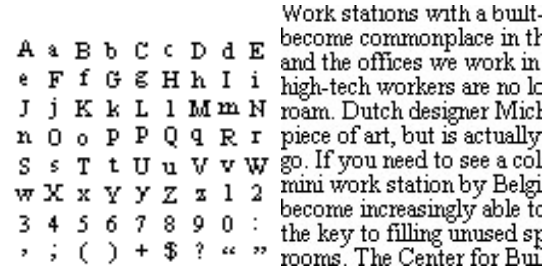
We synthesize a template with all the alphanumeric characters and common punctuation as well as 5 text images from different categories of CNN news. The dictionary is comprised of image patches randomly sampled from the three images and the template, all shown in Fig. 3. The template contains all individual characters, while the training images serve to represent meaningful pairs. We test the deblurring on the other two text images. The data are corrupted by 5×5 Gaussian kernel with $\sigma = 1$ and additive noise whose standard deviation is 5. For direct sparse deblurring, the visual quality is for the most part satisfactory except for the smoothing effects around some letters. This is mostly attributed to the limitations of the dictionary.

We also measure the effect of the number of the elements in the dictionary on the deblurring performance. Fig. 5 shows the results averaged from ten different experiments of randomly sampled elements in the dictionary. In general, increasing the number of elements in the dictionary improves the results, but with a diminishing return.

3.2. General case

As in the case of super-resolution [18], the dictionary consists of random samples from the training images, which have statistics similar to the test one. Here we consider three images: "Rose," "Koala," and "Castle." The training images are taken from the image datasets of flowers, animals and architecture respectively. Fig. 6 shows several examples in each training set. Flower images are from the Internet, while the other images are from the Berkeley Segmentation Dataset [13]. For each category we randomly sample 20,000 patches of size 16×16 to form each dictionary.

The input data are corrupted by convolving with a 9×9 Gaussian kernel of $\sigma = 1$ with additive noise whose standard deviation is 5. As shown in Fig. 7, ROF returns piecewise constant images, while ForWaRD produces noticeable artifacts in the reconstruction.



Work stations with a built-in desk become commonplace in the offices we work in and high-tech workers are no longer roam. Dutch designer Michiel van der Kluis designed a piece of art, but is actually a functional work station by Belgian designer. The Center for Business Architecture has become increasingly able to fill unused spaces with the key to filling unused spaces. The Center for Business Architecture has become increasingly able to fill unused spaces with the key to filling unused spaces. The Center for Business Architecture has become increasingly able to fill unused spaces with the key to filling unused spaces.

Figure 3. The training data. The dictionary is obtained by randomly sampling raw 10×10 patches from the text images as well as the template shown on the top left. All the text images are from different categories of CNN news.

We also conduct an experiment of texture deblurring, in which both training and test images have the same structure. We cut a texture image into half, one as training and the other as testing. We construct the dictionary to be 20,000 random samples from the training image. We blur the test image with Gaussian kernel of $\sigma = 2$ plus additive noise. Fig. 8 shows that our method is a significant improvement over the traditional ones.

Finally, the quantitative measurement, RMS, of each method for all the experiments is listed in Table. 1.

Image	ROF	ForWaRD	Our method
Text 1	65.55	63.15	15.13
Text 2	64.03	62.72	14.94
Rose	5.47	4.95	4.74
Koala	9.80	9.05	8.69
Castle	14.41	12.93	11.99
Texture	14.52	14.44	8.84

Table 1. RMS errors for different methods.

