# WAVELET-DOMAIN REGULARIZED DECONVOLUTION FOR ILL-CONDITIONED SYSTEMS

*Ramesh Neelamani, Hyeokho Choi and Richard Baraniuk*

Department of Electrical and Computer Engineering, Rice University, Houston, TX 77251–1892, USA

## ABSTRACT

We propose a hybrid approach to wavelet-based image deconvolution that comprises Fourier-domain system inversion followed by wavelet-domain noise suppression. In contrast to conventional wavelet-based deconvolution approaches, the algorithm employs a *regularized inverse filter*, which allows it to operate even when the system is non-invertible. Using a mean-square-error metric, we strike an optimal balance between Fourier-domain regularization that is matched to the system and wavelet-domain regularization that is matched to the signal. Theoretical analysis reveals that the optimal balance is determined by economics of the input signal wavelet representation and the operator structure. The resultant algorithm is fast, $O(N \log_2^2 N)$ where $N$ denotes the number of samples, and is well-suited to data with spatially-localized phenomena such as edges. In addition to enjoying asymptotically near-optimal rates of error decay for some systems, the algorithm also achieves excellent performance at fixed data lengths. In simulations with real data, the algorithm outperforms the conventional LTI Wiener filter and other wavelet-based deconvolution algorithms in terms of both visual quality and MSE performance.

## 1. INTRODUCTION

Deconvolution is a recurring theme in a wide variety of signal and image processing problems such as image restoration [1]. For example, satellite images obtained in practice are often blurred due to limitations such as aperture effects of the camera, camera motion, or atmospheric turbulence. Deconvolution becomes necessary if we wish a crisp, deblurred image for viewing or further processing.

In its simplest form, the 1-d deconvolution problem runs as follows. The desired signal $x$ is input to a known linear time-invariant (LTI) system $\mathcal{H}$ having impulse response $h$. Independent identically distributed (i.i.d.) samples of Gaussian noise $\gamma$ with variance $\sigma^2$ corrupt the output samples of the system $\mathcal{H}$. The observations at discrete points $t_n$ are given by $y(t_n) := (x \circledast h)(t_n) + \gamma(n)$, where $n = 0, \ldots, N-1$. Given $y$, we seek to estimate $x$. For simplicity, we assume circular convolution, denoted by $\circledast$. In the discrete Fourier transform (DFT) domain, we equivalently have $Y(f_n) = H(f_n)\,X(f_n) + \Gamma(f_n)$. The $f_n := 2\pi n/N$ denote the normalized frequencies in the DFT domain. The problems formulation trivially extends to multidimensions.

If the system frequency response $H(f_n)$ has no zeros, then we can obtain an unbiased estimate of $x$ as $\widehat{X}(f_n) := H^{-1}(f_n)Y(f_n) = X(f_n) + H^{-1}(f_n)\Gamma(f_n)$. However, if $H(f_n)$ is small at any frequency, then enormous noise amplification results, yielding an infinite-variance, useless estimate.

In situations involving such ill-conditioned systems, some amount of regularization becomes essential. Regularization reduces the variance of the signal estimate (noise reduction) in exchange for an increase in bias (signal distortion). The LTI Wiener filter exploits Fourier domain noise attenuation to estimate the signal from $\widehat{X}(f_n)$. When $x$ is wide-sense stationary (WSS), the LTI *Wiener filter* provides the optimal regularization in the minimum mean-squared-error (MSE) sense [1].

Though the Fourier domain is ideal to represent to the noise colored by $H^{-1}$, it is not appropriate to represent many common signals such as images that contain spatially localized phenomena such as edges because the Fourier basis functions have support that extend over the entire spatial domain. Scalar processing of the Fourier components employed by the LTI Wiener filter results in lack of spatial selectivity, consequently important spatial features such as edges get distorted.

Over the last decade, the wavelet transform has proven to be an invaluable tool for dealing with a wide class of signal including signals with spatially localized features. Wavelets provide economical representations to a large class of signals including many real-world images [2]. This property of wavelets has been leveraged into powerful, spatially adaptive, signal estimation algorithms that are based on simply shrinking the wavelet coefficients of the noisy signal.

