Image Deblurring

Brandon Smith Mikola Lysenko

March 29, 2009

1 The Problem

Blur is a common and unwanted artifact of image acquisition. There are many reasons why images become blurred such as movement, slow shutter speed or incorrect focal distance. Because blurry images are confusing, they are often less appealing and so it is desirable to deblur images for purely aesthetic reasons. Outside of photography, variable blurring creates problems for feature tracking algorithms which rely on the existence of consistent image structures. Blur removal may be used to precondition these algorithms and greatly improve performance.

The most common approach of blur removal is to treat blur as a noisy convolution operation. This model has both its advantages and disadvantages which we discuss in greater detail in the subsequent sections. From the perspective of convolution, there are two basic types of blur removal: non-blind and blind, with the main difference being that in the case of non-blind deconvolution the blur kernel is known.

For our project, we researched multiple blind and non-blind deconvolution methods to gain an understanding of the current state of the art. Given the time limitations of the class project, we were not able to probe much beyond this point, but we did learn of a few interesting and somewhat counter-intuitive properties of the convultion blur model. We also performed some experiments to acquire kernels directly from blurred images using cheap hardware. Finally, we implemented and compared several recent methods which we discuss in this report.

1.1 History

The earliest formulation of this problem was given by Wiener along with a solution in the case of noise with a known power spectrum [8]. Later, Richardson and Lucy proposed an alternative method robust to arbitrary Poisson noise [5, 6]. Of course these methods fail when the measured noise does not match the model, resulting in unpleasant ringing artifacts. To address this, there have been many proposed regularization approaches based on techniques like the bilateral filter [9] or edge statistics [7].

Unlike non-blind deconvolution, blind deconvolution is an ill-posed problem in the sense that a single image does not contain enough information to reconstruct both an image and an arbitrary kernel. To make this problem well-posed, it necessary to add some type of regularization. One of the best priors in practice is to minimize the support of the kernel in the space domain, which results in as small a kernel as possible. This technique was introduced by Levin to solve for motion blur, though the same idea may be used in conjunction with non-blind deconvolution to obtain deblurred image/kernel pairs [4].

