

Image Deblurring and Denoising using Color Priors

Neel Joshi^{†*}

C. Lawrence Zitnick[†]

Richard Szeliski[†]

David J. Kriegman^{*}

[†]Microsoft Research

^{*}University of California, San Diego

Abstract

Image blur and noise are difficult to avoid in many situations and can often ruin a photograph. We present a novel image deconvolution algorithm that deblurs and denoises an image given a known shift-invariant blur kernel. Our algorithm uses local color statistics derived from the image as a constraint in a unified framework that can be used for deblurring, denoising, and upsampling. A pixel's color is required to be a linear combination of the two most prevalent colors within a neighborhood of the pixel. This two-color prior has two major benefits: it is tuned to the content of the particular image and it serves to decouple edge sharpness from edge strength. Our unified algorithm for deblurring and denoising out-performs previous methods that are specialized for these individual applications. We demonstrate this with both qualitative results and extensive quantitative comparisons that show that we can out-perform previous methods by approximately 1 to 3 DB.

1. Introduction

Two of the most common problems in photography are image blur and noise, which can be especially significant in light limited situations, resulting in a ruined photograph.

Image deconvolution in the presence of noise is an inherently ill-posed problem. The observed blurred image only provides a partial constraint on the solution—there exist many “sharp” images that when convolved with the blur kernel can match the observed blurred and noisy image. Image denoising presents a similar problem due to the ambiguity between the high-frequencies of the unobserved noise-free image and those of the noise. Thus, the central challenge in deconvolution and denoising is to develop methods to disambiguate solutions and bias the processes toward more likely results given some prior information.

In this paper, we propose using priors derived from image color statistics for deconvolution and denoising. We also perform up-sampling as a case of deblurring and denoising, using a minor extension, as discussed in Section 3. Our prior models an image as a per-pixel linear combination of two color layers, where the layer colors are expected to vary more slowly than the image itself. Edges and textures in an image are accounted for by blending the colors of the two layers. Our approach places priors on the values used for blending between the two colors, in essence creating a robust edge prior that is independent of gradient magnitude.

A central part of deconvolution is properly handling image noise. As a result, denoising can be considered a sub-problem of deblurring, and deconvolution methods can be used purely for denoising by considering the blurring kernel to be a delta function. Thus, we treat denoising as a sub-problem of deconvolution and present a non-blind deconvolution algorithm that can be used for both applications.

Our unified algorithm for deblurring and denoising out-performs previous methods that are specialized for these individual applications. We demonstrate this with both qualitative results and quantitative comparisons that show that we can out-perform previous methods by around 1 to 3 DB at high noise levels.

2. Related Work

Image deblurring and denoising have received a lot of attention in the computer graphics and vision communities. Basic approaches for denoising, such as Gaussian and median filtering, have a tendency to over-smooth edges and remove image detail. More sophisticated approaches use the properties of natural image statistics to enhance large intensity edges and suppress lower intensity edges. This property has been used by wavelet based methods [16], anisotropic diffusion [15], bilateral filtering [21], and Field of Experts models [19]. Our method shares some similarities with these approaches in that we consider natural image statistics in the form of a prior on the distribution of image gradients [12]; however, we go beyond this and additionally incorporate a prior derived from local color statistics. This allows us to avoid the over-smoothing that can occur with gradient-based methods.

The denoising aspect of our work is most similar to that of Liu *et al.* [14] in that we both use a local-linear color model. Where our methods differ is that we build a color model per pixel, while Liu *et al.* segment the image first and then build the model per segment. A further distinction between our work and previous work is that we address denoising in the larger context of image deconvolution, while most denoising work considers this problem in isolation.

To address image deblurring, some researchers have modified the image capture process [3] or used multiple images [2, 22] to aid in deblurring. Image blur due to limited resolution has led to the development of up-sampling algorithms [9, 7]. Determining the blur kernel from a single image, which is a critical sub-problem for deblurring nat-

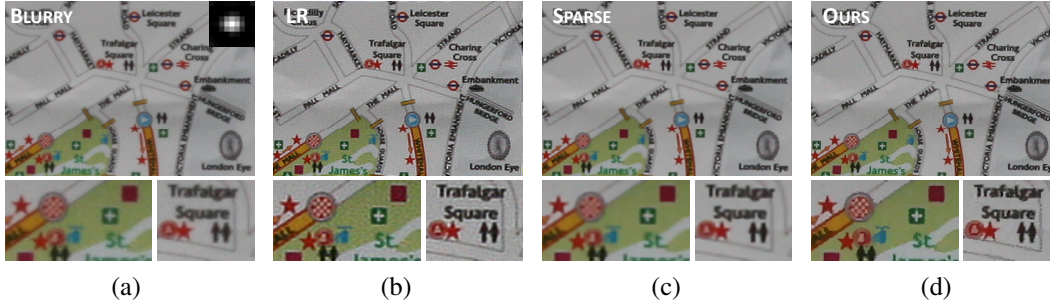


Figure 1: (a) Input blurry image (the 9×9 PSF is displayed (enlarged by $4 \times$) in the top right corner, the noise level is estimated per-pixel with the median $\sigma = 0.0137$), (b) image deblurred using the Lucy-Richardson algorithm, (c) using a sparse gradient prior, and (d) using our method ($\lambda_2 = 0.5$). Result (b) is sharp but noisy and has ringing artifacts; (c) is overly-smoothed; our result (d) has the sharpness of Lucy-Richardson (most notably in the red/white checkerboard) with significantly reduced noise and ringing.