Motivated by the ability of wavelets to provide efficient, economical representations to a wide class of signals, Donoho proposed the wavelet-vaguelette decomposition algorithm to solve some special linear inverse problems [3]. With a slight abuse of notation, we will refer to the wavelet-vaguelette decomposition applied to deconvolution as wavelet-vaguelette deconvolution (WVD). In contrast to the Wiener filter, the WVD employs *wavelet-domain denoising* to estimate the signal in the presence of noise colored by $H^{-1}$. In fact, for special classes of blurring opera-

This work was supported by the National Science Foundation, grant no. MIP–9457438, DARPA, grant no. DARPA/AFOSR F49620–97–1–0513, ONR, grant no. N00014–99-1–0813, Texas Instruments, and the Rice Consortium for Computational Seismic/Signal Interpretation.

Email: neelsh@rice.edu, choi@ece.rice.edu, richb@rice.edu

Web: www.dsp.rice.edu

tors such as the Radon transform, WVD possesses asymptotically (as $N \to \infty$) near-optimal rates of error decay for a wide class of input signals as the number of observation samples increases [3].

Though the wavelet transform efficiently represents the input signal, it is not well-suited to represent general convolution operators. Systems with zeroes in the frequency response will result in infinite variance noise in many wavelet coefficients, consequently deeming wavelet domain estimation ineffective.

Motivated by the fact that the Fourier domain matches the convolution operator while the wavelet domain matches the input signal, we propose an improved hybrid wavelet-based regularized deconvolution (WaRD) algorithm suitable for use with any ill-conditioned system. The basic idea is simple: employ *both* Fourier-domain (Wiener-like) regularized inversion and wavelet-domain signal estimation. This tandem processing exploits Fourier-domain regularization adapted to the convolution system to control the noise but uses it sparingly to keep the accompanying smearing distortions to the minimum required. The bulk of the noise removal and signal estimation is achieved using wavelet shrinkage.

By optimizing over an MSE metric,[1] we find the optimal balance between local processing with wavelet basis and global processing with Fourier basis.

Interestingly, one extreme of the balance is to perform no Fourier-domain regularization; this is similar to the WVD approach of Donoho [3] and the mirror wavelet basis of Kalifa et al. [4]. Since WaRD subsumes WVD, the proposed WaRD technique also possesses asymptotically near-optimal error decay rates for special operators such as the Radon transform. In addition, WaRD also improves on the performance of WVD at any fixed samples by choosing the optimal amount of regularization. Further, unlike WVD and the mirror wavelet basis approach, WaRD is applicable to any convolution operator.

Nowak et al. [5] have employed an under-regularized system inverse and subsequently used wavelet domain signal estimation. However they neither studied the implications of using the regularization nor the choice of the optimal amount of regularization. Banham et al. apply a multiscale Kalman filter to the deconvolution problem [6]. The amount of regularization chosen for each wavelet scale is the lower bound that allows for reliable edge classification. While similar in spirit to the multiscale Kalman filter approach, in WaRD, the amount of regularization is chosen to optimize the overall MSE performance of the deconvolved estimate.

## 2. REGULARIZED INVERSE FILTERS

Consider a zero-mean, wide-sense-stationary (WSS) signal $x$ with power spectral density (PSD) $P_x(f)$. Given the general deconvolution problem from the Introduction, a general form for a Fourier-domain-regularized signal estimate is given by

$$\widehat{X}_\alpha(f) \;:=\; G_\alpha(f)\,Y(f) \qquad (1)$$

[1] While the MSE metric does not capture the visual appeal of images in general, we employ it for tractability reasons. In practice, we have found that it yields satisfactory results.

with

$$G_\alpha(f) \;:=\; \left(\frac{1}{H(f)}\right)\left(\frac{|H(f)|^2\,P_x(f)}{|H(f)|^2\,P_x(f) \;+\; \alpha\sigma^2}\right). \qquad (2)$$

The regularization parameter $\alpha$ controls the tradeoff between the amount of noise suppression and the amount of signal distortion. Setting $\alpha = 0$ gives an unbiased but noisy estimate. Setting $\alpha = \infty$ completely suppresses the noise, but also totally distorts the signal ($\widehat{x}_\infty = 0$). For $\alpha = 1$, (2) corresponds to the LTI Wiener filter, which is optimal in the MSE sense for a Gaussian input signal $x$.