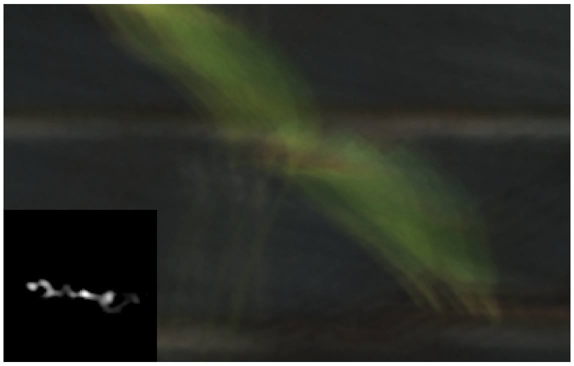

(a) Original image

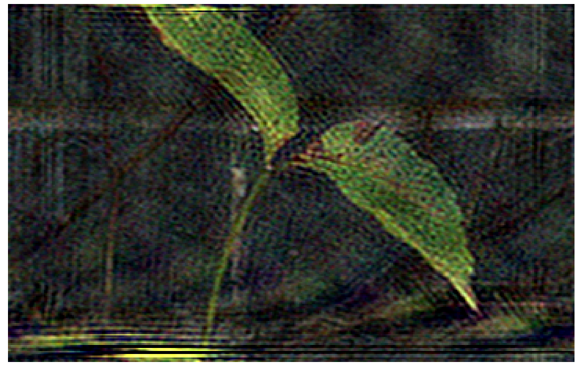

(b) Wiener filter

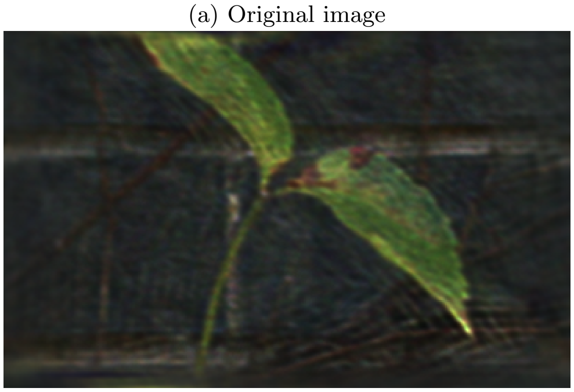

(c) Richardson-Lucy

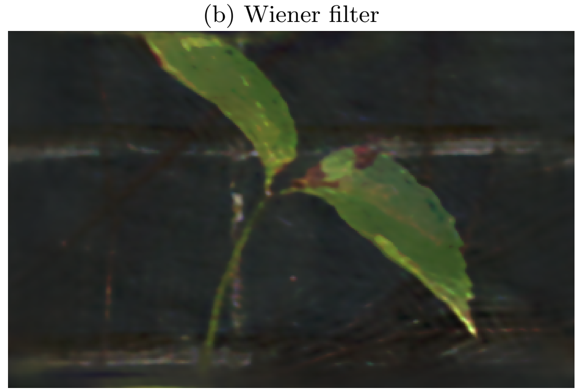

(d) Bilateral Richardson-Lucy

(e) Shan *et al.*

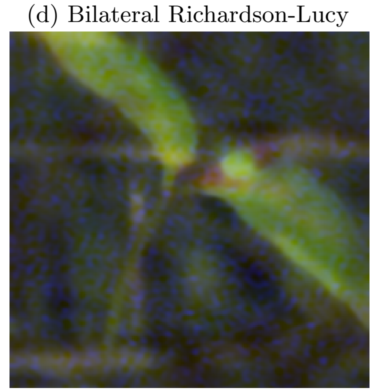

(f) Total Variation Minimization

Figure 1: Shown above is a comparison of results using five non-blind deconvolution methods on a blurry image of a plant. The kernel used for deconvolution is shown in the lower left corner of the original image (a). Despite the severe blur in the original image, this is not necessarily a very difficult case. If the blur kernel is accurate and the image has low noise, which we believe to be the case here, then the size of the blur kernel is largely irrelevant.

(a) Original image

(b) Wiener filter

(c) Richardson-Lucy

(d) Bilateral Richardson-Lucy

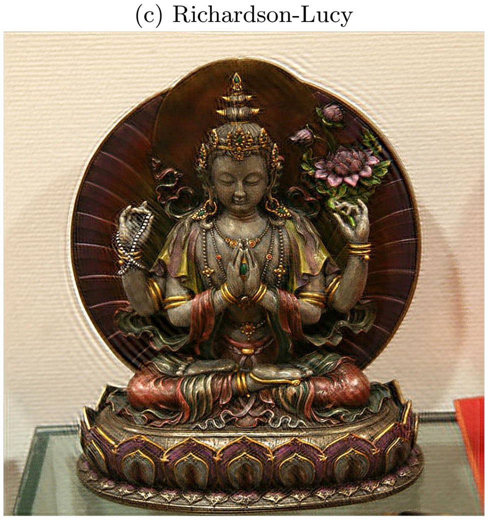

(e) Shan *et al.*

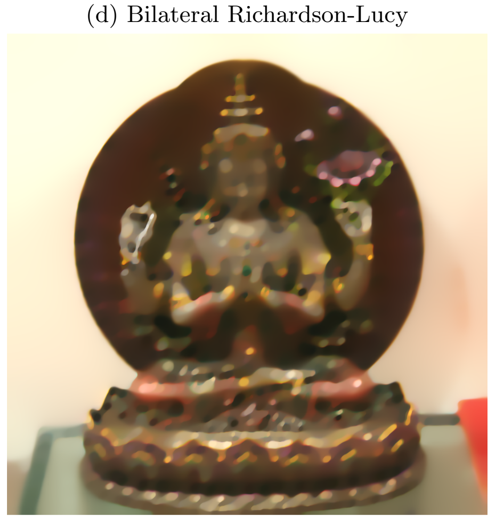

(f) Total Variation Minimization

Figure 2: Shown above is a comparison of results using five non-blind deconvolution methods on a blurry image of a statue of the Hindu goddess Lakshmi. The kernel used for deconvolution is shown in the lower left corner of the original image (a). All methods perform reasonably well with the exception of the total variation minimization method, which over smooths regions of high variation.