4. Discussion

We have presented a simple deblurring/deconvolution algorithm that exploits the assumption of sparse representation for natural image statistics. It is a "direct" algorithm, in the sense that it performs inference by synthesis relative to an explicit generative model (which acts as a regularizer), without the need to solve an ill-conditioned inverse prob-

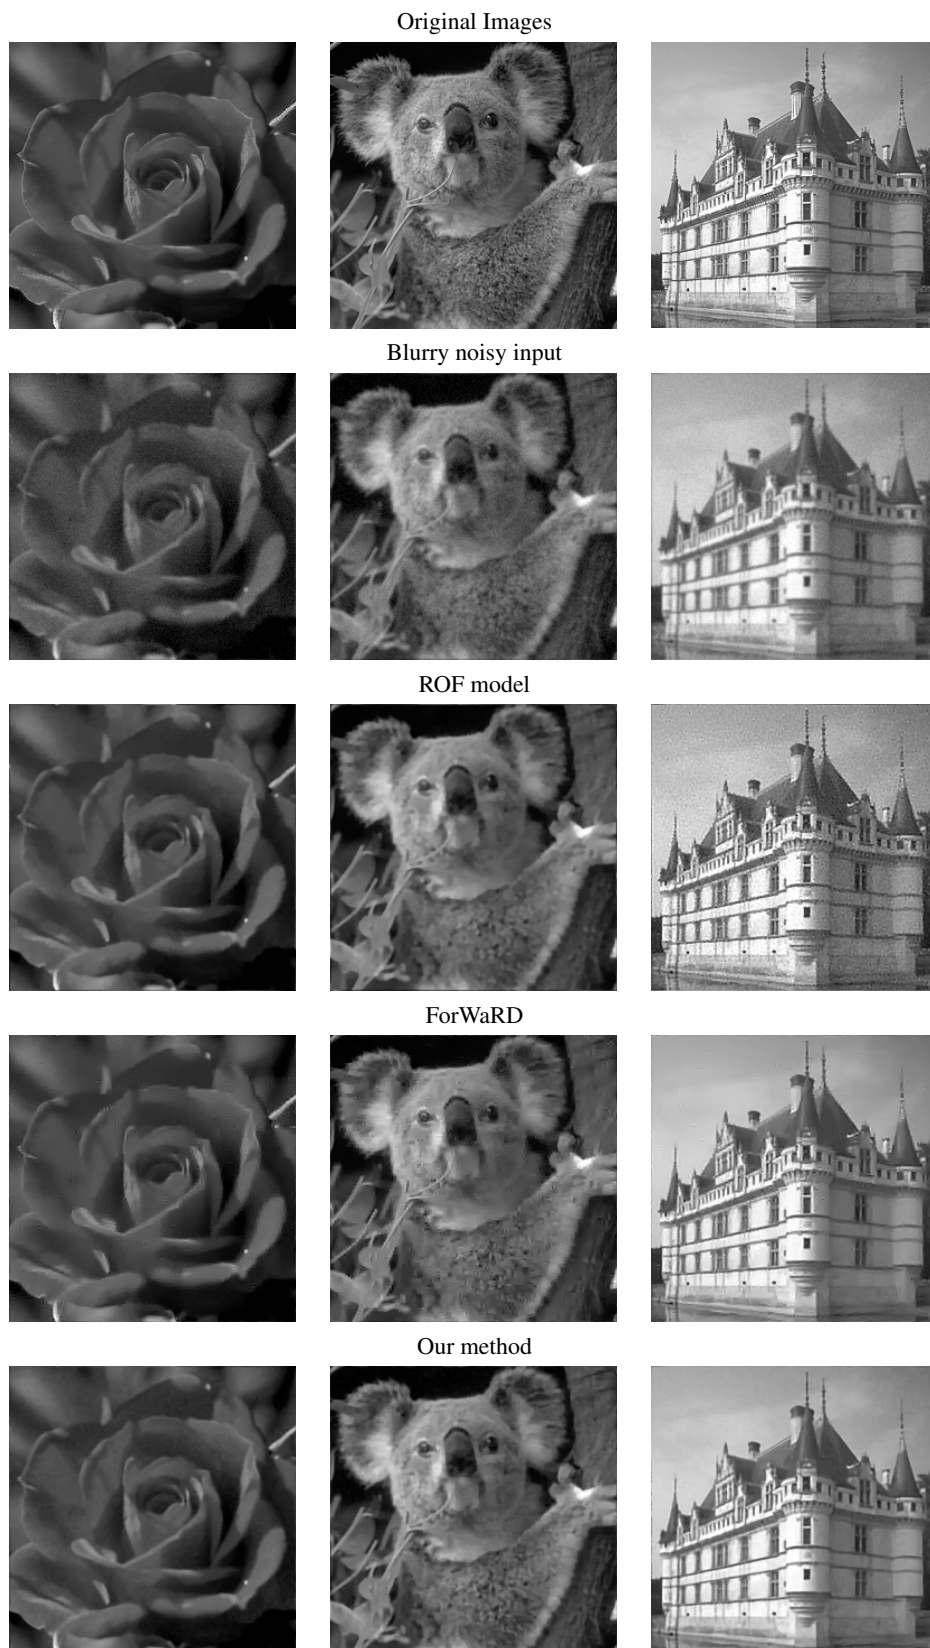


Figure 7. Grayscale image deblurring with 20,000 dictionary elements.

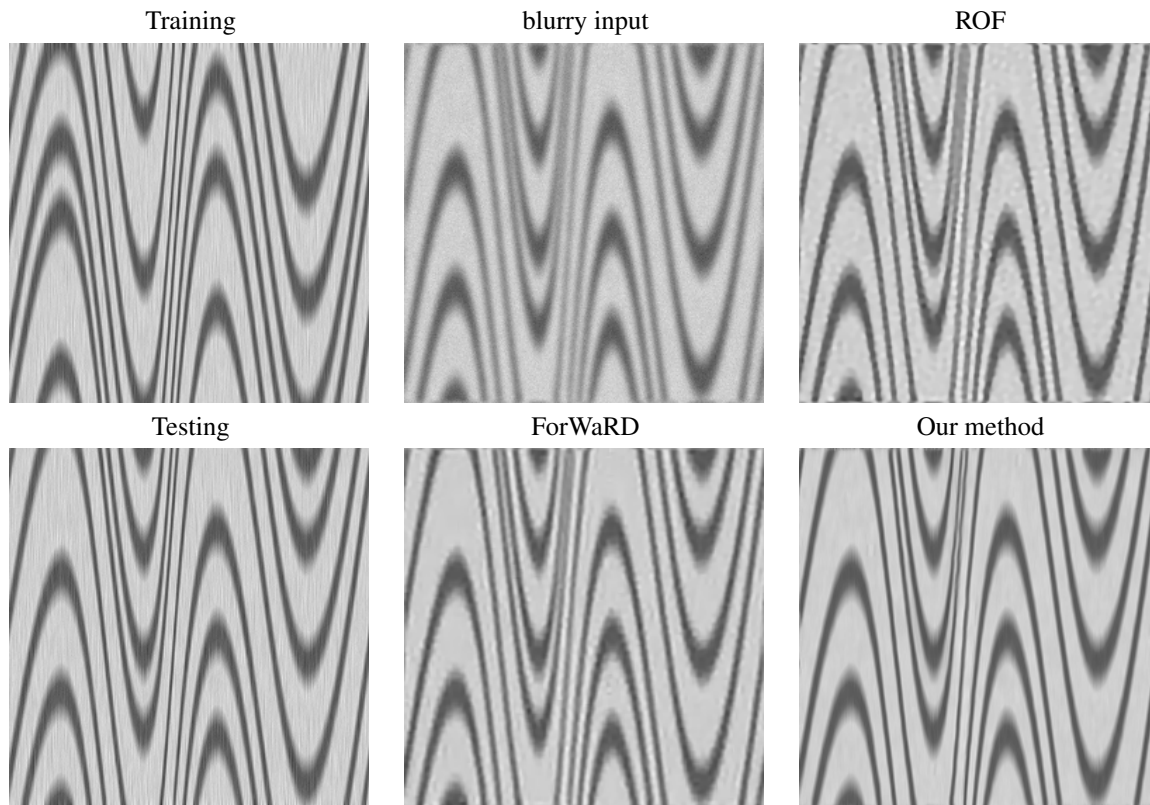


Figure 8. The dictionary is comprised of random samples from the training image, which has the same structure as the testing one.

lem. The algorithm is competitive with the state of the art.