The LTI Wiener filter does not provide a good estimate when the input signal comprises of spatially localized phenomena such as edges. The supports of the Fourier basis functions extend over the entire spatial domain. So, scalar operations on the Fourier coefficients lack spatial localization. Hence, all spatial components of the input signal are processed uniformly to result in a substantially distorted estimate.

## 3. WAVELETS AND DECONVOLUTION

The joint time-frequency analysis of the wavelet basis efficiently captures non-stationary signal features. The discrete wavelet transform (DWT) represents a 1-d signal $z$ in terms of shifted versions of a low-pass scaling function $\phi$ and shifted and dilated versions of a prototype bandpass wavelet function $\psi$ [7]. For special choices of $\phi$ and $\psi$, the functions $\psi_{j,k}(t) := 2^{j/2}\,\psi\big(2^j t - k\big)$, $\phi_{j,k}(t) := 2^{j/2}\,\phi\big(2^j t - k\big)$, $j, k \in \mathbb{Z}$ form an orthonormal basis, and we have the representation [7]

$$x(t) \;=\; \sum_k u_{j_0,k}\,\phi_{j_0,k}(t) \;+\; \sum_{j=j_0}^{J}\sum_k w_{j,k}\,\psi_{j,k}(t), \qquad (3)$$

with $u_{j,k} := \int x(t)\,\phi^*_{j,k}(t)\,dt$ and $w_{j,k} := \int x(t)\,\psi^*_{j,k}(t)\,dt$. The parameter $J$ controls the resolution of the wavelet representation. For brevity, we will collectively refer to the set of scaling and wavelet coefficients as $\{\theta_{j,k}\} := \{u_{j_0,k}, w_{j,k}\}$. Multidimensional DWTs are easily obtained by alternately wavelet-transforming along each dimension.

The DWT enjoys an enviable ability to represent many real world signals economically [2]: Most of the energy of many real-world signals is captured by just a few large wavelet coefficients. However, white noise remains disbursed over a large number of small coefficients. This disparity can be exploited to distinguish signal from noise and has given rise to a number of powerful denoising techniques based on simple thresholding [7, 8, 9] that can suppress noise while preserving time-localized signal structures.

Wavelet denoising figures prominently in a number of recent advanced deconvolution algorithms [3, 7, 10]. All three methods have the same two basic steps in common:

**Inversion:**

Compute the noisy estimate $\widehat{X}_0 = H^{-1}(f)\,Y(f)$. This inversion necessarily amplifies noise components at frequencies where $H(f)$ is small.

**Regularization by wavelet denoising:**

Compute the DWT of $\widehat{x}_0$, then denoise using thresholding and finally invert the DWT to obtain the final signal estimate $\widetilde{x}$. Note that the Inversion step colors the white corrupting noise $\gamma$; hence scale-dependent thresholds [9] should be employed.

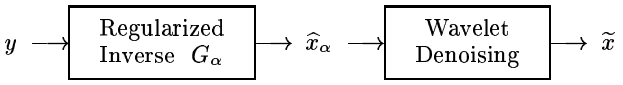


Figure 1: *Wavelet-based regularized deconvolution (WaRD): partially regularized inverse filtering following by wavelet denoising.*

Donoho showed that such an approach possesses asymptotically near-optimal rates of error decay for a wide class of signals (e.g., those in Besov spaces) when the linear operator $\mathcal{H}$ satisfies $h(t) \circledast y(at) = a^\beta\, h(at) \circledast y(t)$, for some exponent $\beta$ e.g., the Radon transform [3].

However, such a wavelet-based approach has its limitations. The noise variance $\sigma_j^2$ at scale $j$ can be approximated by[2]

$$\sigma_j^2 \approx \frac{1}{2^{(j-1)}} \sum_{n=2^{j-1}}^{2^{(j)}-1} \frac{1}{|H(f_n)|^2}\, \sigma^2, \qquad (4)$$

where $H(f_n)$ are the discrete Fourier coefficients of the blurring operator $\mathcal{H}$, and $\sigma^2$ is the variance of the AWGN $\gamma$.

¿From (4), it is clear that even if $H(f_n)$ is small at any $f_n$, the variance of the colored noise in the entire corresponding wavelet scale explodes rendering wavelet denoising ineffective.

