ARTICLE

Ramesh Neelamani and Richard Baraniuk
WInHD: Wavelet-based Inverse Halftoning via Deconvolution
***Rejecta Mathematica*, Vol. 1, No. 1, pp. 84-103, July 2009**

math.rejecta.org

An open letter concerning WInHD: Wavelet-based Inverse Halftoning via Deconvolution

Ramesh Neelamani and Richard Baraniuk

Birth: The niche problem of inverse halftoning error-diffused halftones has been addressed by a number of solid researchers using several practical and effective methods. However, due to the nonlinearity of the halftoning process and the complexities of the human visual system, the methods proposed to date have been ad hoc.

At first glance, we thought that we had little chance of coming up with even a mediocre solution to the nonlinear inverse halftoning problem. We pursued lines of research from photon-limited imaging and Polya trees, but those approaches lead nowhere. One day, Rob Nowak found some literature on an intriguing linear approximation to halftoning. We were pleasantly surprised when a wavelet-thresholding based estimator based on this linear approximation produced competitive results (not only in terms of the workhorse mean-squared-error (MSE) metric but also in terms of a standard visual quality metric). We called our algorithm Wavelet-based Inverse Halftoning via Deconvolution (WInHD).

We thought that WInHD would be a "slam-dunk" paper that would certainly interest the image processing community, because in addition to presenting competitive results near the state-of-the-art, our insights also reduced the inverse halftoning problem to a well-understood deconvolution problem. Furthermore, assuming that the linear approximation was accurate and that the model noise was Gaussian, we were able to derive and analyze bounds on WInHD's MSE performance as the image resolution increased.

With high optimism, we submitted a paper to a top-tier image processing journal.

Death: But alas, our enthusiasm was deflated due to the following review points, which we disagree with.

- The linear approximation was deemed questionable. Any claims about optimality were deemed to be overstated.
- Our results were deemed to be visually inferior. The metrics used used to evaluate our simulation results did not conform to the quality of the images as perceived. We were urged to seek input from the experts in the field and then publish the results of the survey.

The combination of lukewarm reviews and diverging author interests meant that the paper had to be abandoned.

After-life: With its publication in Rejecta Mathematica, we would like to honestly address some of the issues raised in our paper's day of reckoning.

We believe that while the reviewers raised several valid points, the paper contained several contributions that would benefit the image processing community. Addressing the linear approximation point, we agree that a linear approximation to the halftoning process is not suitable for all purposes. However, our view is that the surprising results obtained using such a model make our paper more, not less, interesting. We do concede that the optimality claims made in the paper need to be taken with a this linear approximation in mind. However, the limitations of our analysis

have been clearly stated in the paper (it was termed as conditional optimality in the paper, but perhaps our analysis required some bigger and bolder disclaimers).

On the visual quality issue, beauty indeed lies in the eyes of the beholder! Like a majority of image processing practitioners, we agree that the MSE may be inadequate to measure the visual quality of an image. However, in our paper, we employed all of the metrics that were accessible in the literature (that is, we did not cherry-pick them) to substantiate that our method provided "superior visual" performance (arguably a strong term to use, but certainly not obviously wrong). Surveys can certainly be an effective approach to analyzing an image processing result. But, while useful, conducting surveys for every image processing paper borders on onerous. As an alternative, we published our code so that our results were reproducible and so that our method could be tested on anyone's images of choice.

The tussle about the visual quality improvement afforded by WInHD seems to have no easy resolution in sight. However an even larger question emerges. Is it really necessary for follow-on papers to always significantly improve upon previous results? Should a paper's publishablity be so heavily reliant on the improved results that it produces? How about insights that may open some closed doors?

WInHD: Wavelet-based Inverse Halftoning via Deconvolution

Ramesh Neelamani[*] *and Richard Baraniuk*

Abstract

We propose the *Wavelet-based Inverse Halftoning via Deconvolution* (WInHD) algorithm to perform inverse halftoning of error-diffused halftones. WInHD is motivated by our realization that inverse halftoning can be formulated as a deconvolution problem under Kite et al.'s linear approximation model for error diffusion halftoning. Under the linear model, the error-diffused halftone comprises the original gray-scale image blurred by a convolution operator and colored noise; the convolution operator and noise coloring are determined by the error diffusion technique. WInHD performs inverse halftoning by first inverting the model-specified convolution operator and then attenuating the residual noise using scalar wavelet-domain shrinkage. Since WInHD is model-based, it is easily adapted to different error diffusion halftoning techniques. Using simulations, we verify that WInHD is competitive with state-of-the-art inverse halftoning techniques in the mean-squared-error sense and that it also provides good visual performance. We also derive and analyze bounds on WInHD's mean-squared-error performance as the image resolution increases.

Key words: inverse halftoning, error diffusion, deconvolution, wavelets, wavelet-vaguelette.

[*]Contact author: R. Neelamani. neelsh@gmail.com.

1 Introduction

Digital halftoning is a common technique used to render a sampled gray-scale image using only black or white dots [1] (see Figures 3(a) and (b)); the rendered bi-level image is referred to as a halftone. *Inverse halftoning* is the process of retrieving a gray-scale image from a given halftone. Applications of inverse halftoning include rehalftoning, halftone resizing, halftone tone correction, and facsimile image compression [2, 3]. In this paper, we focus on inverse halftoning images that are halftoned using popular error diffusion techniques such as those of Floyd et al. [4], and Jarvis et al. [5] (hereby referred to as Floyd and Jarvis respectively).

Error-diffused halftoning is non-linear because it uses a quantizer to generate halftones. Recently, Kite et al. proposed an accurate linear approximation model for error diffusion halftoning (see Figure 4) [6, 7]. Under this model, the halftone $y(n_1, n_2)$ is expressed in terms of the original gray-scale image $x(n_1, n_2)$ and additive white noise $\gamma(n_1, n_2)$ as (see Figure 1)

$$\begin{aligned} y(n_1, n_2) &= \mathcal{P}x(n_1, n_2) + \mathcal{Q}\gamma(n_1, n_2) \\ &= (p * x)(n_1, n_2) + (q * \gamma)(n_1, n_2), \end{aligned} \tag{1}$$

with $*$ denoting convolution and (n_1, n_2) indexing the pixels. The $\mathcal{P}$ and $\mathcal{Q}$ are the linear time-invariant (LTI) systems with respective impulse responses $p(n_1, n_2)$ and $q(n_1, n_2)$ determined by the error diffusion technique.

From (1), we infer that inverse halftoning can be posed as the classical *deconvolution* problem because the gray-scale image $x(n_1, n_2)$ can be obtained from the halftone $y(n_1, n_2)$ by deconvolving the filter $\mathcal{P}$ in the presence of the colored noise $\mathcal{Q}\gamma(n_1, n_2)$. Conventionally, deconvolution is performed in the Fourier domain. The Wiener deconvolution filter, for example, would estimate $x(n_1, n_2)$ by inverting $\mathcal{P}$ and *regularizing* the resulting noise with scalar Fourier shrinkage. As we will see, inverse halftoning using a Gaussian low-pass filter (GLPF) [8] can be interpreted as a naive Fourier deconvolution approach to inverse halftoning.

Unfortunately, all Fourier-based deconvolution techniques induce ringing and blurring artifacts due to the fact that the energy of edge discontinuities spreads over many Fourier coefficients. As a result of this uneconomical representation, the desirable edge Fourier coefficients are easily confounded with those due to the noise [9–11].

In contrast, the wavelet transform provides an economical representation for images with sharp edges [12]. This economy makes edge wavelet coefficients easy to distinguish from those due to the noise and has led to powerful image estimation algorithms based on scalar wavelet shrinkage [11, 13].