2 Analysis

We begin with a formal description of the convolution model for blurred image formation:

$$I = L \otimes f + n, \tag{1}$$

where $I(x, y)$ is the resulting blurred image, $L(x, y)$ is some unknown deblurred image, $f(x, y)$ is the blur kernel and $n(x, y)$ an unknown noise term. There are several problems with this model, such as the fact that it does not account for object motion or parallax effects, but it remains applicable in a wide range of scenarios. For image deconvolution, our goal is to solve for L given I and (optionally) f. If $n = 0$ then by the convolution theorem we have:

$$\begin{aligned} \hat{I} &= \hat{L}\hat{f} \\ \hat{L} &= \frac{\hat{I}}{\hat{f}}, \end{aligned}$$

which has a unique solution for L given any $\hat{f}$ such that $\hat{f}(\omega_x, \omega_y) \neq 0$ for all ω_x, ω_y (with $\hat{f}$ denoting the Fourier transform of f). This fact is easy to check numerically on artificially convolved images.

In the real world, the solution for Eq. (1) is substantially complicated by the existence of noise. In fact, the ringing artifacts associated with most deblurring algorithms are a result of incorrectly modeling noise, not Gibbs phenomena as they are often incorrectly attributed to. We view the way different algorithms deal with noise as their distinguishing characteristic.

2.1 Wiener

One of the earliest attempts at robust deconvolution is due to Wiener[8]. He observed that the mean squared error of some estimated signal, L', may be expressed as:

$$\epsilon(\omega_x, \omega_y) = E(|\hat{L}(\omega_x, \omega_y) - \hat{L}'(\omega_x, \omega_y)|^2)$$

Where $E(\cdot)$ denotes the expectation taken with respect to the set of all possible signals and noise vectors in this case. If we substitute $L' = f' \otimes I$ into the above equation then take the derivative with respect to f, we find that the solution for f' which minimizes this error is:

$$\hat{f}'(\omega_x, \omega_y) = \frac{1}{\hat{f}(\omega_x, \omega_y)} \left(\frac{|\hat{f}(\omega_x, \omega_y)|^2}{|\hat{f}(\omega_x, \omega_y)|^2 + \frac{E(|\hat{n}(\omega_x, \omega_y)|^2)}{E(|\hat{I}(\omega_x, \omega_y)|^2)}} \right)$$

The tricky part is finding a good guess for $\frac{E(|\hat{I}(\omega_x, \omega_y)|^2)}{E(|\hat{n}(\omega_x, \omega_y)|^2)}$, also known as the signal to noise ratio. Once we have an expression for $\hat{f}'$, solving for L' is performed exactly as in the noiseless case as described above.

2.2 Richardson-Lucy

Richardson [6] and Lucy [5] both independently discovered what is now referred to as the *Richardson-Lucy* (RL) deconvolution method (Richardson discovered it first and then Lucy discovered on his own soon after). RL deconvolution, unlike the Wiener method, is more robust to noise because it explicitly models it as a Poisson process. Although this does not very accurately reflect real noise, it works better than the power spectrum used by the Wiener method.

The major benefit of the RL algorithm is that it is simple and fast. The RL method reduces to the following update equation.

$$I^{t+1} = I^t \left[f^* \otimes \frac{L}{I^t \otimes K} \right], \tag{2}$$

where f^* is the adjoint of f and t is the time step. Significant ringing is still present in RL results (shown in the following sections) because of the poor noise model.

