Fast Deconvolution with Color Constraints on Gradients

Ayan Chakrabarti
and
Todd Zickler

TR-06-12

Fast Deconvolution with Color Constraints on Gradients

Ayan Chakrabarti and Todd Zickler

Harvard University, Cambridge, MA, USA
ayanc@eecs.harvard.edu, zickler@seas.harvard.edu

◆

Abstract—In this report, we describe a fast deconvolution approach for color images that combines a sparse regularization cost on the magnitudes of gradients with constraints on their direction in color space. We form these color constraints in a way that allows retaining the computationally-efficient optimization strategy introduced in recent deconvolution methods based on *half-quadratic splitting*. The proposed algorithm is capable of handling a different blur kernel in each color channel, and is used for per-layer deconvolution in our paper: "Depth and Deblurring from a Spectrally-varying Depth-of-Field" [1]. A MATLAB implementation of this method is available at **http://vision.seas.harvard.edu/ccap**, and takes roughly 20 seconds to deconvolve a three-channel 1544×1028 color image, on a Linux-based Intel I-3 2.1GHz machine.

1 INTRODUCTION

Non-blind deconvolution refers to the recovery of a sharp image $x(n)$ from a blurred and noisy observation $y(n)$:

$$y(n) = (x * k)(n) + \epsilon(n), \qquad (1)$$

where the blur kernel k is known, and $\epsilon(n)$ is typically assumed to be white Gaussian noise. A regularized estimation approach recovers the sharp image $x(n)$ as

$$x(n) = \arg\min_x \frac{\mu}{2} \sum_n [y(n) - (x * k)(n)]^2 + \Phi(x), \qquad (2)$$

where the first term measures fidelity to the observation, $\Phi(x)$ is a regularization cost based on statistical properties of sharp natural images, and μ is the relative weight between the two.

A common choice for $\Phi(x)$ is a *smoothness* cost that discourages discontinuities in the sharp image x by penalizing gradient magnitudes [2], [3], [4] as

$$\Phi(x) = \sum_{\nabla \in \mathcal{G}} \sum_n |x_\nabla(n)|^\alpha, \qquad (3)$$

where $\alpha \geq 0$, $\mathcal{G}$ is a set of gradient filters, and $x_\nabla(n) = (\nabla * x)(n)$. For $\alpha = 2$, the solution for (2) has a closed form that can be computed efficiently in the Fourier domain. However, such a squared cost prefers to distribute gradients equally over the recovered image [2], and based on the choice of μ, yields images that are either over-smoothed, or have ringing artifacts.

It is generally understood that a choice of $\alpha \in (0, 1]$ corresponds better to the distribution of gradients in natural images, and yields solutions with sharper edges where gradients are concentrated at a sparse set of pixels. Unfortunately, these values of α do not admit closed-form solutions and require an iterative approach for deconvolution. Levin et al. [2] propose such an approach based on iterative re-weighted least squares (IRLS), where at each iteration the regularizer is approximated with a weighted squared cost with weights based on the current estimate of x. Unfortunately with spatially-varying weights, deconvolution with the approximated squared cost

can not be carried out in the Fourier domain, and the solution at each iteration must in turn be computed iteratively using the conjugate-gradient method.

Recent work [3], [4] demonstrates that one can solve (2) for $\alpha < 2$ with greater efficiency using an optimization approach based on half-quadratic splitting (which we describe in detail in Sec. 2.1). These algorithms involve iterating between per-pixel *shrinkage* operations on gradients, and deconvolution in the Fourier domain. A comparison in [4] shows that this approach is about 400 times faster than IRLS on large "megapixel" images (i.e., those with a resolution greater than 1024×1024). In this work, we extend the half-quadratic splitting approach to include color constraints on gradients in the sharp image x in a way that yields improved results, with roughly the same computational cost.

1.1 Color-based Regularization

The regularization schemes discussed above are defined exclusively in terms of single-channel image statistics, and ignore color information. Recently, Joshi et al. [5] demonstrated that higher quality results can be obtained by using color statistics during deconvolution, even for the case when the blur kernel is spectrally uniform. Color models are even more beneficial in the context of deconvolution in our work in [1], where each color channel is blurred with a different kernel. In this case, observed artifacts from independent deconvolution of color channels correspond to different spatial frequencies in the different channels, and can lead to the creation of spurious chromaticities in the estimated image. Conversely, by incorporating color statistics, one can exploit the fact the observed image contain complementary spatial frequency information in the different channels.

The color-based regularization introduced in [5] is defined in terms of pixel-domain statistics in local neighborhoods, and the corresponding estimation algorithm is developed in the IRLS framework. In this work, we aim to incorporate color statistics in a way that admits use of the fast deconvolution framework of [3], [4]. We reason that with a high weight on the regularizer during estimation, we recover an image X_s that is excessively smoothed but has fewer artifacts due to noise, and can be considered to be a blurred version of the true sharp image X. We use a squared-cost with a low value of μ to estimate this over-regularized estimate X_s in closed-form.