The wavelet transform was first exploited in inverse halftoning by J. Luo et al. [14]. Xiong et al. extended this algorithm using non-orthogonal, redundant wavelets to obtain improved results for error-diffused halftones [15]. Both these algorithms rely on a variety of steps such as clipping and edge-adapted noise attenuation in the wavelet subbands to exploit different empirical observations. However, these steps and their implications are not theoretically well-justified.

To simultaneously exploit the economy of wavelet representations and the interplay between inverse halftoning and deconvolution, we propose the *Wavelet-based Inverse Halftoning via Deconvolution* (WInHD) algorithm (see Figure 2) [16]. WInHD provides robust estimates by first inverting the convolution operator $\mathcal{P}$ determined by the linear model (1) for error diffusion and then effectively attenuating the residual colored noise using wavelet-domain scalar shrinkage operations [13, 17]. Since WInHD is model-based, it easily adapts to different error diffusion halftoning techniques. See Figure 3 for simulation results.

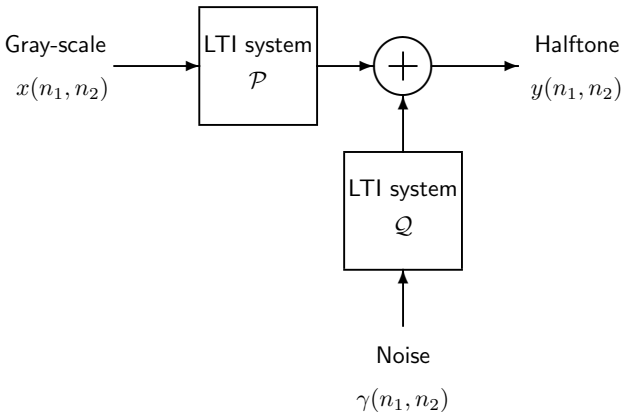

Figure 1: *Linear approximation for error diffusion halftoning. Under the linear model of [6, 7], the error-diffused halftone $y(n_1, n_2)$ comprises the original gray-scale image $x(n_1, n_2)$ passed through an LTI system $\mathcal{P}$ and white noise $\gamma(n_1, n_2)$ colored by an LTI system $\mathcal{Q}$. The systems $\mathcal{P}$ and $\mathcal{Q}$ are determined by the error diffusion technique.*

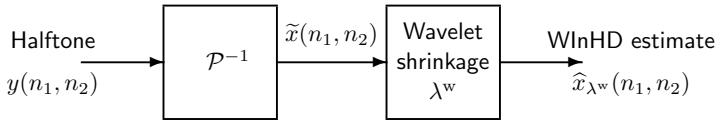

Figure 2: *Wavelet-based Inverse Halftoning via Deconvolution (WInHD). WInHD inverts the convolution operator $\mathcal{P}$ to obtain a noisy estimate $\widetilde{x}(n_1, n_2)$ of the gray-scale image. Subsequent scalar shrinkage with λ^{w} in the wavelet domain (for example, level-dependent hard thresholding) effectively attenuates the residual noise corrupting $\widetilde{x}(n_1, n_2)$ to yield the WInHD estimate $\widehat{x}_{\lambda^{\mathrm{w}}}(n_1, n_2)$.*

Unlike previous inverse halftoning algorithms, we can analyze the theoretical performance of WInHD under certain conditions. For images in a Besov smoothness space, we derive the minimum rate at which the WInHD estimate's mean-squared-error (MSE) decays as the resolution increases; that is, as number of pixels in the gray-scale image tends to infinity. We assume that the linear model for error diffusion (1) is exact and that the noise $\gamma(n_1, n_2)$ is Gaussian. Further, if the gray-scale image $x(n_1, n_2)$ contains some additive noise (say, scanner noise) before halftoning that is Gaussian, then we show that the MSE decay rate enjoyed by WInHD in estimating the noise-free $x(n_1, n_2)$ is optimal; that is, no other inverse halftoning algorithm can have a better error decay rate for every Besov space image as the number of image pixels increases.

Section 2 describes Kite et al.'s linear model for error diffusion halftoning from [6, 7]. Section 3 clarifies the equivalence between inverse halftoning and deconvolution and also analyzes Fourier-domain inverse halftoning. Section 4 presents a brief overview of wavelets. Section 5 discusses the proposed WInHD algorithm and its theoretical performance. Section 6 illustrates the experimental performance of WInHD. Section 7 provides conclusions and future directions. A technical proof in Appendix A completes the paper.

(a) Original $x(n_1, n_2)$ (b) Floyd halftone $y(n_1, n_2)$

(c) Gradient estimate [18] (PSNR = 31.3 dB) (d) WInHD estimate (PSNR = 32.1 dB)

Figure 3: *(a) Original Lena image (512 × 512 pixels). (b) Floyd halftone. (c) Multiscale gradient-based estimate [18], PSNR = 31.3 dB. (d) WInHD yields competitive PSNR performance (32.1 dB) and visual performance. (All documents including the above images undergo halftoning during printing. To minimize the halftoning effect, the images have been reproduced at the maximum size possible.) See Figure 8 for image close-ups.*

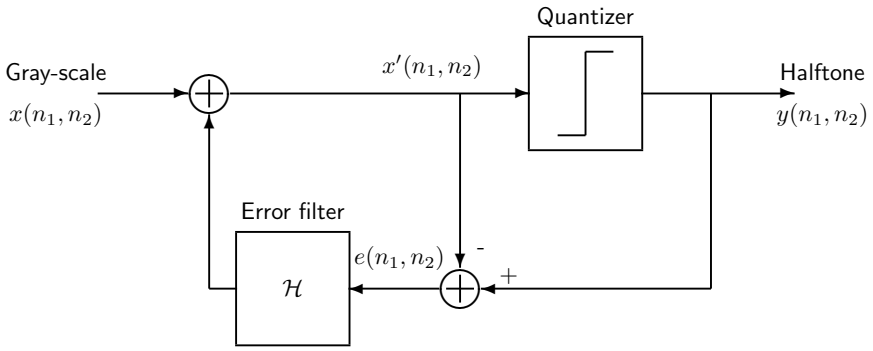

(a) Error diffusion halftoning

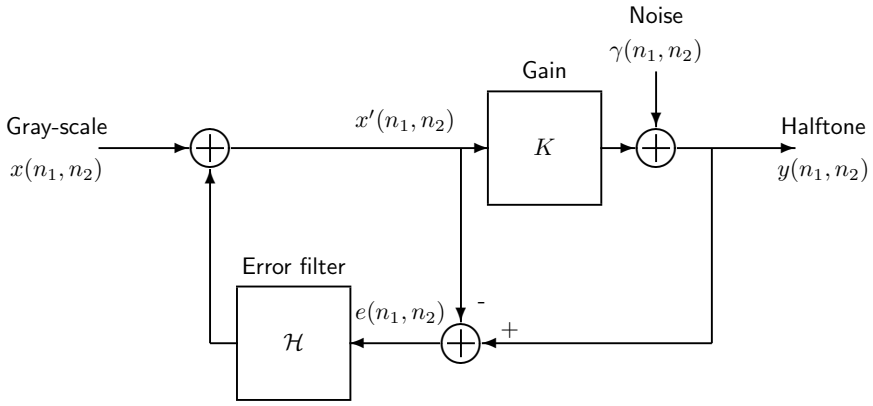

(b) Linear approximation

Figure 4: *(a) Error diffusion halftoning. The gray-scale image pixels $x(n_1, n_2)$ are quantized to yield $y(n_1, n_2)$ and the quantization error $e(n_1, n_2)$ is diffused over a causal neighborhood by the error filter $\mathcal{H}$. (b) The linear model approximates the quantizer with gain K and additive white noise $\gamma(n_1, n_2)$ [6].*

2 Linear Model for Error Diffusion

In this section, we describe the non-linear error diffusion halftoning and the linear approximation proposed in [6, 7].