2.3 Bilateral Richardson-Lucy

A simple extension of the RL method called the *Bilateral Richardson-Lucy* (BRL) method, as described in Yuan *et al.*'s work [9], does not improve the noise model of the RL method, but instead directly tackles the ringing problem by applying a bilateral filter at each iteration. The BRL update equation is simply

$$I^{t+1} = I^t \left[f^* \otimes \frac{L}{I^t \otimes K} \right] \cdot B(I^t), \tag{3}$$

where $B(I^t)$ is the bilateral filter portion. $B(I^t)$ has a regularization weight λ that must be set differently depending on the input image characteristics, which Yuan *et al.* [9] describe in more detail. Intuitively, a bilateral filter smooths nearby and similar pixels, but not distant and/or dissimilar pixels. We show in Section 3 that the ringing artifacts are reduced by applying a bilateral filter at each iteration. However, this comes with a price as some detail is lost and color and intensity boundaries can become unnaturally sharp.

2.4 Joint Bilateral Richardson-Lucy

As a final extension, Yuan *et al.* [9] implemented what they call the *Joint Bilateral Richardson-Lucy* (JBRL) method. Again, this extension does not improve the noise model of the RL method, but it does target the noise problem. JBRL creates a set of progressively more down-sampled versions of the original blurry image. It starts by deconvolving the most down-sampled image, then uses the result to "guide" the deconvolution of the next higher resolution image. This continues until the highest resolution image is deblurred. More specifically, JBRL changes the bilateral filter portion of the BRL update function to include an extra multiplicative penalty dependent on the intensities of the previous, lower resolution image. In practice, they implement JBRL by approximating the down-sampling with a small Gaussian blur. Intuitively, JBRL starts by aggressively removing noise, which results in a poor deconvolved result in terms of lack of sharpness, but one that is relatively free of ringing artifacts; this result then serves to dampen the ringing that arises in the deconvolution of the next higher resolution image.

2.5 Shan *et al.*'s Method

Shan et al. proposed yet another model for the noise within the image based on an extension of the Weiner filter. Like before, they assume that the noise is a collection of independent identically distributed Gaussian functions, though they make the further assumption that within locally constant regions the amount of noise is reduced. This removes ringing artifacts like the BRL method while maintaining the theoretical advantage of being a MAP problem. In practice, the difference between Shan's method and Weiner filtering is only apparent when there exist many constant regions within the blurred image.

2.6 Total Variation Minimization

Using the time we received for an extension, we also tried implementing total variation minimization for deblurring [3]. Total variation minimization is unlike any of the above methods in that it does not model noise, but rather the images. The key idea follows along the lines of compressed sensing: assume that the image has sparse support in some basis, typically the Fourier domain or wavelets, then find the smallest support within said basis that represents the image faithfully. For total variation minimization, this is tweaked slightly to finding the basis with the minimal derivative (in the Sobolev sense), which is also equivalent to minimizing the support of the frequency representation. In the general case, finding an optimal basis is NP-complete, but an approximation via the L1 relaxation problem has been proven to give good results [2] and may be solved in polynomial time