Then, we draw on intuition from the spatio-spectral model in [1] that gradients in local neighborhoods have the same color direction, i.e., they lie on the same line in color space. This means that gradients in X_s, being local averages of the true gradients, are likely to lie on the same color line as those in X, albeit with attenuated magnitudes. Our strategy

is to use the gradient colors from X_s as additional constraints while estimating X with a traditional sparse regularization cost on gradient magnitudes. We find that this simple approach achieves performance comparable to that of the color IRLS-based method in [5], at a fraction of the computational cost.

2 PROPOSED METHOD

We consider the general case where each color channel $y^{\{i\}}(n)$ of $Y(n), i \in \{R, G, B\}$ is blurred with a different kernel $k^{\{i\}}$. We first generate an over-regularized estimate X_s of X with a squared cost and $\mu_s = 0.5$ as:

$$x_s^{\{i\}} = \arg\min_x \frac{\mu_s}{2} \sum_n \left[y^{\{i\}}(n) - (x * k^{\{i\}})(n) \right]^2 + \sum_{\nabla \in \mathcal{G}} \sum_n |x_\nabla(n)|^2, \tag{4}$$

where the set $\mathcal{G}$ contains two gradient filters corresponding to vertical and horizontal finite-differences. This can be solved directly in the Fourier domain as:

$$x_s^{\{i\}} = \mathcal{F}^{-1} \left[\frac{\mu_s \mathcal{F}[k^{\{i\}}]^* \mathcal{F}[y^{\{i\}}]}{\sum_\nabla |\mathcal{F}[\nabla]|^2 + \mu_s |\mathcal{F}[k^{\{i\}}]|^2} \right], \tag{5}$$

where $\mathcal{F}$ and $\mathcal{F}^{-1}$ refer to the 2D Fourier transform and its inverse. Note that this Fourier-domain formulation is based on a periodic-extension assumption on the image. To prevent ringing from sharp discontinuities between pixels at opposite boundaries, we extend the observed image $Y(n)$ by padding it with values that smoothly blend these discontinuities. These extended regions are discarded after deconvolution to yield a result with the original resolution.

$X_s(n)$ serves as the reference for gradient colors in the final estimate X. For notational convenience, we define $X_\nabla(n)$ as a color-gradient vector of $X(n)$ as:

$$X_\nabla(n) = \left[x_\nabla^{\{R\}}(n), x_\nabla^{\{G\}}(n), x_\nabla^{\{B\}}(n) \right], \tag{6}$$

and the gradient color unit-vectors $\hat{V}_\nabla(n)$ from X_s as

$$\hat{V}_\nabla(n) = X_{s\nabla}(n) / \|X_{s\nabla}(n)\|. \tag{7}$$

The final sharp image is then recovered as

$$X = \arg\min_X \frac{\mu}{2} \sum_{n,i} \left[y^{\{i\}}(n) - (x^{\{i\}} * k^{\{i\}})(n) \right]^2 + \sum_{\nabla \in \mathcal{G}} \sum_n |\langle X_\nabla(n), \hat{V}_\nabla(n) \rangle|^\alpha, \tag{8}$$

with the constraint that $X_\nabla(n) \propto \hat{V}_\nabla(n), \ \forall n$. We use $\alpha = 2/3$, and set μ based on the noise level in the observed image ($\mu = 10^4$ for all results reported in [1]).

2.1 Half-quadratic Splitting

To solve the optimization problem in (8), we closely follow the approach in [4] and introduce a new cost-function:

$$C(X(n), W_\nabla(n), \beta) = \frac{\mu}{2} \sum_{n,i} \left[y^{\{i\}}(n) - (x^{\{i\}} * k^{\{i\}})(n) \right]^2 + \frac{\beta}{2} \sum_{\nabla,n} \|W_\nabla(n) - X_\nabla(n)\|^2 + \sum_{\nabla,n} |\langle W_\nabla(n), \hat{V}_\nabla(n) \rangle|^\alpha, \tag{9}$$
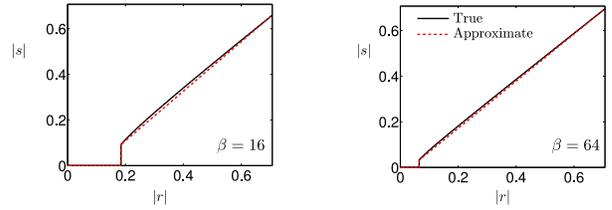

Fig. 1. Gradient magnitude "shrinkage" function for $\alpha = 2/3$ and different values of β. We show the true function, as well as the thresholded-linear approximation used in our implementation.

where $W_\nabla(n)$ are auxiliary variables constrained to lie along $\hat{V}_\nabla(n)$. The solution to (8) is then derived using an iterative algorithm that alternately updates X and $W_\nabla(n)$ to minimize $C(\cdot)$, starting with X_s as the initial estimate for X. The parameter β is increased after each round of updates to W_∇ and X, by a factor of 4 in our implementation. We ensure that the iterations end with $\beta = 100\mu$, and find 8 iterations to be sufficient for convergence.