## 4. WAVELET-BASED REGULARIZED DECONVOLUTION (WaRD)

To simultaneously exploit the diagonalization of the convolution operator in the Fourier domain and the economical representation of the input signal in the wavelet domain, we propose the wavelet-based regularized deconvolution (WaRD) algorithm. Since a minute amount of regularization (small $\alpha$ in (2)) can lead to a huge reduction in the degree of noise amplification — all this at the expense of only a slight increase in the signal distortion — we simply replace the Inversion step of the algorithms of [3, 7, 10] with a regularized inversion step (see Fig. 1)[3] to obtain the WaRD algorithm.

### 4.1. Optimal regularization for each scale

But how to pick the right value for the regularization parameter $\alpha$? The tradeoff is clear: On one hand, since regularization smears spatially localized signal features (bias), we would prefer $\alpha$ as small as possible. On the other hand, large $\alpha$ prevents excessive noise amplification (variance) during inversion which aids the wavelet denoising.

To be more precise, we will determine the *optimal* regularization parameter for the WaRD system by minimizing the overall MSE. The overall MSE is well-approximated by a cost function $\widetilde{\mathrm{MSE}}(\alpha)$ that includes the distortion error due to the Fourier-domain regularized inversion stage and error incurred during wavelet-domain signal estimation stage:

$$\widetilde{\mathrm{MSE}}(\alpha) := \sum_{n=1}^{N} [1 - G_\alpha(f_n)\, H(f_n)]^2\, P_x(f_n) + \sum_{j,k} \min\left(|\theta_{j,k}(\widehat{x}_\alpha)|^2,\, \sigma_j^2(\alpha)\right). \qquad (5)$$

[2]This approximation is valid only when the zeros in the frequency response $H$ are of sufficiently low order.

[3]Note that the Fourier-domain-regularized inverse (2) was derived for WSS signals $x$ only. With deterministic $x$, we set $P_x(f_n) = |X(f_n)|^2$.

Here $\theta_{j,k}(\widehat{x}_\alpha)$ denotes the wavelet coefficients of $\widehat{x}_\alpha$, $\sigma_j^2(\alpha)$ is the variance of the wavelet-domain noise at scale $j$.

The first term in $\widetilde{\mathrm{MSE}}(\alpha)$ is an estimate of the distortion in the input signal due to the regularized Fourier-domain inverse. This distortion is an increasing function of $\alpha$. The second term is an estimate of the error due to ideal wavelet domain hard thresholding [8]. Ideal thresholding consists of keeping a noisy wavelet coefficient only if the signal power in that coefficient is greater than the noise power. Otherwise, the coefficient is set to zero. Ideal thresholding assumes that the signals under consideration are known. This error is a decreasing function of $\alpha$. The optimal regularization parameter, denoted by $\alpha^\star$, corresponds to the minimum of $\widetilde{\mathrm{MSE}}(\alpha)$.

A useful generalization has different regularization parameters for each wavelet scale. Such a generalization takes advantage of the fact that the cost function is separable with respect to each scale $j$. By minimizing the cost function with respect to $\alpha$, we have shown in [11] that the optimal regularization parameter $\alpha_j^\star$ for the scale $j$ satisfies

$$\alpha_j^\star \approx \frac{1}{2^j} \#\left(\left|\theta_{j,k}(\widehat{x}_{\alpha_j^\star})\right|^2 > \sigma_j^2(\alpha_j^\star)\right), \qquad (6)$$

where $\theta_{j,k}(\widehat{x}_{\alpha_j^\star})$ denotes the wavelet coefficients of the distorted input signal $x_{\alpha^\star}$ at scale $j$, and $\sigma_j^2(\alpha_j^\star)$ is the variance of colored noise at scale $j$. In (6), $\#\left(|\theta_{j,k}(\widehat{x}_{\alpha_j^\star})|^2 > \sigma_j^2(\alpha_j^\star)\right)$ denotes the number of wavelet coefficients $\theta_{j,k}(\widehat{x}_{\alpha_j^\star})$ that have energy greater than the noise variance $\sigma_j^2(\alpha_j^\star)$.

Condition (6) suggests that the noise variance be reduced using regularization so that sufficient number of wavelet coefficients of the signal are greater than the noise variance.