ural images, has also been studied [8, 11]. One area that has received less attention, yet is critical for the above techniques, is that of non-blind deconvolution.

Non-blind image deconvolution is the process of recovering a sharp image from an input image corrupted by blurring and noise, where the blurring is due to convolution with a known kernel and the noise level is known. Most deblurring approaches rely on decades old deconvolution techniques such as the Lucy-Richardson algorithm [18], Wiener filtering, and least-squares deconvolution. Many of these algorithms were developed for applications where the images are quite different than those taken by an everyday photographer. Consequently, these methods are not always well suited to the desired task and often generate unwanted artifacts such as ringing. One approach to overcome these shortcomings is that of Levin *et al.* [12], which incorporates priors derived from natural image statistics. Other methods have explored the use of graph cuts to reduce over-smoothing [17], deconvolution using multiple blurs [10], and energy minimization functions using wavelets for deconvolution [6].

Levin *et al.* [12] perform non-blind deconvolution using a prior based on assumptions about the edge content of images. The authors assume that images are piecewise smooth and thus the gradient distribution of an image is zero-peaked with high kurtosis. They enforced this property using a hyper-Laplacian prior on image gradients during deconvolution. In our work, we show that when this prior is used alone it has a tendency to overly smooth results.

Yuan *et al.* [23] use a multi-scale Lucy-Richardson deconvolution combined with a bilateral filter to suppress ringing artifacts. While their method suppresses ringing, the noise sensitivity of Lucy-Richardson remains. A further difficulty is their method’s use of ten parameters of which four are user-specified. In contrast, our method is quite robust to noise and only has three parameters with one user-specified parameter. Yuan *et al.* do not address denoising or up-sampling.

3. Deconvolution and Denoising Overview

We model a blurred, noisy image as the convolution of a latent sharp image with a known shift-invariant kernel plus additive white Gaussian noise, whose result is potentially down-sampled. Specifically, blur formation is modeled as:

$$B = D(I \otimes K) + N, \quad (1)$$

where K is the blur kernel, $N \sim \mathcal{N}(0, \sigma^2)$ is the noise, and σ^2 potentially varies spatially. $D(I)$ down-samples an image by point-sampling $I_L(m, n) = I(sm, sn)$ at a sampling rate s for integer pixel coordinates (m, n) . The down-sampling function allows us to perform up-sampling, by solving on a sub-pixel grid. In regular deblurring, $s = 1$, and for up-sampling, $s > 1$.

Our goal is to recover the unobserved sharp image I from only the observed blurred input image B given the kernel K . We use the method of Liu *et al.* [14] to estimate the level-dependent noise level σ^2 .

We formulate the image deconvolution problem using a Bayesian framework and find the most likely estimate of the sharp image I , given the observed blurred image B , the blur kernel K , and the recovered noise level σ^2 using a *maximum a posteriori* (MAP) technique.

We express this as a maximization over the probability distribution of the posterior using Bayes’ rule. The result is a minimization of a sum of negative log likelihoods:

$$P(I|B, K) = P(B|I)P(I)/P(B) \quad (2)$$

$$\arg\max_I P(I|B) = \arg\min_I [L(B|I) + L(I)]. \quad (3)$$

The problem of deconvolution is now reduced to defining and minimizing the negative log likelihood terms. Given the blur formation model (Equation 1), the “data” negative log likelihood is:

$$L(B|I) = \|B - I \otimes K\|^2 / \sigma^2. \quad (4)$$

We incorporate the down-sampling function in Equation 1 by modifying the likelihood term to be $\|B - D(I_H \otimes K_H)\|^2$, where the *up-sampled* latent image and kernel

are I_H and K_H , respectively, and $B_H = I_H \otimes K_H$. $D(I)$ is formed as a simple point-sampling matrix such that $B(m, n) = B_H(sm, sn)$ for a sampling rate s .

To perform denoising alone, we set the kernel K in Equation 4 to a delta function, which reduces the data negative log likelihood to:

$$L(B|I) = \|B - I\|^2 / \sigma^2. \quad (5)$$

The form of the remaining negative log likelihood term, $L(I)$, in Equation 2 depends on the image prior that is used. Defining this term is the main focus of this work.

