

MIT Open Access Articles

A new robust method for two-dimensional inverse filtering

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Fuller, Megan M., and Jae S. Lim. "A New Robust Method for Two-Dimensional Inverse Filtering." Edited by Amir Said, Onur G. Guleryuz, and Robert L. Stevenson. Visual Information Processing and Communication VI (March 4, 2015). © 2015 Society of Photo-Optical Instrumentation Engineers (SPIE)

As Published: <http://dx.doi.org/10.1117/12.2078070>

Publisher: SPIE

Persistent URL: <http://hdl.handle.net/1721.1/100816>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of Use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.



A New Robust Method for Two-Dimensional Inverse Filtering

Megan M. Fuller and Jae S. Lim

Research Laboratory of Electronics
Department of Electrical Engineering and Computer Science
Massachusetts Institute of Technology
Cambridge, MA, 02139, USA

ABSTRACT

In this paper, we present a novel method for inverse filtering a two dimensional (2-D) signal using phase-based processing techniques. A 2-D sequence can be represented by a sufficient number of samples of the phase of its Fourier transform and its region of support. This is exploited to perform deconvolution. We examine the effects of additive noise and incomplete knowledge of the point spread function on the performance of this deconvolution method and compare it with other 2-D deconvolution methods. The problem of finding the region of support will also be briefly addressed. Finally, an application example will be presented.

Keywords: inverse filtering, deconvolution, phase-based processing

1. INTRODUCTION

A common problem that arises in signal processing is that of obtaining an original sequence from a sequence that has been filtered or blurred. This is referred to as the inverse filtering problem. We are interested specifically in inverse filtering 2-D signals.

The model we will consider is of a 2-D signal being convolved with a spatially invariant point spread function (PSF) and then corrupted by additive noise. This model is widely used in practice in fields as diverse as archeology,¹ medical technology (specifically ultrasound,² x-ray and CT scan imaging³), astronomy⁴ and general image restoration.⁵

The most straightforward way to solve this problem is to simply divide the Fourier transform of the degraded signal by the Fourier transform of the PSF. However, this gives poor results in the presence of additive noise. Because filtering happens before the noise is added, the frequencies that have been most attenuated by the filter tend to have the lowest SNR, but these are precisely the frequencies that the inverse filter amplifies.⁶

As an attempt to mitigate this, a cap is sometimes put on the magnitude of the Fourier transform of the inverse PSF. That is, all frequencies with magnitudes above a given threshold are set to be the capped value. This reduces the problem of noise amplification, but it is ad-hoc and still does not produce good results.

In this paper, we propose a new inverse filtering method based on the phase of the degraded signal. Specifically, it has been shown by Hayes⁷ that a typical digital image of $N_1 \times N_2$ pixels can be reconstructed to within a scale factor using $(2N_1 - 1) \times (2N_2 - 1)$ equally spaced samples of the phase of its Fourier transform and its region of support (ROS). The reconstruction can be performed by an iterative algorithm which alternately applies the spatial domain constraints (the ROS) and the Fourier domain constraints (the phase samples). This iteration has been proven to converge to the desired solution^{7,8}.

The method we propose applies this to inverse filtering by subtracting the phase of the PSF from the phase of the degraded signal and then using the reconstruction algorithm to obtain the restored image. We will refer to this method as phase-based deconvolution. When there is no noise, this leads to the same result as dividing the Fourier transform of the degraded image by the Fourier transform of the PSF.

In the presence of noise, the behavior of phase-based deconvolution is very difficult to analyze theoretically. It requires only a subtraction, not a division, and so there is some basis for expecting noise amplification to not be as dramatic as it is when doing direct inverse filtering. However, the relationship between additive noise in the spatial domain and phase errors in the frequency domain is complicated and non-linear, and the effect that

the errors in the “corrected” phase will have on the spatial domain image is similarly complicated. There is no theoretical guarantee that an image even exists which satisfies both the Fourier and spatial domain constraints.

Because of the difficulties in analyzing phase-based deconvolution theoretically, we will do so experimentally by computer simulation using digital images that have been blurred. Such experiments are useful measures of the performance of phase-based deconvolution because image restoration is a common problem and potential application (although phase-based deconvolution is not an image restoration system itself—it is agnostic to the noise and does not perform any noise reduction) and because it provides a good way to see visually the results of performing phase-based deconvolution.