Digital halftoning is a process that converts a given gray-scale digital image (for example, each pixel value $\in [0, 1, \ldots, 255]$) into a bi-level image (for example, each pixel value = 0 or 255) [1]. Error diffusion halftoning is one popular approach to perform digital halftoning. The idea is to take the error from quantizing a gray-scale pixel to a bi-level pixel and diffuse this quantization error over a causal neighborhood. The error diffusion ensures that the spatially-localized average pixel values of the halftone and original gray-scale image are similar; therefore, the halftone visually resembles the gray-scale image. Figure 4(a) illustrates the block diagram for error diffusion halftoning. The $x(n_1, n_2)$ denote the pixels of the input gray-scale image and $y(n_1, n_2)$ denote the bi-level pixels of the output halftone. The $x'(n_1, n_2)$, which yields $y(n_1, n_2)$ after quantization, is obtained by diffusing the quantization error $e(n_1, n_2)$ over a causal neighborhood of $x(n_1, n_2)$ using the error filter $\mathcal{H}$. The quantizer makes error-diffused halftoning a non-linear technique. Error diffusion

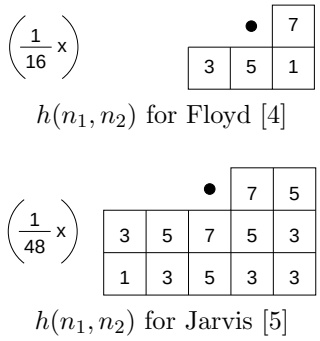

$h(n_1, n_2)$ for Floyd [4]

$h(n_1, n_2)$ for Jarvis [5]

Figure 5: *Error filters $h(n_1, n_2)$ for Floyd [4] and Jarvis [5] error diffusion. The quantization error at the black dot is diffused over a causal neighborhood according the displayed weights.*

techniques such as Floyd [4] and Jarvis [5] are characterized by their choice of $\mathcal{H}$'s impulse response $h(n_1, n_2)$ (see Figure 5).

Recently, Kite et al. proposed an accurate linear model for error diffusion halftoning [6, 7]. This model accurately predicts the "blue noise" (high-frequency noise) and edge sharpening effects observed in various error-diffused halftones. As shown in Figure 4(b), this model approximates the effects of quantization using a gain K followed by the addition of white noise $\gamma(n_1, n_2)$. The halftone $y(n_1, n_2)$ can then be written in terms of the gray-scale image $x(n_1, n_2)$ and the additive white noise $\gamma(n_1, n_2)$ as in (1); the error diffusion technique determines the 2-dimensional (2-D) frequency responses of the LTI filters $\mathcal{P}$ and $\mathcal{Q}$ as

$$P(f_1, f_2) \;:=\; \frac{K}{1 + (K-1)H(f_1, f_2)}, \tag{2}$$

$$Q(f_1, f_2) \;:=\; \frac{1 - H(f_1, f_2)}{1 + (K-1)H(f_1, f_2)} \tag{3}$$

with $P(f_1, f_2)$, $Q(f_1, f_2)$, and $H(f_1, f_2)$ denoting the respective 2-D Fourier transforms of $p(n_1, n_2)$, $q(n_1, n_2)$, and $h(n_1, n_2)$. For any given error diffusion technique, Kite et al. found that the gain K is almost constant for different images. However, the K varied with the error diffusion technique [6]; for example, $K = 2.03$ for Floyd, while $K = 4.45$ for Jarvis. Figure 6 (a) and (b) illustrate the radially-averaged frequency response magnitudes of the filters $\mathcal{P}$ and $\mathcal{Q}$ for Floyd and Jarvis respectively; these responses are obtained by averaging over an annulus of constant radius in the 2-D frequency domain [1]. In [7], Kite et al. further generalized the linear model of (1) by using different gains K_s and K_n in the signal transfer function $P(f_1, f_2)$ and the noise transfer function $Q(f_1, f_2)$ respectively: $P(f_1, f_2) := \frac{K_s}{1+(K_s-1)H(f_1,f_2)}$ and $Q := \frac{1-H(f_1,f_2)}{1+(K_n-1)H(f_1,f_2)}$. In this paper, we assume a single gain factor K for both the signal and noise transfer functions as proposed in [6].

3 Inverse Halftoning ≈ Deconvolution

Given a halftone $y(n_1, n_2)$ (see Figure 4(a)), inverse halftoning aims to estimate the gray-scale image $x(n_1, n_2)$. In the classical deconvolution problem, given the blurred and noisy observation $y(n_1, n_2)$ as in (1) with known LTI filters responses $p(n_1, n_2)$ and $q(n_1, n_2)$, we seek to estimate

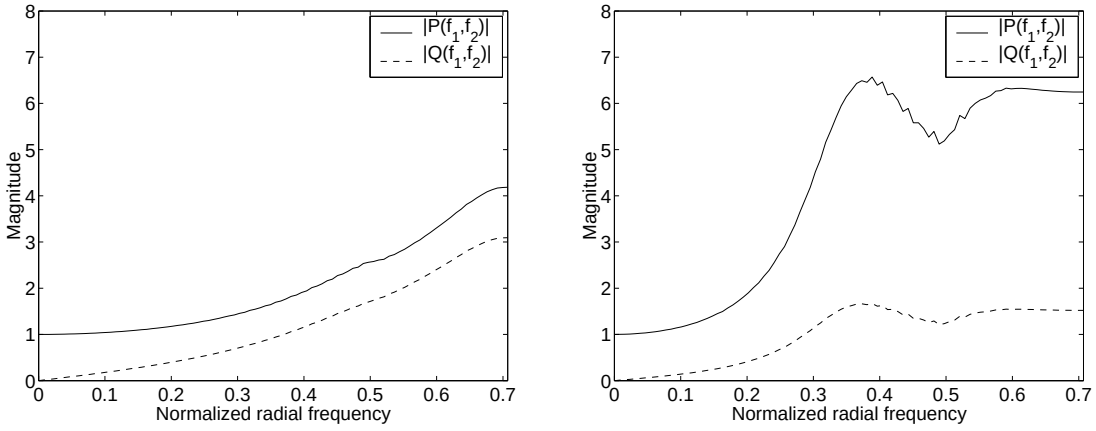

(a) Floyd (b) Jarvis

Figure 6: *Plots (a) and (b) respectively illustrate the radially-averaged frequency response magnitudes $|P(f_1, f_2)|$ (solid line) and $|Q(f_1, f_2)|$ (dotted line) for Floyd and Jarvis. The high-pass $|P(f_1, f_2)|$ explains the sharpened edges, while the high-pass $|Q(f_1, f_2)|$ explains the "blue noise" behavior in the halftones.*

$x(n_1, n_2)$. Thus, under the linear model of [6, 7], inverse halftoning can be posed as a deconvolution problem.

3.1 Deconvolution

Due to the interplay between inverse halftoning and deconvolution, the well-studied deconvolution literature [19–21] can be exploited to understand inverse halftoning as well. Deconvolution algorithms conceptually consist of the following steps:

1. *Operator inversion*
   Invert the convolution operator $\mathcal{P}$ to obtain a noisy estimate $\widetilde{x}(n_1, n_2)$ of the input signal[1]

$$\widetilde{x}(n_1, n_2) := \mathcal{P}^{-1} y(n_1, n_2) = x(n_1, n_2) + \mathcal{P}^{-1}\mathcal{Q}\gamma(n_1, n_2). \tag{4}$$

2. *Transform-domain shrinkage*
   Attenuate the colored noise $\mathcal{P}^{-1}\mathcal{Q}\gamma(n_1, n_2)$ by expressing $\widetilde{x}(n_1, n_2)$ in terms of a chosen orthonormal basis $\{b_k\}_{k=0}^{N-1}$ and shrinking the k-th component with a scalar λ_k, $0 \le \lambda_k \le 1$ [22]

$$\widehat{x}_\lambda := \sum_k \langle \widetilde{x}, b_k \rangle \lambda_k \, b_k = \sum_k \left(\langle x, b_k \rangle + \langle \mathcal{P}^{-1}\mathcal{Q}\gamma, b_k \rangle \right) \lambda_k \, b_k \tag{5}$$

   to obtain the deconvolution estimate $\widehat{x}_\lambda$.