4. Gradient Priors

In image deconvolution, the data likelihood is inherently ambiguous, i.e., there are many “sharp” images that when blurred match the observed blurred image. The range of ambiguity increases with the amount of blur, and image noise further complicates the issue. The role of the image prior is to disambiguate among the set of possible solutions and to reduce over-fitting to the noise. A common approach is to assume that the image is smooth or piecewise smooth, resulting in priors on image gradients. In the following section, we discuss the limitations of gradient priors and present a novel prior derived from image colors statistics.

Gradient priors are typically enforced between neighboring pixels in an image. These interactions can be modeled using a Markov Random Field (MRF) in which the value of an individual pixel is conditionally dependent on the pixel values in a local neighborhood. One possible prior on the gradients is a smoothness prior in which large image gradients are penalized. Thus, neighboring pixels are favored to have values similar to their neighbors. Typically, this is enforced under a assumption of a Gaussian distribution on the image gradients. While such a prior does disambiguate the solution, it can result in an overly-smooth solution and introduce ringing artifacts [12]. This occurs as a result of the quadratic penalty term, which enforces a Gaussian distribution on gradients. Unfortunately, “natural” images have a decidedly non-Gaussian gradient distribution.

Levin *et al.* [12] address this by modifying the gradient penalty to enforce a hyper-Laplacian distribution:

$$L(I) = \lambda \|\nabla I\|^{0.8}. \quad (6)$$

The value ∇I indicates the spatial gradients of the image, and λ is a regularization parameter that controls the weight of the smoothness penalty. This “sparse” gradient prior better models the zero-peaked and heavy tailed gradient distributions seen in natural images. As the penalty function is no longer a quadratic, the minimization is performed using iterative re-weighted least-squares [20].

Deconvolution using a sparse gradient prior is a significant step towards producing more pleasing results, as it reduces ringing artifacts and noise relative to more traditional techniques. However, this prior has some limitations.

While it biases the deconvolution to produce images with a hyper-Laplacian distribution on gradients, this prior is implemented as a penalty on gradient magnitudes. Thus, it is essentially a “smoothness prior” with a robust penalty function. Using this function, larger gradients still incur larger penalties. This results in a preference for finding the lowest intensity edges that are consistent with the observed blurred image. This is particularly an issue with “bar” type edges and high-frequency texture, as illustrated in Figure 2.

The second limitation of the sparse gradient prior arises in the presence of significant image noise. The implication of the Levin *et al.*’s [12] sub-linear gradient penalty function is that a single large gradient is preferred over many small gradients when accounting for intensity variations. This can result in the preservation and sharpening of the noise. As illustrated in Figure 3, the presence of high-frequency noise that varies on a per-pixel level produces mid-frequency texture patterns, Figure 3c, that can be more objectionable than the original noise. The noise may be removed by increasing the weight of the sparse gradient prior, but over-smooths the result, Figure 3d.

5. Color Priors

We propose an additional term in the sparse prior that uses a color model built from local color statistics of the sharp latent image. This overcomes over-smoothing as it allows for sharp edges as long as they are consistent with local color statistics.

5.1. The Two-Color Model

Recently, researchers have further noted that images can locally be described as a mixture of as few as two colors for use in alpha-matting [13, 1], image denoising [14], and Bayer demosaicing [4]. The two-color model states that any pixel color can be represented as a linear combination of two colors, where these colors are piecewise smooth and can be derived from local properties:

$$I = \alpha P + (1 - \alpha)S. \quad (7)$$

P and S are the respective *primary* and *secondary* colors and α is the linear mixing parameter. For notational convenience, the primary color P_i is always assigned to the color that lies closest to the pixel i ’s color I_i . Some pixels may only be described by a single color, in which case $P_i = S_i$.

The two-color model has several benefits when used to provide an image prior for deconvolution. First, given the two colors for a pixel, the space of unknowns is reduced from three dimensions (RGB) to one (α). The second benefit is that the α parameter provides an alternative for parameterizing edges, where the edge sharpness is decoupled from edge intensity—a single pixel transition in α from 1 to 0 indicates a step edge (an single step from primary to secondary) regardless of the intensity of the edge. Thus, we




Figure 2: (a) Many sharp edges can blur to match an observed blurred (and potentially noisy) edge (in tan). The sparse prior always prefers the smallest intensity gradient that is consistent with the observation (in red). Our method picks the edge that is more likely given the dominant primary and secondary colors in the pixel’s neighborhood. (b) The thin blue areas between the white letters are deconvolved to a mid-level blue with the sparse prior. With our method the blue spaces are more distinct and edges are sharper.

can control an edge’s strength with a prior on α while maintaining local smoothness with a separate prior on P and S .