There have been many methods developed for solving the specific problem of restoring blurred, noisy images, and these all rely on assumptions about the noise, the original image, or both. For example, Wiener filtering assumes the power spectral densities of the image and noise are known.⁶ Variations and extensions of Wiener filtering, such as regularized filtering⁵ or the method proposed by Dabov et. al.⁹ assume the noise variance is known. The Richardson-Lucy algorithm assumes the noise is Poisson.¹⁰ More recently, wavelet techniques have been developed that depend on assumptions about the noise distribution.¹¹ Total variation techniques rely on the noise variance.¹² Methods that depend on an assumption that the image is sparse in some domain have also been explored^{13,14} as have methods that assume a particular distribution of the image coefficients¹⁵ or gradients^{16,17,18}. Methods that make use of image dictionaries obtained from a training set have also been studied^{19,20}.

This list is not meant to be exhaustive. Many other image restoration algorithms exist, but the examples given here illustrate the kinds of approaches taken to solving this problem. The results produced by these methods can be quite good if the assumptions they make are correct. However, in some situations, it may be desirable to have a more general method—one that does not rely on any assumptions about the image or the noise but is still robust to noise. For example, this may be the case when the noise statistics are changing or difficult to measure, or when the signal to be restored is not an image. As we will show, phase-based deconvolution is such a method.

This paper is organized as follows. In the next section, we present the performance of phase-based deconvolution in the presence of noise, and in Section 3, we compare phase-based deconvolution to some of the image restoration methods mentioned above as a demonstration of how robust it is. In Section 4, we examine the case when the PSF is not known completely. In Section 5, we briefly explore how the ROS of the image may be found. In Section 6, we look at a specific application example, and we conclude the paper in Section 7.

2. PERFORMANCE IN THE PRESENCE OF NOISE

Before we discuss performance, an error measure is needed. This is not a simple matter, as different types of artifacts can have quite different effects on the viewer, even if they lead to a similar quantitative error. The error measure used here is a variation of mean squared error where the restored image is first scaled by a factor chosen to minimize the mean squared error, and then the mean squared error is computed. This will be referred to as optimal scaling mean squared error (OS-MSE). This measure was chosen both because phase-based deconvolution reconstructs the image only to within a scaling factor and because the measure was reasonably consistent with subjective visual quality. However, it should be thought of as only a rough indicator of quality. Example images will also be included to illustrate the different types of artifacts. The problem of choosing the scaling factor in practice will not be addressed here except to indicate that it can be estimated from the degraded image.

When there is no noise, the phase-based reconstruction algorithm is guaranteed to converge. When there is noise, there is not a theoretical guarantee, but our experiments indicate that it does.²¹ Furthermore, when the iteration is started with a constant initial magnitude, the solution tends to improve as the iteration progresses, and the convergent solution is almost always the optimal solution or very close to the optimal solution in the OS-MSE sense.

We will now compare phase-based deconvolution to other inverse filtering techniques—that is, to methods that do not rely on any assumptions about the noise or image. Direct, thresholded, and phase-based deconvolution are all inverse filtering (or deconvolution) techniques. We will refer to methods that do make assumptions about the noise or image as image restoration techniques.

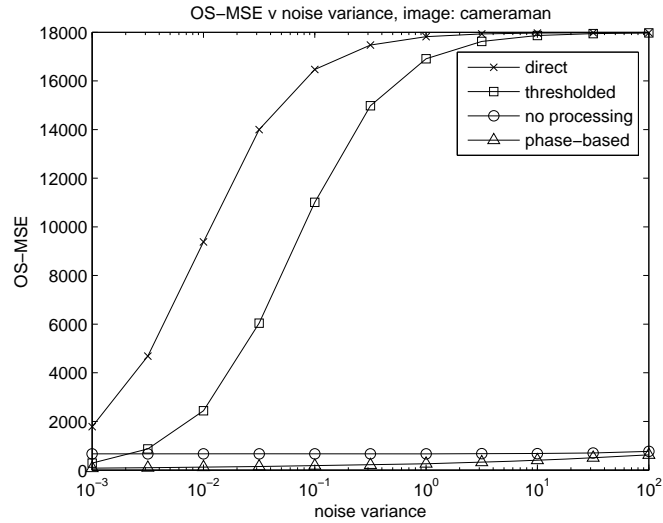


Figure 1. Comparison of inverse filtering methods in the presence of noise.