The optimal regularization parameter $\alpha_j^\star$ is never zero. If $\alpha_j = 0$ satisfies (6), it would imply that the noise variance is greater than the energy of each wavelet coefficient. Hence all wavelet coefficients at scale $j$ would be shrunk to zero during wavelet domain estimation. For most real world signals, a significant proportion of the wavelet coefficients is $\approx 0$; so $\alpha_j = 1$ also will not satisfy (6). Hence $\alpha_j = 1$ also is not an optimal choice. We observed that choosing $\alpha_j$ sufficiently greater than zero ($\alpha_j \approx 0.2$) provided estimates comparable to that obtained by choosing the optimal $\alpha$ when the true $|X(f_n)|^2$ was used.

By construction, WaRD includes the value $\alpha = 0$ in the search-space for the optimal $\alpha^\star$. Hence WaRD enjoys all the desirable properties of Donoho's WVD such as near-optimal rate of error decay for special operators such as the Radon transform. The optimal $\alpha^\star$ is never zero at a finite resolution $N$ (though it may approach zero with increasing $N$). Hence, WaRD will outperform wavelet-based deconvolution methods described in [4, 3, 10] in terms of MSE at a given resolution. (assuming knowledge of $|X(f_n)|^2$). Techniques based on WVD [3, 4] are in general not applicable when $\mathcal{H}$ is not invertible. However, thanks to the optimally regularized inversion, WaRD gives excellent estimates even when $\mathcal{H}$ is not invertible.

### 4.2. WaRD implementation

The WaRD algorithm assumes knowledge of the variance $\sigma^2$ of the additive noise $\gamma$ and the Fourier spectrum $|X(f)|^2$

of the input signal. Since these are typically unknown in practice, we estimate them from $y$ using a median estimator [12] and an iterative Wiener technique [13], respectively. However, since the estimation of $|X(f)|^2$ is not robust at frequencies where $H(f_n) \approx 0$, we terminate the iterative Wiener algorithm after 10 iterations and add a small positive constant to the estimate to boost the estimated $|X(f)|^2$ at high frequencies.

Criterion (6) cannot be used in practice to determine the regularization parameters because it assumes knowledge of the distorted wavelet coefficients of the unknown original signal and ideal thresholding. Since the final performance of WaRD is observed to be quite insensitive to changes in the value of the regularization parameter around the optimal value, we estimate a common regularization parameter for all wavelet scales from a plot of the norm of the WaRD estimate with the amount of regularization (see [11] for further details). This is quite accurate in practice.

A variety of wavelet-domain denoising techniques can be employed. A redundant, shift-invariant DWT will yield both a shift-invariant algorithm and improve the denoising performance substantially [7], all at no significant increase to the overall computational cost. Finally, instead of a threshold, we can apply a Wiener filter to the *wavelet* coefficients, as in [14].[4] Such processing has been shown to outperform simple thresholding for denoising finite samples of data.

The overall computational complexity of the algorithm, given the regularization parameter and $|X(f_n)|^2$, is $O(N \log^2 N)$ [11].

## 5. EXAMPLES

We illustrate the performance of the WaRD algorithm using 2-d simulation described by Banham et al. [6] The input $x$ is the $256 \times 256$ Cameraman image and the discrete-time system response $h$ is a $9 \times 9$-point smoother. Such a response is commonly used as a model for blurring due to a square scanning aperture such as in a CCD camera [1]. The blurred signal-to-noise-ratio (BSNR) is defined as $10 \log_{10} \left( \|x \circledast h\|_2^2 / N\sigma^2 \right)$. The noise variance $\sigma^2$ was set so that the BSNR is 40 dB.

Figure 2 illustrates the original image $x$, the blurred and noisy observed image $y$, the Wiener filter estimate using the estimated $|X(f_n)|^2$, and the WaRD estimate with $\alpha^\star = 0.1$ determined empirically as described in [11]. The methods of [3, 4] are not applicable in this situation, due to the many zeros in frequency response of the blurring operator $H$. The visual quality of the Wiener estimate is severely affected by ripples which result because the Fourier basis have support over the entire spatial domain. In contrast to the Wiener estimate, the smooth regions and most edges are simultaneously well-preserved in the WaRD estimate, thanks to the spatially-localized wavelet basis functions. However, some faint features such as the grass are lost during wavelet-domain estimation in WaRD. The WaRD estimate also outperforms the Wiener estimate in terms of MSE. The improvement in the SNR (ISNR) is defined as $10 \log_{10} \left( \|x - y\|_2^2 / \|x - \widehat{x}\|_2^2 \right)$ where $\widehat{x}$ is the estimate. The Wiener filter provides an ISNR of 8.8 dB. In contrast, the proposed WaRD technique provides an ISNR of 10.6 dB. For the same experiment, Banham et al. [6] report an ISNR of 6.68 dB using their the multiscale Kalman filter. The visual quality of the WaRD estimate as compared to the published multiscale Kalman filtering results also seems much improved (see Figure 7(d) in [6]).