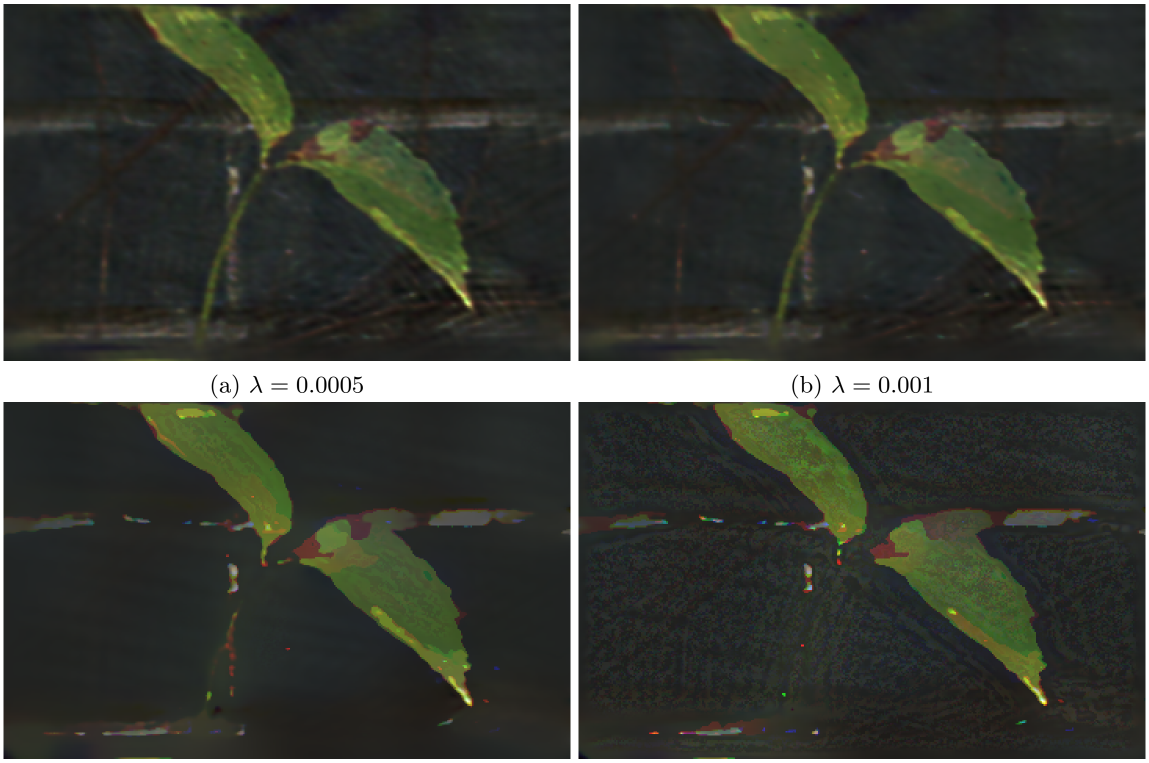

(a) $\lambda = 0.0005$ (b) $\lambda = 0.001$

(c) $\lambda = 0.005$ (d) $\lambda = 0.01$

Figure 3: Shown above is a comparison of results generated using different regularization coefficients, λ, in the BRL method. Small λ behaves like RL (the bilateral filter becomes insignificant), and large λ results in over segmentation.

using second order cone programming. To simplify programming, we used the L1 Magic library [1] for MATLAB.

3 Experimental Results and Discussion

We experimented with several non-blind deconvolution algorithms [5, 6, 8, 9], and one blind deconvolution algorithm [7]. We discuss our results in the following subsections.

3.1 Non-blind Deconvolution

In this section we discuss our non-blind deconvolution results, which are shown in Figures 1, 2, and 4. **These figures are best viewed electronically.**

The simplest method of the five that we experimented with is the Wiener filter [8]. Noise tends to confound the Wiener filter method most, which we discussed in Section 2.1. Figures 1(b), 2(b), and 4(b) show that the Wiener filter produces noticeably more ringing artifacts than the other methods due to its poor handling of noise. However, one benefit of the Wiener filter is its speed and simplicity; it completes in one or two seconds on a medium size image.

Like the Wiener filter the RL method is very simple, but tends to handle noise slightly better at a small price in computation cost. Figures 1(c), 2(c), and 4(c) show that the RL method produces less ringing than the Wiener filter, but the ringing is still quite noticeable.