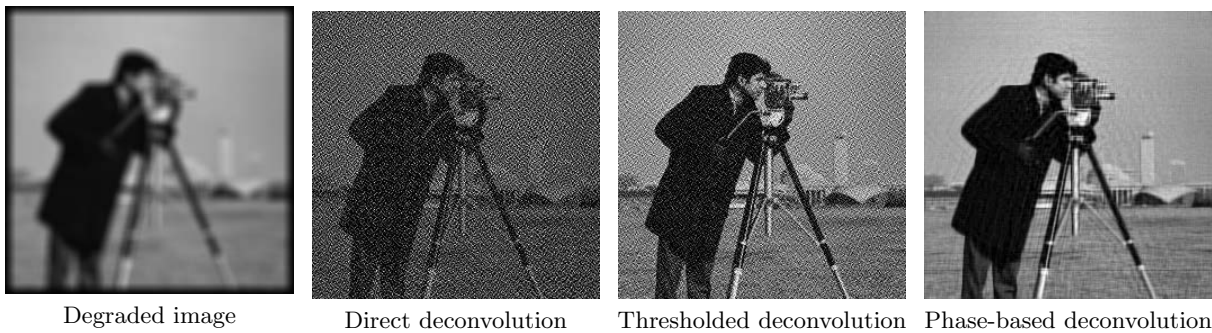


Figure 2. Artifacts in images restored with inverse filtering methods.

To compare the performance of phase-based deconvolution with other deconvolution methods, a variety of images were blurred with an 11x11-point Gaussian filter ($\sigma = 5$) and various levels of white Gaussian noise were added to the blurred images. These images were then deblurred using direct, thresholded, and phase-based deconvolution, and the OS-MSEs between the restored images and the original images were computed.

The results for the cameraman image were typical. The OS-MSE curves for the inverse filtering methods are shown in Fig. 1, and the artifacts in the restored images when the noise variance was 10^{-2} are shown in Fig. 2. In Fig. 2 and all subsequent images, the black borders around the restored and original images are not shown, so, whereas the degraded images are 266x266 pixels, the original and restored images are only 256x256 pixels. As Figs. 1 and 2 show, phase-based deconvolution does significantly better than either direct or thresholded deconvolution.

It is worth noting here that phase-based deconvolution is computationally expensive, with each iteration requiring the computation of a DFT and an inverse DFT of at least twice the size of the signal, and with about a thousand iterations required for convergence. (This varied from one image to the next, but 500 to 1,500 iterations was typical.) However, the gains shown here are certainly worth the computational cost. If phase-based deconvolution were to be used as a step in a full image restoration system, it would be practical only in a setting where computational power was not a major concern.

3. DEMONSTRATION OF ROBUSTNESS

Phase-based deconvolution is so robust to noise that it can be compared to simple image restoration methods that rely on some knowledge of the noise. The results are shown in Fig. 3. The restoration methods used here are

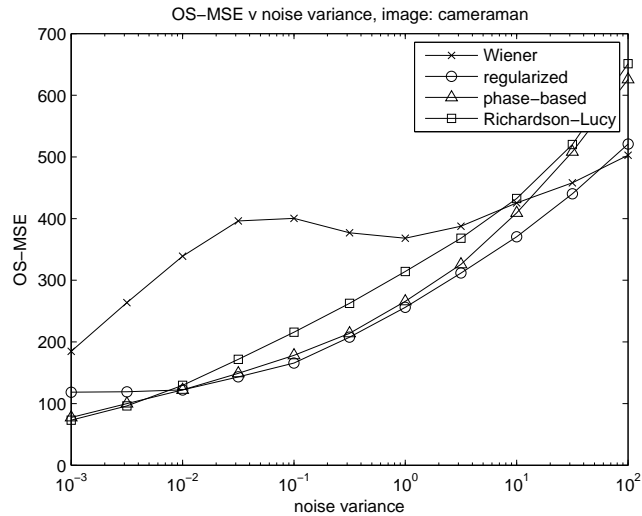


Figure 3. Phase-based deconvolution compared to restoration methods in the presence of noise.



Figure 4. Artifacts in restored images when noise variance was 1.

Wiener filtering, regularized filtering (Wiener filtering with a smoothness constraint), and the Richardson-Lucy iteration. For the purposes of Wiener and regularized filtering, the noise variance was assumed known and the signal to noise ratio was computed from the degraded image. The Richardson-Lucy iteration was stopped at the first local minimum in OS-MSE. In practice, the OS-MSE would not be available and the stopping point would have to be estimated using the noise statistics. Artifacts produced by the various restoration methods are shown in Fig. 4.

As shown in Fig. 3, at low noise levels, phase-based deconvolution performed better than Wiener filtering and was competitive with regularized filtering and the Richardson-Lucy iteration. At high noise levels, it generally performed somewhat worse than the restoration methods. This is expected as those methods attempt to account for (and therefore depend on knowledge of) the noise, whereas phase-based deconvolution does not.