[4] Do not confuse wavelet-domain Wiener filtering with the Fourier-domain Wiener filtering discussed above.

## 6. CONCLUSIONS

In this paper, we have proposed an efficient multiscale deconvolution algorithm WaRD that optimally combines Fourier-domain regularized inversion and wavelet-domain signal estimation. The WaRD could be employed in a wide variety of applications, including satellite imaging, to obtain enhanced deconvolution estimates.

For spatially varying signals, WaRD outperforms the LTI Wiener filter and WVD in terms of both visual quality and MSE performance. Since WaRD subsumes WVD, WaRD also enjoys asymptotically near-optimal rates of error decay with increasing samples for convolution operators such as the Radon transform. In addition, WaRD also improves on the performance of the WVD at any fixed resolution. Furthermore, WaRD continues to provide a good estimate of the original signal even with ill-conditioned systems. The computational complexity of the WaRD algorithm, given the regularization parameters and $|X(f_n)|^2$, is just $O(N \log_2^2 N)$, with $N$ the number of samples.

Theoretical analysis of the ideal WaRD algorithm reveals that the optimal regularization parameter at each wavelet scale is determined by the frugality of the wavelet representation of the input signal *and* the Fourier-domain structure of the convolution operator. From (6) it follows that for finite data samples, using $\alpha = 0$ (no regularization) in wavelet-based deconvolution systems such as WVD is never optimal. Further, using a regularization parameter $\alpha = 1$, which corresponds to employing a Wiener deconvolution filter for inversion, is also sub-optimal for many real-world signals.

We have found the final performance of WaRD quite insensitive to changes in the values of the regularization parameters around the optimal values. So, a near-optimal regularization parameter can be obtained from the norm of the WaRD solution for different regularization parameters. As a guide, in simulations spanning many real-world images and convolution systems, $\alpha^\star$ almost always lay in the range $[0.2, 0.3]$ when the true spectrum $|X(f_n)|^2$ was available. However, if the true spectrum is not available, then the optimal regularization parameter becomes dependent on the quality of the spectral estimate and must be determined empirically.

There are several avenues for future WaRD related research. We are currently working on combining WaRD concepts with hidden Markov model (HMM) tree-based wavelet estimation [15] to exploit statistical dependencies between the wavelet coefficients so that the edges and other spatially localized phenomena are preserved better.[5]

[5] We would like to thank R. Nowak for many productive discussions. A big thanks is also due to J. Romberg for help with implementation.