The $\sum_k \langle x, b_k \rangle \lambda_k \, b_k$ in (5) denotes the *retained part* of the signal $x(n_1, n_2)$ that shrinkage preserves from (4), while $\sum_k \langle \mathcal{P}^{-1}\mathcal{Q}\gamma, b_k \rangle \lambda_k \, b_k$ denotes the *leaked part* of the colored noise

[1] For non-invertible $\mathcal{P}$, we replace $\mathcal{P}^{-1}$ by its pseudo-inverse and $x(n_1, n_2)$ by its orthogonal projection onto the range of $\mathcal{P}$ in (4).

$\mathcal{P}^{-1}\mathcal{Q}\gamma(n_1, n_2)$ that shrinkage fails to attenuate. Clearly, we should set $\lambda_k \approx 1$ if the variance $\sigma_k^2 := \mathbb{E}(|\langle \mathcal{P}^{-1}\mathcal{Q}\gamma, b_k\rangle|^2)$ of the k-th colored noise component is small relative to the energy $|\langle x, b_k\rangle|^2$ of the corresponding signal component and set $\lambda_k \approx 0$ otherwise. The shrinkage by λ_k can also be interpreted as a form of *regularization* for the deconvolution inverse problem [20].

The choice of transform domain to perform the shrinkage in deconvolution (see Step 2 above) critically influences the MSE of the deconvolution estimate. An important fact is that for a given transform domain, even with the best possible λ_k's, the estimate $\widehat{x}_\lambda$'s MSE is lower-bounded within a factor of 2 by [9–11]

$$\sum_k \min\left(|\langle x, b_k\rangle|^2, \sigma_k^2\right). \tag{6}$$

From (6), $\widetilde{x}_\lambda$ can have small MSE only when most of the signal energy ($= \sum_k |\langle x, b_k\rangle|^2$) and colored noise energy ($= \sum_k \sigma_k^2$) is captured by just a few transform-domain coefficients — we term such a representation *economical* — and when the energy-capturing coefficients for the signal and noise are different. Otherwise, the $\widetilde{x}_\lambda$ is either excessively noisy due to leaked noise components or distorted due to lost signal components.

Traditionally, the Fourier domain (with sinusoidal b_k) is used to estimate $x(n_1, n_2)$ from $\widetilde{x}(n_1, n_2)$ because it represents the colored noise $\mathcal{P}^{-1}\mathcal{Q}\gamma(n_1, n_2)$ in (4) most economically. That is, among orthonormal transforms, the Fourier transform captures the maximum colored noise energy using a fixed number of coefficients because it diagonalizes convolution operators [23]. Fourier-based deconvolution performs both the operator inversion and the shrinkage simultaneously in the Fourier domain as

$$\widehat{X}_{\lambda^{\mathrm{f}}} := Y(f_1, f_2)\ \frac{1}{P(f_1, f_2)}\ \lambda^{\mathrm{f}}_{f_1, f_2} \tag{7}$$

with shrinkage

$$\lambda^{\mathrm{f}}_{f_1, f_2} := \frac{|P(f_1, f_2)|^2}{|P(f_1, f_2)|^2\ +\ \Upsilon(f_1, f_2)|Q(f_1, f_2)|^2} \tag{8}$$

at the different frequencies. The $Y(f_1, f_2)$ and $\widehat{X}_{\lambda^{\mathrm{f}}}(f_1, f_2)$ denote the 2-D Fourier transforms of $y(n_1, n_2)$ and the deconvolution estimate $\widehat{x}_{\lambda^{\mathrm{f}}}(n_1, n_2)$ respectively. The $\Upsilon(f_1, f_2)$ in (8) is called the *regularization term* and is set appropriately during deconvolution [20]. For example, using the signal to noise ratio to set $\Upsilon(f_1, f_2) = \frac{\mathbb{E}(|\Gamma(f_1, f_2)|^2)}{|X(f_1, f_2)|^2}$ in (7) yields the Wiener deconvolution estimate [24]; the $\Gamma(f_1, f_2)$ and $X(f_1, f_2)$ denote the respective Fourier transforms of $\gamma(n_1, n_2)$ and $x(n_1, n_2)$. The $\frac{1}{P(f_1, f_2)}\ \lambda^{\mathrm{f}}_{f_1, f_2}$ in (7) constitutes the frequency response of the so-called *deconvolution filter*.

Fourier-based deconvolution suffers from the drawback that its estimates for images with sharp edges are unsatisfactory either due to excessive noise or due to distortions such as blurring or ringing. Since the energy due to the edge discontinuities spreads over many image Fourier coefficients, as dictated by the MSE bound in (6), any estimate obtained via Fourier-domain shrinkage suffers from a large MSE.

3.2 Inverse halftoning via Gaussian low-pass filtering (GLPF)

Conventionally, inverse halftoning has been performed using a finite impulse response (FIR) Gaussian filter with coefficients $g(n_1, n_2) \propto \exp[-(n_1^2 + n_2^2)/(2\sigma_g^2)]$, where $-4 \le n_1, n_2 \le 4$, and σ_g determines the bandwidth [8]. We can interpret inverse halftoning using GLPF as a naive Fourier-domain deconvolution approach to inverse halftoning. This is substantiated by our observation that the deconvolution filter $\frac{1}{P(f_1, f_2)}\ \lambda^{\mathrm{f}}_{f_1, f_2}$ (see (7) and (8)) constructed with the linear model filters $\mathcal{P}$

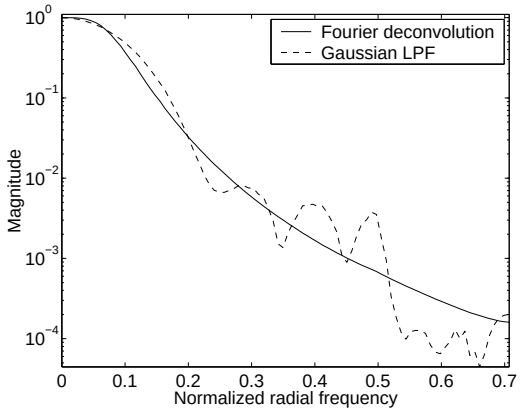

Figure 7: *Comparison of radially-averaged frequency response magnitudes of the FIR GLPF (dashed line) used for inverse halftoning in [8] with the response of the deconvolution filter (solid line) constructed with filters $\mathcal{P}$ and $\mathcal{Q}$ for Floyd and with $\Upsilon(f_1, f_2) \propto \frac{1}{f_1^2+f_2^2}$ (see (7) and (8)). Ripples in the GLPF frequency response result because the filter is truncated in space.*

and $\mathcal{Q}$ for Floyd and with regularization $\Upsilon(f_1, f_2) \propto \frac{1}{f_1^2+f_2^2}$ has a frequency response that closely matches the frequency response of the GLPF (see Figure 7) [8]. The corresponding inverse halftone estimates obtained using simulations are also nearly identical. Predictably, GLPF estimates suffer from the same drawbacks that afflict any Fourier-based deconvolution estimate — excessive noise (when σ_g is small) or significant blurring (when σ_g is large). Exploiting the insights provided by the deconvolution perspective, we can infer that unsatisfactory GLPF estimates result because the Fourier domain does not economically represent images with edges.

4 Background on Wavelets

In contrast to Fourier representations, wavelets provide economical representations for a diverse class of signals including images with edges [11, 12].