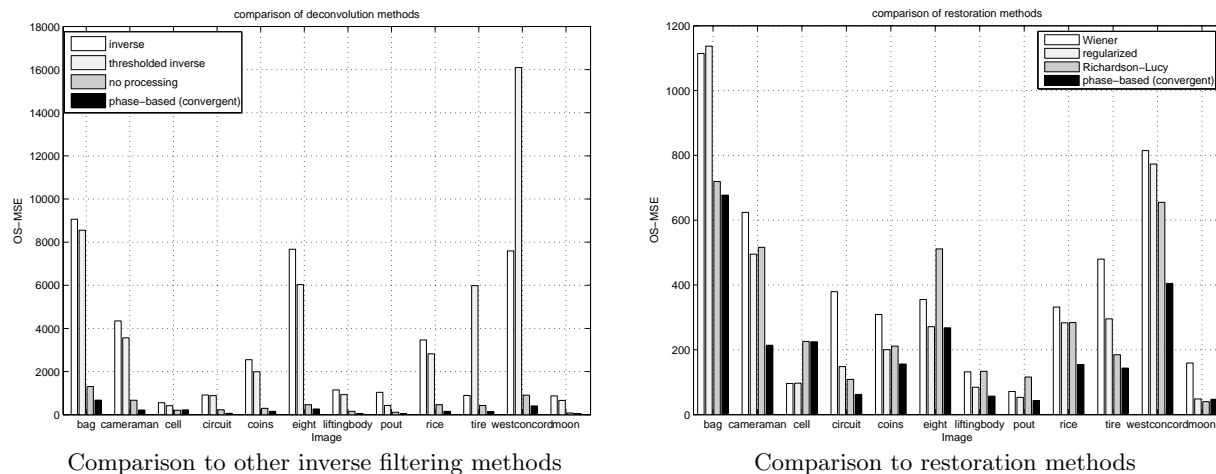


Figure 5. Comparison of methods when knowledge of the PSF was incomplete.

It should be emphasized that this comparison is not meant to imply that phase-based deconvolution should be used as an image restoration system by itself. It is merely a dramatic demonstration of the result that phase-based deconvolution is very robust to noise.

4. INCOMPLETE KNOWLEDGE OF THE PSF

4.1 Known Size

Situations arise in which the PSF is not known completely, if at all. While various methods exist to make some estimate of an unknown PSF, these methods focus on using the statistics of the degraded image and similar images to estimate the magnitude of the PSF.⁶ However, for the large class of low-pass filters that have zero phase, the phase is much easier to estimate. “Zero phase” filters have phases only of 0 or π , and the locations of the phase changes are much more dependent on the size of the filter than on its taper. Therefore, if the size of an unknown PSF can be determined, the phase of a simple averager of the same size may provide an adequate guess of the phase of the true PSF.

Experiments have shown that, when an image is blurred with an 11x11-point Gaussian filter ($\sigma = 5$) and deblurred with an 11x11-point averager, the convergent solution of phase-based deconvolution produced much better results than other inverse filtering methods in all cases. It also generally did better than the restoration methods.

Interestingly, when the PSF was not known completely, the convergent solution was not always the best solution in the OS-MSE sense (though the iteration did always converge). It was sometimes advantageous to stop the reconstruction algorithm before convergence. However, without having the original image, it would be difficult to know when to stop the iteration to achieve an optimal solution, and the convergent solution was also generally acceptable.

These results are illustrated in Fig. 5 for twelve different images. The artifacts present in the restored cameraman images are shown in Fig. 6. The blurring function for these images was an 11x11-point Gaussian filter ($\sigma = 5$), the deblurring function was an 11x11-point averager, and no noise was added. The result for phase-based deconvolution in Fig. 6 is the convergent solution.

When low levels of noise were added to the blurred image, phase-based deconvolution continued to do better than the restoration methods. At high noise levels, the effect of noise began to dominate the effect of the PSF error, and the restoration methods tended to do better than phase-based deconvolution. Phase-based deconvolution always performed much better than any other inverse filtering method. These results are illustrated