(a) Original image

(b) Wiener filter

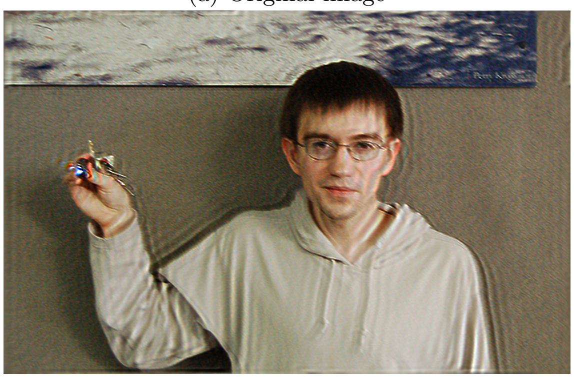

(c) Richardson-Lucy

(d) Bilateral Richardson-Lucy

(e) Shan *et al.*

(f) Total Variation Minimization

Figure 4: Shown above is the result of an experiment to manually estimate the blur kernel from a point light source in the original, blurry image. The blur kernel that we estimated is shown in the lower left corner of the original image (a). The Wiener filter result (b) contains the most ringing, following by the Richardson-Lucy result (c), Shan *et al.*'s result (e), the Bilateral Richardson-Lucy result (d), and finally the total variation minimization result (f). Total variation minimization tends to over smooth the result in regions of high variation. BRL performs well in this case, removing most of the ringing while still producing a crisp result.

The BRL method is more complicated than the RL method and runs *much* slower due to the bilateral filtering, but is effective at reducing ringing artifacts. Figures 1(d), 2(d), 3 and 4(d) show this. Yuan *et al.*seem to suggest applying the bilateral filter at each iteration, but this results in a very slow algorithm and in our experiments it tends to either filter the image too much during early iterations, or not enough overall (depending on the λ regularization weight). Instead, for our experiments, we use 10 iterations of the RL method, following by one iteration of BRL. This is then repeated until a satisfactory result is achieved, usually in less than 10 repetitions.

Shan *et al.*'s results are similar to the Wiener filter results in areas of high contrast, e.g., the leaf in Figure 1 or the face in Figure 4. The ringing is only slightly reduced when compared with the Wiener filter results in these areas. However, Shan *et al.*'s method does a good job of suppressing ringing in locally constant regions, e.g., the background in Figure 1 or the wall in Figure 4. This is due to a novel prior they introduce that constrains the unblurred image gradient to be similar to the blurred image gradient in locally constant regions. We found that, despite a fair amount of circumlocution and complicated math, this simple prior was the bulk of Shan *et al.*'s contribution to deblurring.

Although the problem formulation is interesting, our preliminary results for the total variation minimization method are not very good, as can be seen in Figures 1(f), 2(f), and 4(f). This is likely due in part to a non-ideal setting of the coefficients/weights in the minimization equation. More specifically, the L1 term in our equation may be too strong, which explains the over smoothing in regions of high variation, e.g., the finer details in Figure 2. One of the disadvantages of this method is the significant computation time required, which limited our ability to experiment with and optimize coefficient values. For example, the result shown in Figure 2(f) took 5+ hours of computation time on our unoptimized code. Furthermore, our code currently requires a square input image, hence the cropped results.

3.1.1 BRL Regularization

The results of the BRL method depend heavily on the regularization coefficient λ of the bilateral filter. Figure 3 shows a comparison of results generating using several coefficients. As the regularization coefficient approaches zero, BRL behaves like RL. Large coefficients, on the other hand, tend to over-segment the result.

3.2 Manual Kernel Estimation

In special cases it is possible to manually estimate the blur kernel with high accuracy. For example, if a point light source exists in the scene and the blur in the image is due to camera motion, we can trace the light source in the image to recover the blur pattern. We experimented with this idea by taking a picture of a person holding a small key chain light with a moving camera. Figure 4(a) shows this picture, which contains a noticeable amount of motion blur. The smear pattern of the point light source reflects the blur pattern.

We can extract this pattern to get an accurate estimate of the blur kernel by (1) cropping the point light source, (2) converting the cropped image to grayscale, (3) using a "magic wand" image editing tool to select the bright region, and (4) setting the surrounding region to black. The resulting small image can then be used as the blur kernel for any non-blind deconvolution method. We show the results of four deblurring methods in Figure 4 using this manually estimated kernel.