References

- [1] M. Aharon, M. Elad, and A. Bruckstein. K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation. *IEEE Transactions on Signal Processing*, 54(11):4311–4322, 2006. 2, 3
- [2] H. C. Andrews and B. R. Hunt. *Digital Image restoration*. Prentice-Hall, Englewood Cliffs, NJ, 1977. 2
- [3] D. Datsenko and M. Elad. Example-based single document image super-resolution: a global map approach with outlier rejection. In *Multidim System Signal Processing*, number 18, pages 103–121, 2007. 2
- [4] A. A. Efros and T. K. Leung. Texture synthesis by non-parametric sampling. In *ICCV*, volume 2, pages 1033–1038, 1999. 2
- [5] M. Elad and M. Aharon. Image denoising via sparse and redundant representations over learned dictionaries. *IEEE Transactions on Image Processing*, 15(12):3736–3745, 2006. 2
- [6] P. Favaro, S. Soatto, L. Vese, and S. J. Osher. Shape from anisotropic diffusion. In *Intl. Conf. on Comp. Vis. and Patt. Recog.*, pages I–179–186, June 2003. 1
- [7] A. Foi, K. Dabov, V. Katkovnik, and K. Egiazarian. Shape-adaptive DCT for denoising and image reconstruction. *Proc. of SPIE Electronic Imaging, Image Processing: Algorithms and Systems V*, 6064, 6064N, 2006. 2
- [8] W. T. Freeman, T. R. Jones, and E. C. Pasztor. Example-based super-resolution. *IEEE Computer Graphics and Applications*, 2002. 2
- [9] U. Grenander. *General Pattern Theory: A Mathematical Study of Regular Structures*. Cambridge University Press, England, 1993. 2
- [10] J. Huang and D. Mumford. Statistics of natural images and models. In *Intl. Conf. on Comp. Vis. and Patt. Recog.*, pages 541–547, 1999. 1
- [11] M. S. Lewicki and B. A. Olshausen. Inferring sparse, overcomplete image codes using an efficient coding framework. In *Advances in Neural Information Processing 10 (Proc. NIPS '97)*, pages 815–821, 1998. 1
- [12] J. Mairal, M. Elad, and G. Sapiro. Learning multiscale sparse representations for image and video restoration. *IEEE Transactions on Image Processing*, 17(1):53–69, 2008. 2
- [13] 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 *Int'l Conf. Computer Vision*, volume 2, pages 416–423, July 2001. 5
- [14] R. Neelamani, H. Choi, and R. G. Baraniuk. ForWaRD: Fourier-wavelet regularized deconvolution for ill-conditioned systems. *IEEE Trans. Signal Processing*, 52(2):418–433, Feb. 2004. 2, 5

Original Images
 ation, not the automakers. economy -- not paying to
 ching movies on laptops ir Officials have yet to say w
 hacking their car stereos tc Jintao will propose at the
 ng before they could get a leaders of 20 major econo
 nade LCD screen in a head made the outlines of its str
 ock for their favorite digita announcement of a multibi
 but based on the miles of a to stimulate its economy w
 t this year's Consumer Ele on construction, tax cuts a
 at's about to change, and fi with no mention of efforts
 out the usual CES assortr Prime Minister Gordon Br
 speakers. neon-lit amns an to use its nearly \$2 trillion