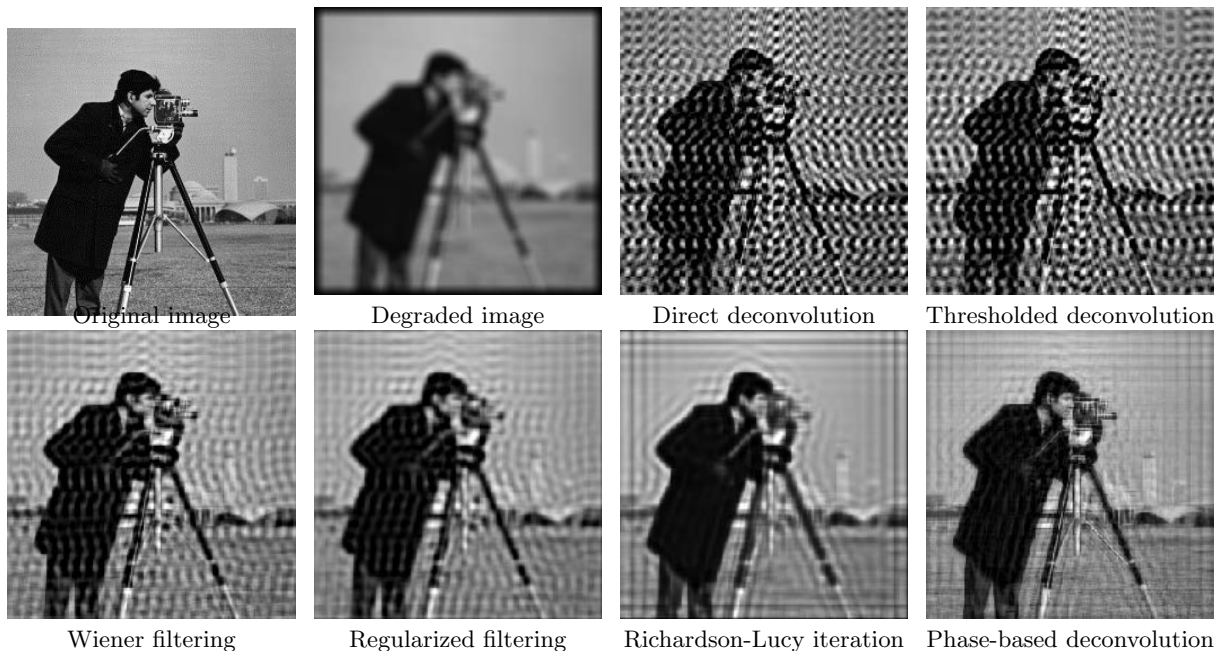


Figure 6. Artifacts in restored images when knowledge of the PSF was incomplete.

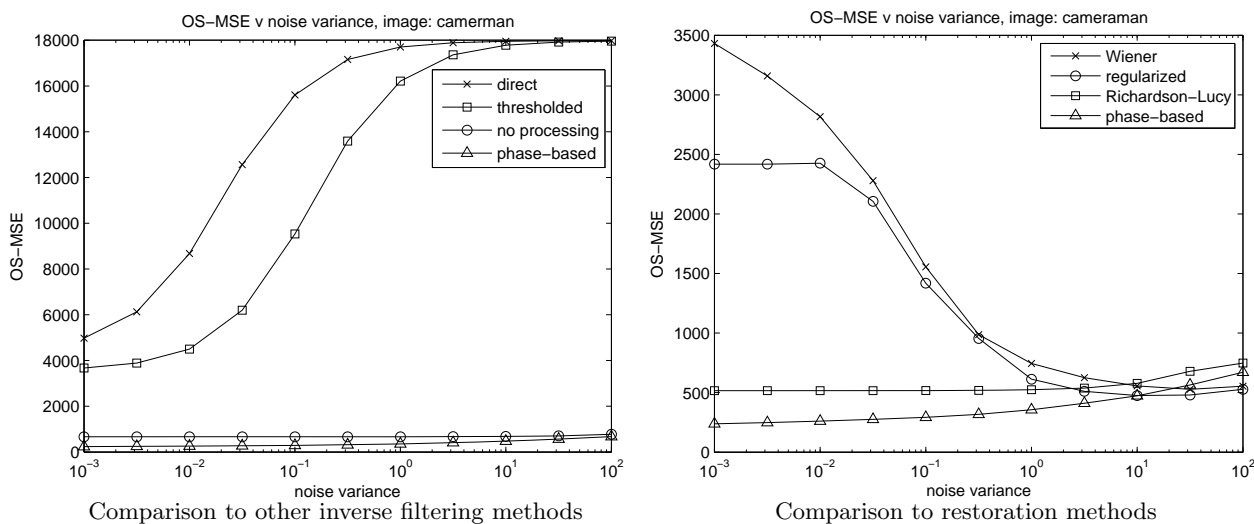


Figure 7. Comparison of methods in the presence of noise and incomplete knowledge of the PSF.

in Fig. 7. Again, the blurring filter was an 11x11-point Gaussian filter ($\sigma = 5$), and the deblurring filter was an 11x11-point averager.