A significant benefit of the two-color model is its ability to capture local color statistics. We observe that local color statistics can provide a strong constraint during deconvolution. These constraints help reduce over-smoothing around “bar edges” and high-frequency texture as shown in Figure 2b. In contrast with a gradient prior, which prefers the lowest intensity edges that are consistent with the observed blurred image, a two-color model can result in higher intensity edges if such edges are more consistent with local color statistics.

Our two-color model is a Gaussian Mixture Model that is a modified version of the approach used by Bennett *et al.* [4]. We model pixel data for a local neighborhood around each pixel as a mixture of four components. A pixel may belong to the primary color, the secondary color, a mixture of the two colors, or be an outlier. The primary and secondary are modeled using a 3D Gaussian. The model for a mixed color is a 1D Gaussian based on the distance between the pixel’s color and the line segment between the primary and secondary colors. The variance is computed as a function of the image noise σ . A small uniform distribution set to 0.1 is used to model outliers.

The primary and secondary Gaussians are initialized by first performing 10 iterations of k-means clustering (with $k=2$). In the maximization step, when each Gaussian’s standard deviation is recomputed, we clamp the minimum for the primary and secondary Gaussians to be the noise’s standard deviation σ . As a result, after several iterations the Gaussians will merge if the standard deviation is less than the noise’s standard deviation. In this case, we consider the pixel to be modeled by one color, otherwise it is marked as a being a two color pixel. A binary variable indicating one vs. two colors is stored for each pixel. We use a 5×5 window around each image pixel and perform 10 iterations of EM clustering.

5.2. Using the Two-Color Model for Deconvolution

The two color model provides a significant constraint for deblurring; there are two ways such a model can be used for deconvolution. The first is to use the model as a hard constraint, where the sharp image I must always be a linear combination of the primary and secondary colors P and S . The second is to use a soft-constraint to encourage I to lie on the line connecting P and S in RGB space (Figure 4a). We believe that the hard-constraint is too limiting and therefore use the soft-constraint.

Our image negative log-likelihood term is defined as:

$$L(I|P, S) = \lambda_1 \|I - [\alpha P + (1 - \alpha)S]\|^2 \quad (8)$$

$$+ \lambda_2 \rho(\alpha) + \lambda_3 \|\nabla I\|^{0.8}. \quad (9)$$

The first likelihood term minimizes the distance between the recovered intensity I and the line defining the space of the two color model (Figure 4a). In the above equation, α is not a free variable and is computed for a pixel i as:

$$\alpha_i = \left(\frac{(P_i - S_i)}{(P_i - S_i)^T (P_i - S_i)} \right)^T (I_i - S_i). \quad (10)$$

For the “one-color” model, in which $P = S$, we do not use the two color model and the negative log likelihood falls back to using the sparse prior only:

$$L(I) = L(I|P) = \lambda_3 \|\nabla I\|^{0.8}. \quad (11)$$

The second likelihood term, $\rho(\alpha)$, allows us to enforce a prior on the distribution of alpha values that are a function of the normalized distances from I to P and S . The shape of this prior plays a crucial role in enforcing sharpness during deconvolution. For a sharp image, we expect the alpha distribution to be peaked around 0 and 1, which enforces sharp transitions between colors by minimizing the number of pixels with partial α values.

To confirm this expectation and to recover the exact shape for the alpha distribution, we measured the distribution of α values for several hundred images. We computed P and S for a set of *sharp* images and then fit a piecewise hyper-Laplacian penalty function, $b|\alpha|^a$, to the shape




Figure 3: Blurred, noisy image (the PSF is 31x31 pixels and $\sigma = 0.01$), deconvolution with Lucy-Richardson, the sparse prior, our result using the two-color prior, and the groundtruth. A smaller weight on the sparse gradient penalty gives noisy results. Increasing the gradient penalty over-smooths the image. In our results the blue spaces are more distinct and edges are sharper. In these results $\lambda_2 = 0.2$.

of the negative log likelihood of the measured α distributions, shown in Figure 4b. The fit parameters are:

$$\begin{aligned} b = 1.6216, a = 0.0867 & \quad \alpha \geq 0 \\ b = 2.2712, a = 0.2528 & \quad \alpha < 0. \end{aligned} \quad (12)$$

Note that we consider α distributions to be symmetric about $\alpha = 0.5$ and accordingly reflect the data in our measurements and penalty function about this value. As expected, the measured distributions and our penalty function reflect that α values near 0 are more likely.

Note that in Equation 8, we retain the sparse prior of Levin *et al.* as it helps remove small spatial variations in α that are otherwise not heavily penalized by our model. Thus we find an image with sparse gradients that is most consistent with the two-color model. For both equations, we set $\lambda_1 = 0.5$, $\lambda_3 = 1$, and λ_2 in the range of $[0.01, 0.5]$.