Blurry noisy input
 ation, not the automakers. economy -- not paying to
 ching movies on laptops ir Officials have yet to say w
 hacking their car stereos tc Jintao will propose at the
 ng before they could get a leaders of 20 major econo
 nade LCD screen in a head made the outlines of its str
 ock for their favorite digita announcement of a multibi
 but based on the miles of a to stimulate its economy w
 t this year's Consumer Ele on construction, tax cuts a
 at's about to change, and fi with no mention of efforts
 out the usual CES assortr Prime Minister Gordon Br
 speakers. neon-lit amns an to use its nearly \$2 trillion

ROF model
 ation, not the automakers. economy -- not paying to
 ching movies on laptops ir Officials have yet to say w
 hacking their car stereos tc Jintao will propose at the
 ng before they could get a leaders of 20 major econo
 nade LCD screen in a head made the outlines of its str
 ock for their favorite digita announcement of a multibi
 but based on the miles of a to stimulate its economy w
 t this year's Consumer Ele on construction, tax cuts a
 at's about to change, and fi with no mention of efforts
 out the usual CES assortr Prime Minister Gordon Br
 speakers. neon-lit amns an to use its nearly \$2 trillion

ForWaRD
 ation, not the automakers. economy -- not paying to
 ching movies on laptops ir Officials have yet to say w
 hacking their car stereos tc Jintao will propose at the
 ng before they could get a leaders of 20 major econo
 nade LCD screen in a head made the outlines of its str
 ock for their favorite digita announcement of a multibi
 but based on the miles of a to stimulate its economy w
 t this year's Consumer Ele on construction, tax cuts a
 at's about to change, and fi with no mention of efforts
 out the usual CES assortr Prime Minister Gordon Br
 speakers. neon-lit amns an to use its nearly \$2 trillion

Our method
 ation, not the automakers. economy -- not paying to
 ching movies on laptops ir Officials have yet to say w
 hacking their car stereos tc Jintao will propose at the
 ng before they could get a leaders of 20 major econo
 nade LCD screen in a head made the outlines of its str
 ock for their favorite digita announcement of a multibi
 but based on the miles of a to stimulate its economy w
 t this year's Consumer Ele on construction, tax cuts a
 at's about to change, and fi with no mention of efforts
 out the usual CES assortr Prime Minister Gordon Br
 speakers. neon-lit amns an to use its nearly \$2 trillion

Figure 4. Text Deblurring with 20,000 dictionary elements.

[15] M. Protter, M. Elad, H. Takeda, and P. Milanfar. Generalizing the non-local-means to super-resolution reconstruction. *IEEE Transactions on Image Processing*, to appear, 2008. 2

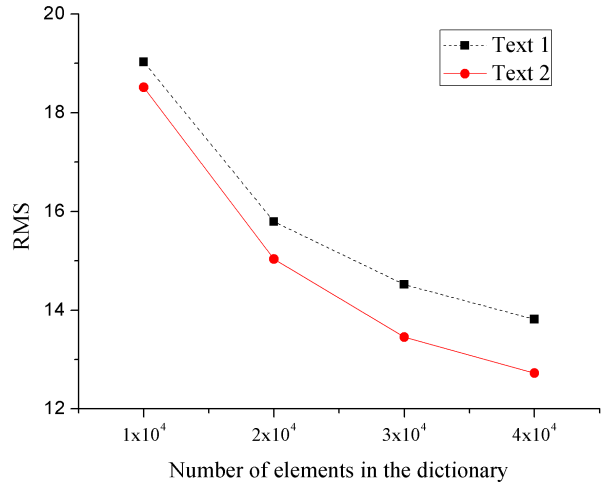


Figure 5. Influence of the number of the elements in the dictionary on the deblurring performance. The average of 10 experiments for each column is reported.



Figure 6. Examples in the training images.

[16] L. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D*, 60:259–268, 1992. 2, 5

[17] A. Tichonov and V. Arsenin. *Solution of Ill-Posed Problems*. New York: Wiley, 1977. 1

[18] J. Yang, J. Wright, T. Huang, and Y. Ma. Image super-resolution as sparse representation of raw image patches. In *Intl. Conf. on Comp. Vis. and Patt. Recog.*, 2008. 2, 5