4.1 Wavelet transform

The 2-D discrete wavelet transform (DWT) represents a spatially-continuous image $x(t_1, t_2) \in L^2([0,1)^2)$ in terms of shifted versions of a low-pass scaling function ϕ and shifted and dilated versions of prototype bandpass wavelet functions $\{\psi^{LH}, \psi^{HL}, \psi^{HH}\}$ [11, 25]. For special choices of ϕ and ψ's, the shifted and dilated functions $\phi_{j,k_1,k_2}(t_1, t_2) := 2^j\phi(2^j t_1 - k_1, 2^j t_2 - k_2)$, and $\psi^b_{j,k_1,k_2} := 2^j\psi^b(2^j t_1 - k_1, 2^j t_2 - k_2)$ with $b \in \mathcal{B} := \{LH, HL, HH\}$, where the LH, HL, and HH denote the *subbands* of the wavelet decomposition, form an orthonormal basis. The j parameter corresponds to the *scale* of the analysis, while the k_1, k_2 parameters correspond to the *location*. A finite-resolution approximation $x^J(t_1, t_2)$ to $x(t_1, t_2)$ is given by

$$x^J(t_1, t_2) = \sum_{k_1,k_2\in\mathbb{Z}} s_{j_0,k_1,k_2}\phi_{j_0,k_1,k_2}(t_1, t_2) + \sum_{b\in\mathcal{B}}\sum_{j=j_0}^{J}\sum_{k_1,k_2\in\mathbb{Z}} w^b_{j,k_1,k_2}\psi^b_{j,k_1,k_2}(t_1, t_2),$$

with scaling coefficients $s_{j_0,k_1,k_2} := \langle x, \phi_{j_0,k_1,k_2} \rangle$ and wavelet coefficients $w^b_{j,k_1,k_2} := \langle x, \psi^b_{j,k_1,k_2} \rangle$. The parameter J controls the resolution of the wavelet reconstruction $x^J(t_1,t_2)$ of $x(t_1,t_2)$; in fact, the L_2 error $\|x^J - x\|_2 \to 0$ as $J \to \infty$.

The DWT can be extended to transform sampled images as well. Consider, for example, a sampled image obtained by sampling $x(t_1,t_2)$ uniformly as

$$x(n_1,n_2) = N \int_{\frac{n_2}{\sqrt{N}}}^{\frac{n_2+1}{\sqrt{N}}} \int_{\frac{n_1}{\sqrt{N}}}^{\frac{n_1+1}{\sqrt{N}}} x(t_1,t_2)\, dt_1\, dt_2, \qquad 0 \le n_1, n_2 \le \sqrt{N} - 1. \tag{9}$$

For such N-pixel images, the N wavelet coefficients can be efficiently computed in $O(N)$ operations using a filter bank consisting of low-pass filters, high-pass filters, and decimators [11].

Purely for notational convenience, we henceforth refer to the location parameters k_1, k_2 by k and do not explicitly specify the different wavelet subbands: w^b_{j,k_1,k_2} and ψ^b_{j,k_1,k_2} for $b \in \mathcal{B} := \{LH, HL, HH\}$ will be referred to simply as $w_{j,k}$ and $\psi_{j,k}$. Further, we discuss the processing of only the wavelet coefficients, but all steps are replicated on the scaling coefficients as well.

4.2 Economy of wavelet representations

Wavelets provide economical representations for images in smoothness spaces such as Besov spaces [9, 12]. Roughly speaking, a Besov space $B^s_{p,q}$ contains functions with "s derivatives in L_p," with q measuring finer smoothness distinctions [12]. Besov spaces with different s, p, and q characterize many classes of functions; for example, $B^1_{1,\infty}$ contains piece-wise polynomial images [26]. If a continuous-space image $x(t_1,t_2) \in B^s_{p,q}$, $s > \frac{2}{p} - 1$, $1 \le p,q \le \infty$, then the DWT coefficients computed using the image samples (see (9)) satisfies (for all N)

$$\frac{1}{\sqrt{N}} \left(\sum_j 2^{jq(s+1-\frac{2}{p}))} \left(\sum_k |w_{j,k}|^p \right)^{\frac{q}{p}} \right)^{\frac{1}{q}} < \infty, \tag{10}$$

assuming the underlying wavelet basis functions are sufficiently smooth [10, 17, 27].[2] From (10), we can infer that the wavelet coefficients of Besov space images decay exponentially fast with increasing scale j. Further, among all orthogonal transforms, the wavelet transform captures the maximum (within a constant factor) signal energy using a fixed number of coefficients for the worst-case Besov space signal [9].

4.3 Wavelet-based signal estimation

The wavelet transform's economical representations have been exploited in many fields [11]. For example, wavelets provide an effective solution to the problem of estimating signal samples $x(n_1,n_2)$ from additive white Gaussian noise (AWGN) corrupted observations [17, 27, 28]

$$\widetilde{x}(n_1,n_2) = x(n_1,n_2) + \gamma(n_1,n_2), \tag{11}$$

[2]The traditional Besov space characterizing equation in [10, 17, 27] assumes L_2-normalized wavelet coefficients $w_{j,k}$; that is, $\sum_{j,k} |w_{j,k}|^2 = \|x(t_1,t_2)\|^2_2$. Because the $w_{j,k}$ used in (10) are computed using signal samples $x(n_1,n_2)$ that satisfy $\sum_{j,k} |w_{j,k}|^2 = \sum_{n_1,n_2} |x(n_1,n_2)|^2 \approx N \|x(t_1,t_2)\|^2_2$, a normalization factor of $\sqrt{N}$ is required.

with $\gamma(n_1, n_2)$ denoting AWGN samples of variance σ^2. Such a setup is similar to estimating $x(n_1, n_2)$ from (4) but with $\mathcal{P}^{-1}\mathcal{Q}$ equal to identity. Simple shrinkage in the wavelet domain with scalars λ^{w} can provide excellent estimates of $x(n_1, n_2)$. This shrinkage is illustrated by (5) with wavelet basis functions as b_k's and with identity $\mathcal{P}^{-1}\mathcal{Q}$. For example, *hard thresholding* shrinks the wavelet coefficients of $\widetilde{x}(n_1, n_2)$ with scalars

$$\lambda^{\mathrm{w}}_{j,k} \quad = \quad \begin{cases} 1, & \text{if } |\widetilde{w}_{j,k}| > \rho_j \sigma_j, \\ 0, & \text{if } |\widetilde{w}_{j,k}| \leq \rho_j \sigma_j, \end{cases} \tag{12}$$

with $\widetilde{w}_{j,k} := \langle \widetilde{x}\ , \psi_{j,k} \rangle$, σ_j^2 the noise variance at wavelet scale j, and ρ_j the scale-dependent threshold factor (for examples, see [11, p. 442]) . When the pixels $x(n_1, n_2)$ arise from a continuous-space image $x(t_1, t_2) \in B^s_{p,q}$ with $s > \frac{2}{p} - 1$ and $1 \leq p, q \leq \infty$, hard thresholding (with judiciously chosen ρ_j [28]) provides estimates whose MSE-per-pixel decays at least as fast as $N^{\frac{-s}{s+1}}$ as $N \to \infty$ [17, 27]. Further, no estimator can achieve a better error decay rate for every $x(t_1, t_2) \in B^s_{p,q}$. If the threshold factor ρ_j is chosen to be scale-independent, then the MSE decay rate is decelerated by an additional $\log N$ factor.