6. Recovering the Sharp Image

In the previous sections, we derived negative log-likelihood terms that when minimized allow us to recover a latent sharp image using a prior derived from a two-color model. Unfortunately, minimizing this error function in one step is not straightforward due to the interdependence of I , P , and S . Fortunately, the problem can be decomposed into two more easily solvable sub-problems, which we minimize using an EM-style method.

The minimization is initialized by computing a deconvolution of the image I_0 using the sparse gradient prior alone.

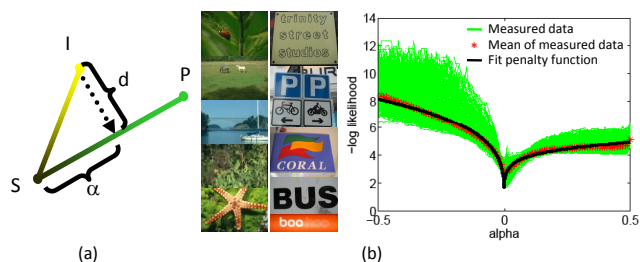


Figure 4: (a) One prior is on the perpendicular distance from a sharp pixel’s color value to the (3D) line defined by the two colors, and a second prior is on the distribution of α ; (b) the measured negative log distributions of alpha values for images from the Berkeley and ICDAR 2003 databases and our piecewise hyper-Laplacian fit to the mean distribution.

We reduce the amount of regularization used by the sparse prior for this initialization to preserve sharpness ($\lambda_3 = 0.25$ for this initialization step and $\lambda_3 = 0.5$ for the subsequent steps). The noise artifacts that result from the reduced regularization are later suppressed when the color model is built. We then iterate the following two steps:

- Estimate P_j and S_j from I_{j-1} by local EM clustering (using a 5x5 window for each pixel).
- Deconvolve the blurred image using P_j and S_j to get I_j by minimizing Equation 2 with Equation 8 as the prior on I_j (using I_{j-1} as an initial guess).

The full definition of our model has two minima in the α penalty function—one at 0 and the other at 1. The only way to properly minimize this exact error function is using a costly non-linear optimization.

Given a few observations about our error function and one approximation, we can use a much faster iterative-least-squares (IRLS) approach. As in the Levin *et al.*’s [12] work, we minimize the hyper-Laplacian α prior using a re-weighting function. Furthermore, due to our alternating minimization that recovers P and S separately from I , terms consisting of P and S can be considered constant and pulled out of the α penalty.

Lastly, instead of using the full bimodal alpha prior, we make an approximation that assigns the identities of P_j and S_j such that P_j is the closest to the initial guess I_{j-1} . We then enforce a unimodal α penalty that biases the values of I_j to be close to P_j . While this now prevents the optimization from allowing the values of I_j to flip between being closer to P_j or S_j , since we perform an alternating minimization, this flip can still happen during the alternation between color estimation and deconvolution. The final alpha penalty term for the IRLS minimization is:

$$\rho(\alpha_i) = \left\| \left[\frac{b|\alpha_i|^{a-2} (P_i - S_i)(P_i - S_i)^T}{((P_i - S_i)^T(P_i - S_i))^2} \right] (I_i - S_i) \right\|^2. \quad (13)$$

The square brackets indicate the “re-weighting” term, which is a 3×3 matrix, for a pixel i . It weights the distance from I_i to S_i by the distance between the P_i and S_i and enforces the shape of our alpha-penalty function. The re-weighting matrix is constant and is pre-computed when




Figure 5: A sweater (the PSF is 30x30 pixels and the median noise is $\sigma = 0.0127$) from Yuan *et al.* and a fountain (the PSF is 39x39 pixels and the median noise is $\sigma = 0.0104$) from Fergus *et al.* We show the blurred image and results from Lucy-Richardson, the sparse prior, and our method. Here $\lambda_2 = 0.05$. Both our results are sharper (e.g. the sweater holes and fountain water streams).

solving for I_i , as it only depends on the values of P_i and S_i . We have found no significant difference between using a full non-linear optimization and using IRLS. Two to three iterations (between color estimation and deconvolution) are sufficient for convergence. In most cases, the perceptual difference after the second iteration is minimal.

7. Results

To validate our deconvolution algorithm, we tested our method on synthetic cases, where the blur kernel and ground-truth sharp image are known, as well as several real images, where the blur kernel was estimated using previously developed methods [8, 22, 11]. For each result, we compare our method to Lucy-Richardson and deconvolution using the sparse-prior. For the sparse prior method, we used the code available online by Levin *et al.* [12].

Figure 3 shows results for the CVPR 2009 logo and text synthetically blurred with a “squiggly” motion-blur kernel with 5% additive white Gaussian noise. These images exhibit numerous “bar edge” type features between the let-

ters. Our method separates and sharpens the letters much more than the Lucy-Richardson and the sparse prior alone. We also perform the same experiment for several images from the Berkeley and ICDAR Databases. Table 2 shows that the PSNR values for our results are consistently higher than those of the other methods. For natural images (the Berkeley database) we out-perform the sparse prior on average by 0.41 DB and up to 0.59 DB for the highest noise case. For text-like images from the ICDAR database, we have an average of a 2.39 DB improvement with our largest improvement at 3.43 DB, which is quite significant. These results further show that our method deblurs better than the previous methods as amount of noise increases.