Updating X: Given the current estimates of $W_\nabla(n)$, X is updated to minimize the first two terms of (9) as

$$x^{\{i\}} = \mathcal{F}^{-1} \left[\frac{\sum_\nabla \mathcal{F}[\nabla]^* \mathcal{F}[w_\nabla^{\{i\}}[n]] + \mu/\beta \mathcal{F}[k^{\{i\}}]^* \mathcal{F}[y^{\{i\}}]}{\sum_\nabla |\mathcal{F}[\nabla]|^2 + \mu/\beta |\mathcal{F}[k^{\{i\}}]|^2} \right]. \tag{10}$$

Note that this update requires only one forward and one inverse Fourier transform per channel at each iteration, since the second term of the numerator needs to be only computed once for an input image, and the two terms of the denominator can be pre-computed with knowledge of the kernels and the resolution of $Y(n)$.

Updating W_∇: Correspondingly, each of the auxiliary variables $W_\nabla(n)$ can be independently updated based on the current estimate of $X(n)$, to minimize the latter two terms of (9) and satisfy the color-direction constraint as:

$$\text{(Project)} \quad r = \langle X_\nabla(n), \hat{V}_\nabla(n) \rangle, \tag{11}$$

$$\text{(Shrink)} \quad s = \arg\min_s |s|^\alpha + \frac{\beta}{2} |s - r|^2, \tag{12}$$

$$W_\nabla(n) = s \, \hat{V}_\nabla(n). \tag{13}$$

Krishnan and Fergus [4] describe an analytical solution for (12) when $\alpha = 2/3$, as well as a faster look-up table-based approach. As illustrated in Fig. 1, we find that the relationship between $|s|$ and $|r|$ can be well approximated by a thresholded-linear function (for each value of β). We use this approximation in our implementation since it is faster to compute, and find that it gives near identical results as the analytical solution.

3 COMPARISONS

The proposed algorithm is used to generate all the deblurred results in [1]. In Fig. 2, we show a close-up of one of these results as well as a comparison to the deblurred image that would have been recovered using the baseline greyscale half-splitting algorithm [4], with the same value of α. The latter essentially corresponds to deconvolving each color channel independently, and while the two approaches have nearly identical running times, the addition of color-gradient constraints removes spurious chromatic effects caused by ringing artifacts at different spatial frequencies in different color channels.
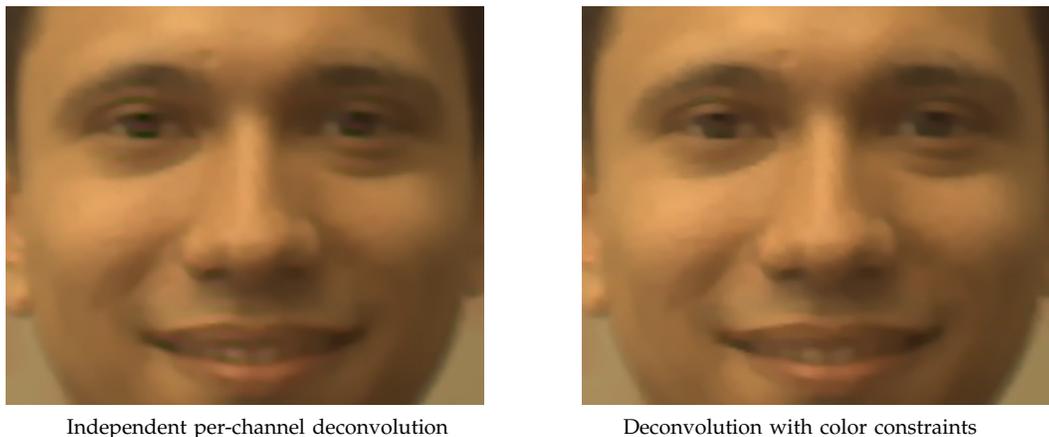
Independent per-channel deconvolution | Deconvolution with color constraints