In practice, the *Wavelet-domain Wiener Filter* (WWF) improves on the MSE performance of hard thresholding by employing Wiener estimation on each wavelet coefficient [29, 30]. WWF chooses

$$\lambda^{\mathrm{w}}_{j,k} = \frac{|w_{j,k}|^2}{|w_{j,k}|^2 + \sigma_j^2}. \tag{13}$$

However, the coefficients $w_{j,k}$ required to construct the $\lambda^{\mathrm{w}}_{j,k}$ are unknown. Hence, a "pilot" estimate of the unknown signal is first computed using hard thresholding. Then, using λ^{w} constructed with the pilot estimate's wavelet coefficients in (13), WWF shrinkage is performed. Sufficiently different wavelet basis functions must be used in the two steps [29, 30].

5 Wavelet-based Inverse Halftoning Via Deconvolution (WInHD)

To simultaneously exploit the economy of wavelet representations and our realization about the interplay between inverse halftoning and deconvolution, we propose the WInHD algorithm [16]. WInHD adopts the wavelet-based deconvolution approach of [10] to perform inverse halftoning.

5.1 WInHD algorithm

WInHD employs scalar shrinkage in the wavelet domain to perform inverse halftoning as follows (see Figure 2):.

1. *Operator inversion*
   As in (4), obtain a noisy estimate $\widetilde{x}(n_1, n_2)$ of the input image by inverting $\mathcal{P}$.

2. *Wavelet-domain shrinkage*
   Employ scalar shrinkage in the wavelet domain to attenuate the noise $\mathcal{P}^{-1}\mathcal{Q}\gamma(n_1, n_2)$ in $\widetilde{x}(n_1, n_2)$ and obtain the WInHD estimate $\widehat{x}_{\lambda^{\mathrm{w}}}(n_1, n_2)$ as follows:

   (a) Compute the DWT of the noisy $\widetilde{x}$ to obtain $\widetilde{w}_{j,k} := \langle \widetilde{x}\ , \psi_{j,k} \rangle$.

(b) Shrink the noisy $\widetilde{w}_{j,k}$ with scalars $\lambda^{\mathrm{w}}_{j,k}$ (using (12) or (13)) to obtain $\widehat{w}_{j,k;\lambda^{\mathrm{w}}} := \widetilde{w}_{j,k}\,\lambda^{\mathrm{w}}_{j,k}$. The colored noise variance at each scale j determining the $\lambda^{\mathrm{w}}_{j,k}$ is given by $\sigma_j^2 := \mathbb{E}\left(\left|\langle \mathcal{P}^{-1}\,\mathcal{Q}\,\gamma\,, \psi_{j,k}\rangle\right|^2\right)$.

(c) Compute the inverse DWT with the shrunk $\widehat{w}_{j,k;\lambda^{\mathrm{w}}}$ to obtain the WInHD estimate $\widehat{x}_{\lambda^{\mathrm{w}}}(n_1, n_2)$.

For error diffusion systems, $\mathcal{P}^{-1}$ is an FIR filter. Hence, the noisy estimate $\widetilde{x}(n_1, n_2)$ obtained in Step 1 using $\mathcal{P}^{-1}$ is well-defined. The subsequent wavelet-domain shrinkage in Step 2 effectively extracts the few dominant wavelet components of the desired gray-scale image $x(n_1, n_2)$ from the noisy $\widetilde{x}(n_1, n_2)$ because the residual noise $\mathcal{P}^{-1}\,\mathcal{Q}\,\gamma(n_1, n_2)$ corrupting the wavelet components is not excessive.

WInHD can be easily adapted to different error diffusion techniques simply by choosing the gain K recommended by [6] and the error filter response $h(n_1, n_2)$ for the target error diffusion technique. K and $h(n_1, n_2)$ determine the filters $\mathcal{P}$ and $\mathcal{Q}$ (see (2) and (3)) required to perform WInHD. In contrast, the gradient-based inverse halftoning method [18] adapts to a given error diffusion technique by employing a set of smoothing filters that need to be designed carefully.

5.2 Asymptotic performance of WInHD

With advances in technology, the spatial resolution of digital images (controlled by the number of pixels N) has been steadily increasing. Hence any inverse halftoning algorithm should not only perform well at a fixed resolution but should also guarantee good performances at higher spatial resolutions. In this section, under some assumed conditions, we deduce the rate at which the per-pixel MSE for WInHD decays as number of pixels $N \to \infty$.

Invoking established results in wavelet-based image estimation in Gaussian noise, we prove the following proposition in Appendix A about the asymptotic performance of WInHD.

Proposition 1 *Let $x(n_1, n_2)$ be a N-pixel gray-scale image obtained as in (9) by uniformly sampling a continuous-space image $x(t_1, t_2) \in B^s_{p,q}$ with $t_1, t_2 \in [0, 1)$, $s > \frac{2}{p} - 1$, and $1 \le p, q, \le \infty$. Let $p(n_1, n_2)$ and $q(n_1, n_2)$ denote known filter impulse responses that are invariant with N and with Fourier transform magnitudes $|P(f_1, f_2)| \ge \epsilon > 0$ and $|Q(f_1, f_2)| < \infty$. Let $y(n_1, n_2)$ be observations obtained as in (1) with $\gamma(n_1, n_2)$ zero-mean AWGN samples with variance σ^2. Then, the per-pixel MSE of the WInHD estimate $\widehat{x}(n_1, n_2)$ obtained from $y(n_1, n_2)$ using hard thresholding behaves as*

$$\frac{1}{N}\mathbb{E}\left(\sum_{n_1,n_2} |\widehat{x}(n_1, n_2) - x(n_1, n_2)|^2\right) \le C\, N^{\frac{-s}{s+1}}, \qquad N \to \infty, \tag{14}$$

with constant $C > 0$ independent of N.

The above proposition affirms that the per-pixel MSE of the WInHD estimate decays as $N^{\frac{-s}{s+1}}$ with increasing spatial resolution ($N \to \infty$) under the mild assumptions discussed below.

The central assumption in Proposition 1 is that the linear model (1) for error diffusion is accurate. This is well-substantiated in [6, 7]. The conditions $|P(f_1, f_2)| \ge \epsilon > 0$ and $|Q(f_1, f_2)| < \infty$ respectively ensure that $\mathcal{P}$ is invertible and that the variance of the colored noise $Q\gamma(n_1, n_2)$ is bounded. We have verified that for common error diffusion halftoning techniques such as Floyd and Jarvis, the filters $\mathcal{P}$ and $\mathcal{Q}$ recommended by the linear model of Kite et al. satisfy these

(a) Floyd halftone (b) Gradient estimate [18] (c)WInHD estimate

Figure 8: *Close-ups (128 × 128 pixels) of (a) Floyd halftone, (b) Gradient estimate, and (c) WInHD estimate.*

conditions (see Figure 6). The final assumption is that the noise $\gamma(n_1, n_2)$ is Gaussian; this is required to invoke the established results on the asymptotics of wavelet-based estimators [17]. However, recently, wavelet-domain thresholding has been shown to be optimal for many other noise distributions as well [31, 32]. Hence the noise Gaussianity assumption in Proposition 1 could be relaxed.

Often, gray-scale digital images are corrupted with some noise before being subjected to halftoning. For example, sensor noise corrupts images captured by charged coupled device (CCD) based digital cameras. In such cases as well, WInHD can effectively estimate the noise-free gray-scale image with an MSE decay rate of $N^{\frac{-s}{s+1}}$ as in Proposition 1. Further, WInHD's MSE decay rate can be shown to be optimal. The noise-free gray-scale image and resulting halftone can be related using the linear model of [6, 7] as

$$y(n_1, n_2) = \mathcal{P}\left[x(n_1, n_2) + \beta(n_1, n_2)\right] + \mathcal{Q}\gamma(n_1, n_2), \tag{15}$$

with $\beta(n_1, n_2)$ denoting the noise corrupting the gray-scale image before halftoning. If the $\beta(n_1, n_2)$ is AWGN with non-zero variance, then we can easily infer that the residual noise after inverting $\mathcal{P}$ in Step 1 of WInHD can be analyzed like white noise because its variance is bounded but non-zero [10]. Hence we can invoke well-established results on the performance of wavelet-based signal estimation in the presence of white noise [17, 27, 28] to conclude that no estimator can achieve a better error decay rate than WInHD for every Besov space image. Thus, WInHD is an optimal estimator for inverse halftoning error-diffused halftones of noisy images.