3.3 Kernel Estimation and Blind Deconvolution

Levin [4] use the original blurry image and a deconvolved version of it to estimate the blur kernel. The basic idea in Levin's method is to solve for the kernel by minimizing its support within the space basis. In a sense, this is dual to the total variation minimization method in that it performs the minimization of the space support rather than frequency. Though we were able to achieve reasonable

(a) Original image

(b) Shan *et al.*result

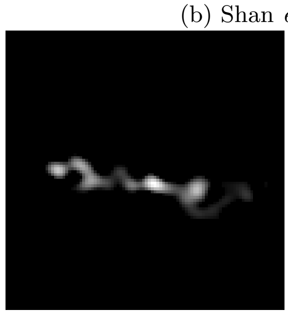 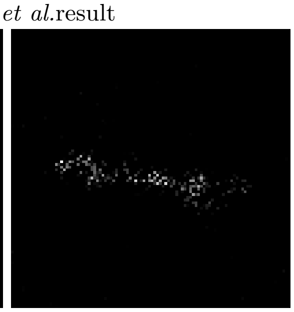

(c) Groundtruth kernel (d) Estimated kernel

Figure 5: Here we show the results of our implementation of Levin's kernel estimation method [4]. The original blurry image is shown in (a) and the deblurred image is shown in (b), which was recovered using the groundtruth kernel shown in (c). Given (a) and (b), Levin's method gives the kernel estimate shown in (d). The quality of the kernel estimate is dependent on the quality of the deblurred result, which itself is dependent on the quality of the blur kernel estimate.

results from images which we deblurred using a known good kernel, we found that the deblurring algorithms were too sensitive to the initial choice of kernel for this to be of much use in practice.

4 Conclusions

In this work we have investigated the characteristics of six different deblurring methods and one kernel estimation method. The deblurring methods can perhaps be categorized along three lines: (1) those based on the Wiener filter, (2) those based on the classic Richardson-Lucy method, and (3) total variation minimization. We have found that state-of-the-art methods are typically founded on one of these three basic approaches, with a set of heuristics-based extensions added on. The extensions typically target undesirable artifacts in the deblurred result (i.e., ringing) as opposed to the noise problem that underlies these artifacts. If the kernel is known, noise is the greatest limiting factor when it comes to achieving a good result. Therefore, as part of future work, it would be interesting to further investigate image noise in terms of its cause and how it might be better modeled. With a better noise model, it is very likely that better image deblurring is possible.

References

[1] Emmanuel Candès and Justin Romberg. L1-magic: Recovery of sparse signals via convex programming, http://www.acm.caltech.edu/l1magic/.

[2] Emmanuel Candès and Terence Tao. Near-optimal signal recovery from random projections and universal encoding strategies. In *IEEE Transactions on Information Theory*, November 2004.

[3] T. Chan, S. Esedogluy, F. Parky, and A. Yipy. Recent developments in total variation image restoration. *Mathematical Models in Computer Vision: The Handbook*, December 2004.

[4] Anat Levin. Blind motion deblurring using image statistics. *Advances in Neural Information Processing Systems*, December 2006.

[5] L.B. Lucy. An iterative technique for the rectification of observed distributions. *Astronomy Journal*, 79:745, 1974.

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

[7] Qi Shan, Jiaya Jia, and Aseem Agarwala. High-quality motion deblurring from a single image. *ACM Transactions on Graphics (SIGGRAPH)*, 2008.

[8] Norbert Wiener. *Extrapolation, Interpolation, and Smoothing of Stationary Time Series*. The MIT Press, 1964.

[9] Lu Yuan, Jian Sun, Long Quan, and Heung-Yeung Shum. Progressive inter-scale and intra-scale non-blind image deconvolution. In *SIGGRAPH '08: ACM SIGGRAPH 2008 papers*, pages 1–10, New York, NY, USA, 2008. ACM.