Fig. 2. Comparison of deconvolution with and without color constraints, for a real captured image from [1]. We see that adding color constraints removes various spurious colors (near the eyes and mouth) that appear in the case of independent per-channel deconvolution.
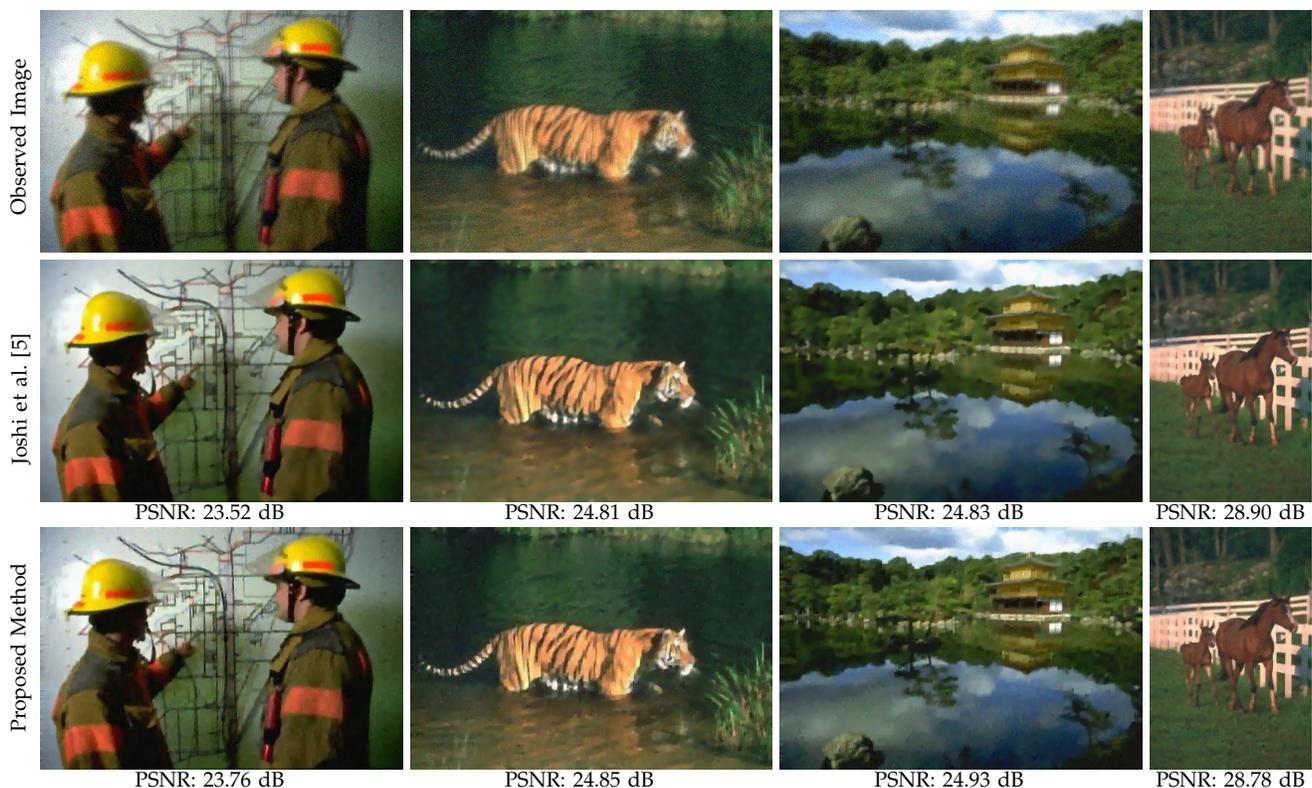

Fig. 3. Comparison to results reported in [5], for a spectrally-uniform kernel. Despite being significantly less expensive computationally, the proposed algorithm yields results with similar quality to the color ILRS-based method of Joshi et al. [5].

We next show results on images from the Berkeley segmentation database [6], that were synthetically blurred with a spectrally-uniform kernel and used for evaluation in [5]. Figure 3 compares the deconvolution results from our method to those reported for the color IRLS-based approach of [5], with PSNR values for both. The proposed method is found to have equivalent performance to [5], while offering a substantial advantage in terms of computational efficiency.

REFERENCES

[1] A. Chakrabarti and T. Zickler, "Depth and deblurring from a spectrally-varying depth-of-field," in *Proc. ECCV*, 2012.

[2] A. Levin, R. Fergus, F. Durand, and W. T. Freeman, "Image and depth from a conventional camera with a coded aperture," *ACM Transactions on Graphics (SIGGRAPH)*, 2007.

[3] Y. Wang, J. Yang, W. Yin, and Y. Zhang, "A new alternating minimization algorithm for total variation image reconstruction," *SIAM J. on Imag. Sci.*, 2008.

[4] D. Krishnan and R. Fergus, "Fast image deconvolution using Hyper-Laplacian priors," in *NIPS*, 2009.

[5] N. Joshi, C. Zitnick, R. Szeliski, and D. Kriegman, "Image deblurring and denoising using color priors," in *Proc. CVPR*, 2009.

[6] 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 *Proc. ICCV*, 2001.