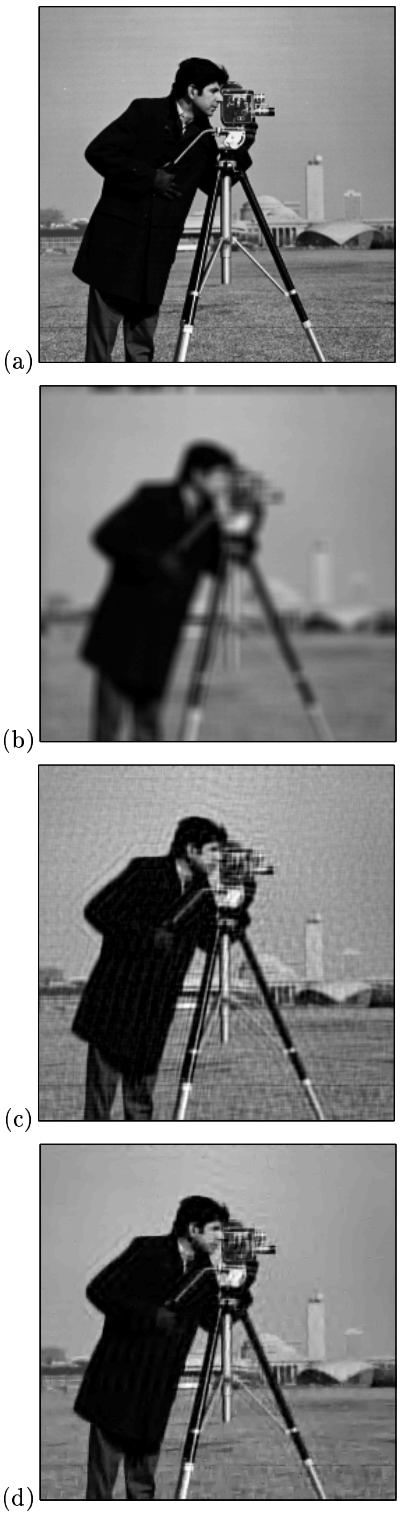

Figure 2: *(a) Cameraman image $x$. (b) Observed image $y$. (c) Wiener filter estimate (ISNR = 8.8 dB). (d) WaRD estimate (ISNR = 10.6 dB).*

## 7. REFERENCES


[1] A. K. Jain, *Fundamentals of Digital Image Processing*, Prentice-Hall, Englewood Cliffs, NJ, 1989.

[2] A. Cohen and J. P. D'Ales, "Nonlinear approximation of random functions," *SIAM J. App. Math.*, vol. 57, no. 2, pp. 518–540, Apr. 1997.

[3] D. L. Donoho, "Nonlinear solution of linear inverse problems by Wavelet-Vaguellete Decomposition," *App. Comp. Harmonic Anal.*, vol. 2, pp. 101–126, 1995.

[4] J. Kalifa, S. Mallat, and B. Rougé, "Image deconvolution in mirror wavelet bases," in *Proc. IEEE Int. Conf. Image Processing — ICIP '98*, Chicago, Oct. 1998, pp. 565–569.

[5] R. D. Nowak and M. J. Thul, "Wavelet-Vaguelette restoration in photon-limited imaging," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Processing — ICASSP '98*, Seattle, WA, 1998, pp. 2869–2872.

[6] M. R. Banham and A. K. Katsaggelos, "Spatially adaptive wavelet-based multiscale image restoration," *IEEE Trans. Image Processing*, vol. 5, pp. 619–634, April 1996.

[7] S. Mallat, *A Wavelet Tour of Signal Processing*, Academic Press, 1998.

[8] D. L. Donoho, "Nonlinear wavelet methods for recovery of signals, densities, and spectra from indirect and noisy data," in *Different Perspectives on Wavelets.* 1993, vol. 47 of *Proc. Symp. Appl. Math.*, pp. 173–205, American Mathematical Society.

[9] I. M. Johnstone and B. W. Silverman, "Wavelet threshold estimators for data with correlated noise," *J. Royal Stat. Soc. B*, , no. 59, pp. 319–351, 1997.

[10] R. D. Nowak, "A fast wavelet-vaguelette algorithm for discrete LSI problems," Tech. Rep., Michigan State University, Aug. 1997.

[11] R. Neelamani, "Wavelet-based deconvolution for ill-conditioned systems," M.S. thesis, Dept. of ECE, Rice Univ., in *www-dsp.rice.edu/publications/*, 1999.

[12] D. L. Donoho and I. Johnstone, "Adapting to unknown smoothness via wavelet shrinkage," *J. Amer. Stat. Assoc.*, vol. 90, pp. 1200–1224, Dec. 1995.

[13] A. D. Hillery and R. T. Chin, "Iterative Wiener filters for image restoration," *IEEE Trans. Signal Processing*, vol. 39, no. 8, pp. 1892–1899, Aug. 1991.

[14] S. Ghael, A. M. Sayeed, and R. G. Baraniuk, "Improved wavelet denoising via empirical Wiener filtering," in *Proc. SPIE Int. Soc. Opt. Eng.*, 1997, vol. 3169, pp. 389–399.

[15] M. S. Crouse, R. D. Nowak, and R. G. Baraniuk, "Wavelet-based statistical signal processing using hidden Markov models," *IEEE Trans. Signal Processing*, vol. 46, no. 4, Apr. 1998, (Special Issue on Wavelets and Filter Banks).