In Figure 1, we show a result using a real image of a map and PSF from Joshi *et al.* [11]. Our result has the sharpness of the Lucy-Richardson result and is not overly smoothed, as in the result when using the sparse-prior alone. If the regularization weight for using the sparse prior alone is reduced, more textured noise appears creating an effect similar to that in Figure 3. Our result has minimal noise artifacts and no ringing.

In Figure 5, we show results for two real images. The first blurry image and kernel are from the work by Yuan *et al.* [22]. In their work, the authors obtain accurate PSFs for blurry images using a sharp, noisy image and a blurry image. We have used their PSFs for deblurring the blurry images alone. The second result is using an image and PSF from the work of Fergus *et al.* [8]. As in the previous results, in comparison with the other methods, our results are sharper.

To test denoising with our algorithm, we added 5% and 10% additive gaussian noise to several images from the Berkeley Image Database, shown in Figure 6, and ran our algorithm. Table 1 shows a comparison of the PSNR values for our results compared to the results of using Liu *et al.*'s [14] and Portilla *et al.*'s [16] methods – two recent,



Figure 6: Berkeley and ICDAR Images used for testing.

| Image | Denoising $\sigma = 5\%$ | | | | | Denoising $\sigma = 10\%$ | | | | |
|---------|--------------------------|-------|-------|--------|-------|---------------------------|-------|-------|--------|-------|
| | 0th | 1st | Port. | Sparse | Ours | 0th | 1st | Port. | Sparse | Ours |
| B22013 | 32.17 | 32.33 | 31.31 | 32.50 | 32.54 | 27.14 | 28.84 | 27.12 | 27.53 | 27.91 |
| B23084 | 32.04 | 32.64 | 32.14 | 33.01 | 33.15 | 27.04 | 29.23 | 27.24 | 27.20 | 27.74 |
| B35008 | 34.84 | 35.97 | 35.74 | 35.88 | 35.67 | 30.24 | 33.27 | 31.23 | 32.16 | 32.49 |
| B65010 | 31.99 | 32.18 | 30.95 | 32.26 | 32.36 | 27.18 | 28.41 | 26.73 | 27.07 | 27.55 |
| B100075 | 31.69 | 31.68 | 31.27 | 32.36 | 32.41 | 28.14 | 28.96 | 28.31 | 28.51 | 28.80 |
| B105053 | 33.77 | 34.02 | 34.01 | 34.79 | 34.66 | 30.63 | 31.95 | 31.41 | 31.73 | 31.91 |
| B108073 | 31.94 | 31.98 | 31.48 | 32.49 | 32.54 | 28.33 | 29.21 | 27.94 | 27.94 | 28.31 |
| B113009 | 32.31 | 32.61 | 32.89 | 33.62 | 33.64 | 28.89 | 30.19 | 29.91 | 29.87 | 30.27 |
| B134052 | 32.58 | 32.88 | 32.09 | 33.01 | 33.02 | 28.61 | 29.55 | 28.20 | 28.15 | 28.68 |
| Mean | 32.59 | 32.92 | 32.43 | 33.32 | 33.33 | 28.47 | 29.96 | 28.68 | 28.91 | 29.30 |

Table 1: For denoising, our PSNR value are higher than when using Portilla *et al.*'s method, Liu *et al.*'s 0th order denoising, and the sparse prior. Liu *et al.*'s 1th order method results in slightly higher PSNR values than our method for the 10% noise images; however, Liu *et al.*'s results have artifacts (see Figure 7). For all of our results $\lambda_2 = 0.02$.

state of the art color denoising methods. Figure 7 shows a visual comparison for two images. Table 1 also compares our work to denoising using the sparse prior alone. (Note that while denoising using the sparse prior alone is the same as Levin *et al.*'s deconvolution with the data term in Equation 5, Levin *et al.* never specifically address denoising.) PSNR values are consistently higher with our method than when using Portilla *et al.*'s method, Liu *et al.*'s 0th order denoising, and the sparse prior in the higher noise case. Liu *et al.*'s 1th order method results in slightly higher PSNR

values than our method for the 10% noise images; however, Liu *et al.*'s results show significant blocking artifacts, as shown in Figure 7.

Given our formulation, it is also possible to do more traditional up-sampling. We up-sample a low-resolution image by deconvolving it on an up-sampled grid where the PSF is a down-sampling (7-tap binomial filter) anti-aliasing filter. The application of our method to upsampling shares similarities with the work of Fattal [7] and Dai *et al.* [5], in that we all consider alpha priors. Figure 8 shows a 4x up-sampling results using data from Fattal. Our result is significantly sharper than the result of bi-cubic interpolation and is similar to Fattal's result.