It is interesting to note that Wiener and regularized filtering give better results as the noise increases. This happens because they take the noise statistics as inputs and in this case a larger noise variance is a better description of the error due to the incorrect blurring filter than a smaller noise variance. This demonstrates an advantage of inverse filtering techniques over image restoration techniques: Because inverse filtering techniques do not rely on any noise statistics, there is no penalty for not knowing the noise accurately.

Again, these results do not imply that phase-based deconvolution should be used as an image restoration system. It is simply an inverse filtering technique that is extremely robust to a certain class of errors in the estimate of the PSF.

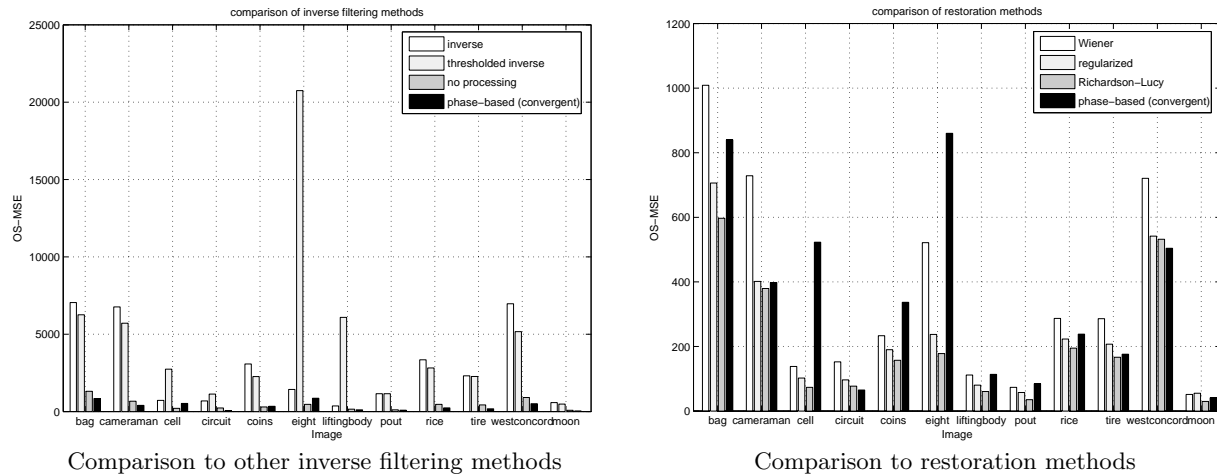


Figure 8. Comparison of methods when knowledge of the PSF was incomplete.

4.2 Unknown Size

To demonstrate the importance of knowledge of the size of the PSF (rather than taper) to phase-based deconvolution, the same experiments were performed with a 9x9-point averager as the estimate of the blurring function, which was still the 11x11-point Gaussian filter used previously. These results are shown in Fig. 8. Comparing Fig. 8 with Fig. 5, it is clear that phase-based deconvolution does not perform nearly as well when the estimated PSF is of the wrong size, though it still performs much better than any other inverse filtering method. Interestingly, in some cases the restoration methods perform better when the estimated PSF was a 9x9-point averager than when it was an 11x11-point averager. This implies that there is some sense in which the 9x9-point averager is a better overall estimate of the 11x11-point Gaussian filter, even though it is a worse estimate of the phase.

5. FINDING THE REGION OF SUPPORT

In the previous two sections, it has been assumed that the ROS of the image was known. However, for the class of pictures where the image is on a large black field (for example, astronomical images), this may not be the case.

Experiments were performed to determine the effect of an incorrect ROS on the result produced by phase-based deconvolution when the PSF was known and there was no noise. These experiments indicated that a ROS estimated too large leads to a much smaller error than a ROS estimated too small. This may be because a ROS estimated too small enforces the zeroing of pixels that should be non-zero, whereas a ROS estimated too large merely allows non-zero values in pixels that should be zero. The frequency domain aspect of the iteration may or may not try to make those pixels non-zero. This should be a guiding principle when looking at methods to determine the ROS—there should be a bias towards methods that overestimate the ROS.

Several such methods have been considered. Results indicate that morphological processing is quite promising. Using fairly low thresholds, morphological processing can generally determine the ROS to within a few pixels even at high noise levels, and using a ROS determined in this way gives similar results to using the correct ROS under the conditions considered in the previous two sections.

6. APPLICATION EXAMPLE

In this section, we discuss an application example. An original image was processed with a low-pass filter in discrete space. The resulting blurred image was quantized and the original was lost. It is now desirable to retrieve the original image from the blurred, quantized image. The blurring function will be assumed known.