6 Results

We illustrate WInHD's performance using 512×512-pixel *Lena* and *Peppers* test images halftoned using the Floyd algorithm [4] (see Figure 3 and 8). All WInHD estimates and software are available at `www.dsp.rice.edu/software`. We set the gain $K = 2.03$, as calculated for Floyd in [6, 7], and use the Floyd error filter response $h(n_1, n_2)$ (see Figure 5) to characterize the impulse responses $p(n_1, n_2)$ and $q(n_1, n_2)$. Inverting the operator $\mathcal{P}$ (Step 2) requires $O(N)$ operations and memory

Table 1: *PSNR and computational complexity of inverse halftoning algorithms (N pixels).*

Inverse halftoning algorithm	PSNR (dB)		Computational complexity
	Lena	*Peppers*	
Gaussian [8]	28.6	27.6	$O(N)$
Kernel [2]	32.0	30.2	$O(N)$
Gradient [18]	31.3	31.4	$O(N)$
Wavelet denoising [15]	31.7	30.7	$O(N \log N)$
WInHD	32.1	31.2	$O(N)$

for a N-pixel image since $\mathcal{P}^{-1}$ is FIR. To perform the wavelet-domain shrinkage (Step 2), we choose the WWF because it yields better estimates compared to schemes such as hard thresholding.

Estimates obtained by shrinking DWT coefficients are not shift-invariant; that is, translations of $y(n_1, n_2)$ will result in different estimates. Hence, we exploit the *complex wavelet transform* (CWT) instead of the usual DWT to perform the WWF. The CWT expands images in terms of shifted and dilated versions of *complex-valued* basis functions instead of the real-valued basis functions used by the DWT [33, 34]; the expansion coefficients are also complex-valued. Wavelet-domain shrinkage using WWF on the CWT coefficient magnitudes yields significantly improved near shift-invariant estimates with just $O(N)$ operations and memory. (The redundant, shift-invariant DWT can also be used instead of the CWT to obtain shift-invariant estimates [11], but the resulting WInHD algorithm requires $O(N \log N)$ operations and memory.) The standard deviation of the noise $\gamma(n_1, n_2)$, which is required during wavelet shrinkage, is calculated using the standard deviation of $y(n_1, n_2)$'s finest scale CWT coefficients.

Figures 3 and 8 compares the WInHD estimate with the multiscale gradient-based estimate [18] for the Lena image. We quantify the WInHD's performance by measuring the peak signal-to-noise ratio PSNR $:= 20 \log_{10} \frac{512 \times 255}{\|\widehat{x} - x\|_2}$ (for 512×512-pixel images with gray levels $\in [0, 1, \ldots, 255]$) with $\widehat{x}(n_1, n_2)$ the estimate. Table 1 summarizes the PSNR performance and computational complexity of WInHD compared to published results for inverse halftoning with Gaussian filtering [8], kernel estimation [2], gradient estimation [18], and wavelet denoising with edge-detection [15]. We can see that WInHD is competitive with the best published results.

The WInHD estimate yields competitive visual performance as well. We quantify visual performance using two metrics: weighted SNR (WSNR) [35, 36] and the *Universal Image Quality Index* (UIQI) [37]. Both metrics were computed using the halftoning toolbox of [38]. The WSNR is obtained by weighting the SNR in the frequency domain according to a linear model of the human visual system [35, 36]. The WSNR numbers in Table 2 are calculated at a spatial Nyquist frequency of 60 cycles/degree. The recently proposed UIQI metric of Wang et al. effectively models image distortion with a combination of correlation loss, luminance distortion, and contrast distortion [37]; UIQI $\in [-1, 1]$ with larger values implying better image quality. For the Lena image, WInHD's performance in terms of both the visual metrics is competitive with the gradient estimate's performance (see Table 2).

Table 2: *Visual metrics for inverse halftoned estimates of Lena.*

Algorithm	WSNR (dB)	UIQI
Gradient [18]	34.0	0.62
WInHD	35.9	0.62

7 Conclusions

Using the linear error diffusion model of [6, 7], we have demonstrated that inverse halftoning can be posed as a deconvolution problem in the presence of colored noise. Exploiting this new perspective, we have proposed the simple *Wavelet-based Inverse Halftoning via Deconvolution* (WInHD) algorithm based on wavelet-based deconvolution to perform inverse halftoning. Since WInHD is model-based, it is easily tunable to the different error diffusion halftoning techniques. WInHD yields state-of-the-art performance in the MSE sense and visually.

WInHD also enjoys desirable theoretical properties under certain mild conditions. For images in a Besov space, WInHD estimate's MSE is guaranteed to decay rapidly as the spatial resolution of the input gray-scale image increases. Further, if the gray-scale image lies in a Besov space and is noisy before halftoning, then WInHD's MSE decay rate cannot be improved upon by any estimator.

We have assumed *a priori* knowledge of the error diffusion filter in this paper. However, the error diffusion filter is not always known. Under such circumstances, the error diffusion filter coefficients could be estimated by integrating adaptive techniques such as the one proposed by Wong [39] into our algorithm. However, this remains a topic of future study.

To facilitate efficient hardware implementation, in addition to requiring minimal memory and computations, an inverse halftoning algorithm should also be compatible with fixed-point digital signal processors. For example, the gradient-based algorithm [18] is optimized for hardware implementation while still obtaining good inverse halftoning results. Since our focus in this paper has been primarily theoretical, we have not specifically addressed any hardware optimization issues. The design of a hardware-compatible inverse halftoning algorithm based on WInHD is a topic of interesting future study.

A Decay Rate of WInHD's MSE

We deduce the asymptotic performance of WInHD as claimed in Proposition 1.

Instead of analyzing the problem of estimating $x(n_1, n_2)$ from $y(n_1, n_2)$, we can equivalently analyze the estimation of $x(n_1, n_2)$ from the noisy observation $\widetilde{x}(n_1, n_2)$ obtained after inverting $\mathcal{P}$ (see (4)). The reduction is equivalent because $P(f_1, f_2)$ is known and invertible (since $|P(f_1, f_2)| \geq \epsilon > 0$).[3]

The frequency components of the colored noise $\mathcal{P}^{-1}\mathcal{Q}\gamma(n_1, n_2)$ corrupting the $\widetilde{x}(n_1, n_2)$ in (4) is given by $\frac{Q(f_1,f_2)\Gamma(f_1,f_2)}{P(f_1,f_2)}$. These frequency components are independent and Gaussian because the Fourier transform diagonalizes convolution operators. Since $|P(f_1, f_2)|$ is strictly non-zero and $|Q(f_1, f_2)|$ is bounded, the variance of $\frac{Q(f_1,f_2)\Gamma(f_1,f_2)}{P(f_1,f_2)}$ is uniformly bounded — say with variance ς^2

[3]Since the filter $\mathcal{P}^{-1}$ is FIR for error diffusion systems, boundary effects are negligible asymptotically because only a finite number of boundary pixels are corrupted.

— at all frequencies.

Because the estimation error due to wavelet-domain hard thresholding is monotone with respect to noise variance [10], the error in estimating $x(n_1, n_2)$ from (4) using wavelet-domain hard thresholding is less than the error in estimating $x(n_1, n_2)$ observed in white noise as in (11) but with variance ς^2. Hence the per-pixel MSE in estimating $x(n_1, n_2)$ from (4) can be bounded with the decay rate $N^{\frac{-s}{s+1}}$ established for the white noise setup (see Section 4.3) to yield (14) with a constant $C > 0$ independent of N [27, 28]. □

Acknowledgments

We thank Dr. Brian Evans for his many constructive comments.