To view full resolution versions of our results, including additional examples, visit <http://research.microsoft.com/ivm/twocolordeconvolution/>.

8. Discussion and Future Work

We have developed a unified framework to deblur, denoise, and up-sample images using a novel prior that incorporates local color statistics. Our method produces sharper results with less noise and quantitatively out-performs previous methods.

Our results suggest several areas for future work. Perhaps the largest limitation of our method is not in its theory, but in its practice. For deblurring, recovering both the color model and deconvolving the image is a non-linear problem. Furthermore, the error function we use can suffer from local minima due to the necessary existence of two minima in the alpha prior. Two obvious routes exist for improvement, the first being to investigate alternative optimization techniques. The second is to improve the initialization for the color model; we have experimented with both the Lucy-Richardson algorithm and using the sparse prior alone. We found that using the sparse prior alone with a low weighting provides a good initial guess; however, due to our current optimization method, the quality of our results is somewhat bound by this initialization. Other choices may yield improved results.

| Image | Deblurring $\sigma = 2\%$ | | | Deblurring $\sigma = 5\%$ | | |
|---------|---------------------------|--------|-------|---------------------------|--------|-------|
| | LR | Sparse | Ours | LR | Sparse | Ours |
| B22013 | 24.26 | 27.87 | 28.19 | 19.10 | 25.08 | 25.51 |
| B23084 | 22.88 | 26.16 | 26.60 | 28.49 | 22.93 | 23.52 |
| B35008 | 26.93 | 33.70 | 33.86 | 19.08 | 31.08 | 31.57 |
| B65010 | 24.11 | 26.54 | 26.87 | 18.61 | 24.52 | 24.83 |
| B100075 | 25.70 | 28.90 | 29.20 | 18.85 | 27.13 | 27.34 |
| B105053 | 27.00 | 32.80 | 32.98 | 18.84 | 30.78 | 31.01 |
| B108073 | 24.79 | 27.24 | 27.58 | 18.77 | 24.32 | 24.81 |
| B113009 | 26.62 | 31.18 | 31.31 | 19.22 | 28.56 | 28.90 |
| B134052 | 25.39 | 27.72 | 28.31 | 18.71 | 24.80 | 25.44 |
| I2460 | 27.17 | 33.22 | 33.03 | 20.13 | 28.09 | 30.59 |
| I2468 | 22.29 | 26.66 | 27.77 | 18.36 | 22.80 | 24.89 |
| I2469 | 25.98 | 31.69 | 32.49 | 19.34 | 26.35 | 29.38 |
| I2471 | 24.54 | 29.37 | 30.09 | 19.09 | 24.39 | 27.31 |
| I2483 | 27.33 | 35.55 | 34.02 | 19.54 | 31.79 | 33.13 |
| I2485 | 25.41 | 30.93 | 31.45 | 19.51 | 26.41 | 28.81 |
| I2506 | 22.56 | 32.75 | 32.62 | 18.62 | 28.87 | 30.25 |
| I2530 | 24.80 | 34.35 | 34.00 | 19.23 | 28.96 | 31.72 |
| I2531 | 26.42 | 31.83 | 32.02 | 19.54 | 26.33 | 29.76 |
| I2542 | 22.93 | 26.19 | 27.54 | 18.84 | 22.60 | 24.61 |
| B Mean | 25.30 | 29.12 | 29.43 | 19.96 | 26.58 | 26.99 |
| I Mean | 24.94 | 31.25 | 31.50 | 19.22 | 26.66 | 29.05 |

Table 2: For deblurring, our PSNR values are higher than the results of using the Lucy-Richardson algorithm. For natural images from the Berkeley database (images starting with B) we out-perform the sparse prior on average by 0.41 DB for the highest noise case. For text-like images from the ICDAR database (images starting with I), we have an average of a 2.39 DB improvement. For the "B" results $\lambda_2 = 0.02$ and for the "I" results $\lambda_2 = 0.2$.

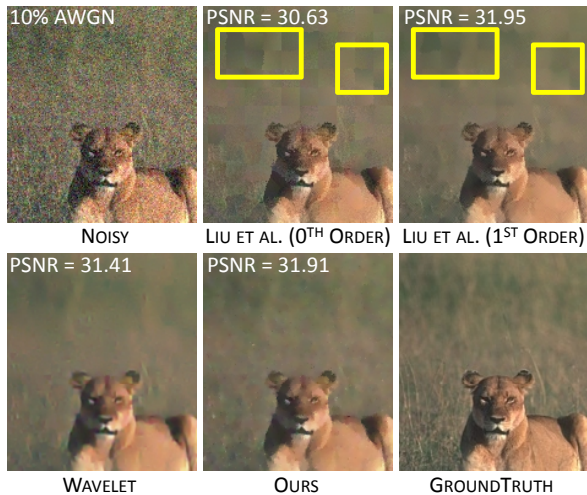


Figure 7: Our result has the sharpness of the Liu *et al.* results in the face region, but does not have the blocking artifacts in the low-frequency background region (see the areas indicated with yellow rectangles). For all of the denoising results $\lambda_2 = 0.02$.