Table 1. Number of times each method produced the optimal restoration

Method	coarse	fine	total
Wiener filtering	1	0	1
regularized filtering	14	13	27
Richardson-Lucy iteration	7	17	24
phase-based	20	24	44
thresholded inverse filtering	0	1	1
no processing	2	0	2

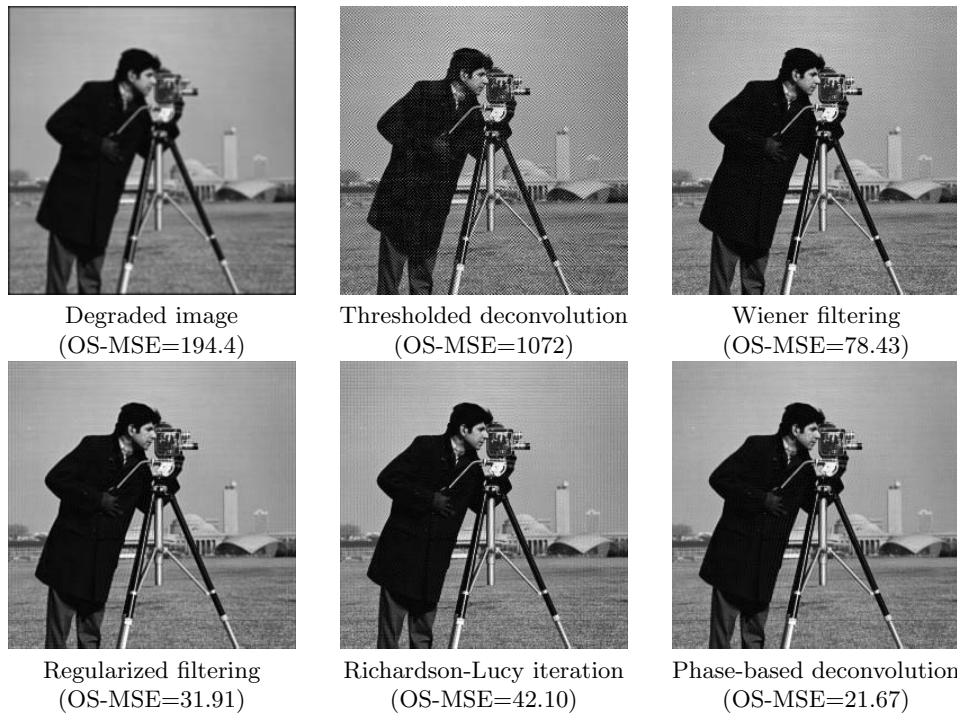


Figure 9. Artifacts in restored images when the blurred image was quantized to 8 bits.

While admittedly somewhat contrived, this example will serve to illustrate the utility of an inverse filtering method that is robust to noise.

Restoration was attempted on eleven different images, with quantization levels ranging from four to twelve bits. The images considered were fairly “typical” black and white images with a variety of characteristics (horizontal and vertical edges, details, flat areas, etc.) Phase-based deconvolution was generally much better than any other inverse filtering method in the OS-MSE sense, and it was generally also competitive with the restoration methods considered thus far. This is illustrated by Table 1, which shows the number of times each restoration method produced the best result in the OS-MSE sense when the blurred images were quantized coarsely (4-7 bits) and finely (8-12 bits).

As can be seen from Table 1, phase-based deconvolution does as well as or better than any other restoration method. This demonstrates the advantage of having a robust inverse filtering method—it does not depend on any description of the noise, and so performs well when the noise is difficult to describe accurately, as quantization noise is.

Additionally, phase-based deconvolution produces different types of artifacts than the other methods. Which types of artifacts are most objectionable will likely depend on application. The restored cameraman images are shown in Fig. 9 to give an example of what these artifacts look like.

7. CONCLUSION

In this paper, we have presented a novel method of two-dimensional inverse filtering using only the phase and region of support of the degraded signal and the phase of the PSF. It has been shown that this method performs much better than other inverse filtering methods in the presence of noise and is even comparable to some basic restoration methods at low noise levels. When only the size of the PSF is known, the phase can be estimated from an averager of that size, and phase-based deconvolution can produce results that are generally better than what is produced by other inverse filtering or restoration methods. The problem of finding the ROS of the image was addressed, and results indicate that morphological processing can perform quite well. An application example was presented which highlighted the utility of an inverse filtering method (as opposed to a restoration method) when the noise was due to quantization of the blurred image.