References

[1] R. Ulichney, *Digital Halftoning*. Cambridge, MA: MIT Press, 1987.

[2] P. Wong, "Inverse halftoning and kernel estimation for error diffusion," *IEEE Trans. Image Processing*, vol. 6, pp. 486–498, Apr. 1995.

[3] M. Y. Ting and E. A. Riskin, "Error-diffused image compression using a binary-to-gray scale decoder and predictive pruned tree-structured vector quantization," *IEEE Trans. Image Processing*, vol. 3, pp. 854–857, Nov. 1994.

[4] R. W. Floyd and L. Stienberg, "An adaptive algoritm for spatial grayscale," *Proc. Soc. Image Display*, vol. 17, no. 2, pp. 75–77, 1976.

[5] J. Jarvis, C. Judice, and W. Ninke, "A survey of techniques for the display of continuous tone pictures on bilevel displays," *Comput. Graph and Image Process.*, vol. 5, pp. 13–40, 1976.

[6] T. D. Kite, B. L. Evans, A. C. Bovik, and T. L. Sculley, "Digital halftoning as 2-D delta-sigma modulation," *Proc. IEEE ICIP '97*, vol. 1, pp. 799–802, Oct. 26–29 1997.

[7] T. D. Kite, B. L. Evans, A. C. Bovik, and T. L. Sculley, "Modeling and quality assessment of halftoning by error diffusion," *IEEE Trans. Image Processing*, vol. 9, pp. 909–922, May 2000.

[8] S. Hein and A. Zakhor, "Halftone to continuous-tone conversion of error-diffusion coded images," *IEEE Trans. Image Processing*, vol. 4, pp. 208–216, Feb. 1995.

[9] D. L. Donoho, "Unconditional bases are optimal bases for data compression and for statistical estimation," *Appl. Comput. Harmon. Anal.*, vol. 1, pp. 100–115, Dec. 1993.

[10] D. L. Donoho, "Nonlinear solution of linear inverse problems by Wavelet-Vaguelette Decomposition," *Appl. Comput. Harmon. Anal.*, vol. 2, pp. 101–126, 1995.

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

[12] R. A. DeVore, B. Jawerth, and B. J. Lucier, "Image compression through wavelet transform coding," *IEEE Trans. Inform. Theory*, vol. 38, pp. 719–746, Mar. 1992.

[13] D. L. Donoho and I. M. Johnstone, "Ideal spatial adaptation via wavelet shrinkage," *Biometrika*, vol. 81, pp. 425–455, 1994.

[14] J. Luo, R. de Queiroz, and Z. Fan, "A robust technique for image descreening based on the wavelet transform," *IEEE Trans. Signal Processing*, vol. 46, pp. 1179–1184, Apr. 1998.

[15] Z. Xiong, M. T. Orchard, and K. Ramchandran, "Inverse halftoning using wavelets," *IEEE Trans. Signal Processing*, vol. 8, pp. 1479–1482, Oct. 1999.

[16] R. Neelamani, R. D. Nowak, and R. G. Baraniuk, "Model-based inverse halftoning with Wavelet-Vaguelette Deconvolution," in *Proc. IEEE ICIP '00*, (Vancouver, Canada), pp. 973–976, Sept. 2000.

[17] D. L. Donoho, "De-noising by soft-thresholding," *IEEE Trans. Inform. Theory*, vol. 41, pp. 613–627, May 1995.

[18] T. D. Kite, N. Damera-Venkata, B. L. Evans, and A. C. Bovik, "A fast, high-quality inverse halftoning algorithm for error diffused halftones," *IEEE Trans. Image Processing*, vol. 9, pp. 1583–1592, Sept. 2000.

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

[20] A. K. Katsaggelos (Ed.), *Digital Image Restoration*. New York: Springer-Verlag, 1991.

[21] R. Neelamani, H. Choi, and R. G. Baraniuk, "Wavelet-based deconvolution using optimally regularized inversion for ill-conditioned systems," in *Wavelet Applications in Signal and Image Processing VII, Proc. SPIE*, vol. 3813, pp. 58–72, July 1999.

[22] W. James and C. Stein, "Estimation with quadratic loss," in *Proc. Fourth Berkeley Symp. Math. Statist. Probab.*, vol. 1, pp. 361–380, Univ. California Press, 1961.

[23] G. Davis and A. Nosratinia, "Wavelet-based image coding: An overview," in *Appl. Comput. Control Signals Circuits* (B. N. Datta, ed.), vol. 1, Birkhauser, 1999.

[24] K. R. Castleman, *Digital Image Processing*. New Jersey: Prentice Hall, 1996.

[25] C. S. Burrus, R. A. Gopinath, and H. Guo, *Introduction to Wavelets and Wavelet Transforms: A Primer*. Prentice-Hall, 1998.

[26] J. Kalifa and S. Mallat, "Thresholding estimators for linear inverse problems," *Ann. Statist.*, vol. 31, Feb. 2003.

[27] D. L. Donoho and I. M. Johnstone, "Asymptotic minimaxity of wavelet estimators with sampled data," *Statist. Sinica*, vol. 9, no. 1, pp. 1–32, 1999.

[28] D. L. Donoho and I. Johnstone, "Minimax estimation by wavelet shrinkage," *Ann. Statist.*, vol. 26, pp. 879–921, 1998.

[29] S. Ghael, A. M. Sayeed, and R. G. Baraniuk, "Improved wavelet denoising via empirical Wiener filtering," in *Wavelet Applications in Signal and Image Processing V, Proc. SPIE*, vol. 3169, pp. 389–399, Oct. 1997.

[30] H. Choi and R. G. Baraniuk, "Analysis of wavelet domain Wiener filters," in *IEEE Int. Symp. Time-Frequency and Time-Scale Analysis*, (Pittsburgh), Oct. 1998.

[31] R. Averkamp and C. Houdre, "Wavelet thresholding for non (necessarily) Gaussian noise: Functionality," *Ann. Statist.*, vol. 31, Feb. 2003.

[32] H.-Y. Gao, "Choice of thresholds for wavelet shrinkage estimate of the spectrum," *Journal of Time Series Analysis*, vol. 18, pp. 231–251, 1997.

[33] N. G. Kingsbury, "Complex wavelets for shift invariant analysis and filtering of signals," *Appl. Comput. Harmon. Anal.*, vol. 10, pp. 234–253, May 2001.

[34] J. K. Romberg, H. Choi, R. G. Baraniuk, and N. G. Kingsbury, "A hidden Markov tree model for the complex wavelet transform," *IEEE Trans. Signal Processing*, 2002. Submitted.

[35] T. Mitsa and K. Varkur, "Evaluation of contrast sensitivity functions for the formulation of quality measures incorporated in halftoning algorithms," in *Proc. IEEE ICASSP '93*, vol. 5, pp. 301–304, 1993.

[36] T. D. Kite, N. Damera-Venkata, B. L. Evans, and A. C. Bovik, "Image quality assessment based on a degradation model," *IEEE Trans. Image Processing*, vol. 9, pp. 636–650, Apr. 2000.

[37] Z. Wang and A. C. Bovik, "A universal image quality index," *IEEE Signal Processing Lett.*, vol. 9, pp. 81–84, Mar. 2002.

[38] V. Monga, N. Damera-Venkata, and B. L. Evans., *Halftoning Toolbox for MATLAB*. 2002. Available: www.ece.utexas.edu/~bevans/projects/halftoning/toolbox.

[39] P. W. Wong, "Inverse halftoning and kernel estimation for error diffusion," *IEEE Trans. Image Processing*, vol. 4, pp. 486–498, Apr. 1995.