In our experience, we have found that varying the weighting of the alpha priors can help create better results. Most often, text-like images require a higher weight than natural images. We are interested in exploring the alpha prior and weighting values in a class-specific way. We believe that there may be a consistent, but different, set of weights for text versus natural images.

Acknowledgements: We thank the anonymous reviewers for their comments. This work was partially completed while the first author was an intern at Microsoft Research and student at UCSD.

References

[1] Y. Bando, B.-Y. Chen, and T. Nishita. Extracting depth and matte using a color-filtered aperture. *ACM Transactions on Graphics*, 27(5):134:1–134:9, Dec. 2008.



Figure 8: Our formulation also allows us to perform up-sampling of low-resolution images.

[2] B. Bascle, A. Blake, and A. Zisserman. Motion deblurring and super-resolution from an image sequence. In *ECCV '96*, pages 573–582, 1996.

[3] M. Ben-Ezra and S. K. Nayar. Motion-based motion deblurring. *PAMI*, 26(6):689–698, 2004.

[4] E. Bennett, M. Uyttendaele, C. Zitnick, R. Szeliski, and S. Kang. Video and image bayesian demosaicing with a two color image prior. In *ECCV '06*, pages 508–521, 2006.

[5] S. Dai, M. Han, W. Xu, Y. Wu, and Y. Gong. Soft edge smoothness prior for alpha channel super resolution. In *CVPR '07*, pages 1–8, 2007.

[6] P. de Rivaz and N. Kingsbury. Bayesian image deconvolution and denoising using complex wavelets. *ICIP 2001*, 2:273–276 vol.2, 2001.

[7] R. Fattal. Image upsampling via imposed edge statistics. *ACM Transactions on Graphics*, 26(3):95:1–95:8, July 2007.

[8] R. Fergus, B. Singh, A. Hertzmann, S. T. Roweis, and W. T. Freeman. Removing camera shake from a single photograph. *ACM Transactions on Graphics*, 25(3):787–794, July 2006.

[9] W. T. Freeman, T. R. Jones, and E. C. Pasztor. Example-based super-resolution. *IEEE CGA*, 22(2):56–65, March-April 2002.

[10] Y. Harikumar, G.; Bresler. Exact image deconvolution from multiple fir blurs. *IEEE TIP*, 8(6):846–862, 1999.

[11] N. Joshi, R. Szeliski, and D. Kriegman. PSF estimation using sharp edge prediction. In *CVPR '08*, pages 1–8, 2008.

[12] 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*, 26(3):70:1–70:9, July 2007.

[13] A. Levin, D. Lischinski, and Y. Weiss. A closed form solution to natural image matting. *PAMI*, 30(2):228–242, 2008.

[14] C. Liu, R. Szeliski, S. B. Kang, C. L. Zitnick, and W. T. Freeman. Automatic estimation and removal of noise from a single image. *PAMI*, 30(2):299–314, 2008.

[15] P. Perona and J. Malik. Scale-space and edge detection using anisotropic diffusion. *PAMI*, 12(7):629–639, 1990.

[16] J. Portilla, V. Strela, M. Wainwright, and E. Simoncelli. Image denoising using scale mixtures of gaussians in the wavelet domain. *IEEE TIP*, 12(11):1338–1351, 2003.

[17] A. Raj and R. Zabih. A graph cut algorithm for generalized image deconvolution. In *ICCV '05*, pages 1048–1054, 2005.

[18] W. H. Richardson. Bayesian-based iterative method of image restoration. *Journal of the Optical Society of America (1917–1983)*, 62:55–59, 1972.

[19] S. Roth and M. J. Black. Fields of experts: A framework for learning image priors. In *CVPR '05*, pages 860–867, 2005.

[20] C. V. Stewart. Robust parameter estimation in computer vision. *SIAM Reviews*, 41(3):513–537, September 1999.

[21] C. Tomasi and R. Manduchi. Bilateral filtering for gray and color images. *Computer Vision, 1998. Sixth International Conference on*, pages 839–846, 1998.

[22] L. Yuan, J. Sun, L. Quan, and H.-Y. Shum. Image deblurring with blurred/noisy image pairs. *ACM Transactions on Graphics*, 26(3):1:1–1:10, July 2007.

[23] L. Yuan, J. Sun, L. Quan, and H.-Y. Shum. Progressive inter-scale and intra-scale non-blind image deconvolution. *ACM Transactions on Graphics*, 27(3):74:1–74:10, Aug. 2008.