Finally, in this paper, we considered using only the phase for the image reconstruction. If any reliable magnitude information is available, we note that such information could easily be incorporated in the iterative algorithm.

REFERENCES

- [1] Zunino, A., Benvenuto, F., Armadillo, E., Bertero, M., and Bozzo, E., "Iterative deconvolution and semiblind deconvolution methods in magnetic archaeological prospecting," *GEOPHYSICS* **74**(4), L43–L51 (2009).
- [2] Jirik, R. and Taxt, T., "Two-dimensional blind iterative deconvolution of medical ultrasound images," in [*Ultrasonics Symposium, 2004 IEEE*], **2**, 1262–1265 Vol.2 (Aug 2004).
- [3] Stayman, J. W., Zbijewski, W., Tilley, S., and Siewerdsen, J., "Generalized least-squares ct reconstruction with detector blur and correlated noise models," (2014).
- [4] Hanisch, R. J., White, R. L., and Gilliland, R. L., "Deconvolution of hubble space telescope images and spectra," in [*Deconvolution of Images and Spectra*], Jansson, P. A., ed., **2**, Academic Press (1997).
- [5] Gonzalez, R. and Woods, R. E., [*Digital Image Processing*], Addison-Wesley (1993).
- [6] Lim, J. S., [*Two-Dimensional Signal and Image Processing*], Prentice Hall (1990).
- [7] Hayes, M. H., "The reconstruction of a multidimensional sequence from the phase or magnitude of its Fourier transform," *IEEE Trans. Acoust., Speech, and Signal Proc.* **30**, 140–154 (April 1982).
- [8] Tom, V. T., Quatieri, T. F., Hayes, M. H., and McClellan, J. H., "Convergence of iterative nonexpansive signal reconstruction algorithms," *IEEE Trans. Acoust., Speech, and Signal Proc.* **29** (October 1981).
- [9] Dabov, K., Foi, A., Katkovnik, V., and Egiazarian, K., "Image restoration by sparse 3d transform-domain collaborative filtering," in [*Proc. SPIE*], **6812** (2008).
- [10] Richardson, W. H., "Bayesian-based iterative method of image restoration," *Journal of the Optical Society of America* **62**, 55–59 (January 1972).
- [11] Kalifa, J., Mallat, S., and Rouge, B., "Deconvolution by thresholding in mirror wavelet bases," *IEEE Trans. Image Proc.* **12** (April 2003).
- [12] Wang, Y., Yang, J., Yin, W., and Zhang, Y., "A new alternating minimization algorithm for total variation image reconstruction," *SIAM J. Imaging Sciences* **1**, 248–272 (August 2008).
- [13] Nowak, R. and Figueiredo, M., "Fast wavelet-based image deconvolution using the em algorithm," in [*Proc. of the 35th Conference on Signals, Systems and Computers*], 371–375 (November 2001).
- [14] Dong, W., Zhang, D., and Guangming, S., "Centralized sparse representation for image restoration," in [*Proc. ICCV '11*], 1259–1266 (November 2011).
- [15] Krishnan, D. and Fergus, R., "Fast image deconvolution using hyper-laplacian priors," in [*Advances in Neural Information Processing Systems 22*], Bengio, Y., Schuurmans, D., Lafferty, J., Williams, C., and Culotta, A., eds., 1033–1041, Curran Associates, Inc. (2009).
- [16] Cho, T. S., Zitnick, C. L., Joshi, N., Kang, S. B., and Freeman, W. T., "Image restoration by matching gradient distributions," *IEEE Trans. Pattern Analysis and Machine Intelligence* **34** (April 2012).
- [17] Joshi, N., Zitnick, C., Szeliski, R., and Kriegman, D., "Image deblurring and denoising using color priors," in [*Proc. CVPR '09*], 1550–1557 (June 2009).
- [18] Cho, T., Joshi, N., Zitnick, C., Kang, S., Szeliski, R., and Freeman, W., "A content-aware image prior," in [*Proc. CVPR '10*], 167–176 (June 2010).

- [19] Zoran, D. and Weiss, Y., “From learning models of natural image patches to whole image restoration,” in [*Proc. ICCV '09*], 479–486 (November 2011).
- [20] Dong, W., Zhang, D., Shi, G., and Wu, X., “Image deblurring and super-resolution by adaptive sparse domain selection and adaptive regularization,” *IEEE Trans. Image Proc.* **20** (July 2011).
- [21] Fuller, M. M., *Inverse Filtering by Signal Reconstruction from Phase*, Master’s thesis, Massachusetts Institute of Technology (June 2014).