

PERCEPTUAL SOFT THRESHOLDING USING THE STRUCTURAL SIMILARITY INDEX

Sumohana S. Channappayya, Alan C. Bovik, and Robert W. Heath Jr.

The University of Texas at Austin
Department of Electrical and Computer Engineering
1 University Station C0803, Austin TX - 78712-0240 USA

ABSTRACT

In this paper, we present a novel algorithm for wavelet domain image denoising using the soft thresholding function. The thresholds are designed to be locally optimal with respect to the structural similarity (SSIM) index. The SSIM Index is first expressed in terms of wavelet transform coefficients of orthogonal wavelet transforms. The wavelet domain representation of the SSIM Index, along with the assumption of a Gaussian prior for the wavelet coefficients is used to formulate the soft thresholding optimization problem. A locally optimal solution is found using a quasi-Newton approach. This solution is applied to denoise images in the wavelet domain. The visual quality of the images denoised using the proposed algorithm is shown to be higher compared to the MSE-optimal soft thresholding denoising solution, as measured by the SSIM Index.

Index Terms— Image denoising.

1. INTRODUCTION

Image denoising is an important image processing problem. The literature is rich with several excellent denoising solutions such as the sparse 3-D collaborative filtering [1], presence of signal of interest based algorithm [2], Gaussian scale mixture (GSM) based minimum mean squared error (MSE) solution [3], the non-local (NL) means method [4] to name a few. A majority of denoising solutions use distortion measures that are not perceptually motivated, and more often than not use the MSE. It has been shown however, that the MSE is not the best metric either for quality assessment or for optimizing image processing algorithms [5]. The MSE is popular because it lends itself well to analysis, and due to a lack of competitive image quality assessment (IQA) algorithms. Recent advances in full-reference IQA have resulted in a number of powerful new algorithms such as the SSIM Index [6].

Based on the performance of the SSIM Index as a powerful IQA algorithm, using it as the objective function in optimizing image processing algorithms appears very promising. This optimization is not straightforward, however, given the form of the SSIM Index; algorithms that explicitly optimize for it are only recently being developed [7, 8]. In this paper, we propose a soft thresholding algorithm based on the SSIM Index and apply it to image denoising. Soft thresholding is considered not only due to its strength as a denoising solution [9] but also its relatively simple mathematical form.

We begin with a brief overview of the space domain SSIM Index. The SSIM Index is then expressed in terms of wavelet coefficients (of orthogonal wavelet transforms), and the soft thresholding optimization problem is formulated using the wavelet domain definition. The optimization problem is non-linear and non-convex, thereby making it a non-trivial one. We present a locally optimal solution to the problem and apply the same to denoise images in

the wavelet domain. The denoising results demonstrate that the proposed solution gives a higher perceptual quality to the denoised images when compared to the traditional MSE-optimal solution.

2. THE SSIM INDEX

The most general form of the metric that is used to measure the structural similarity between two signal vectors \mathbf{x} and \mathbf{y} in \mathbf{R}^n is

$$SSIM(\mathbf{x}, \mathbf{y}) = [l(\mathbf{x}, \mathbf{y})]^\alpha [c(\mathbf{x}, \mathbf{y})]^\beta [s(\mathbf{x}, \mathbf{y})]^\gamma. \quad (1)$$

The term $l(\mathbf{x}, \mathbf{y}) = \frac{2\mu_x\mu_y + C_1}{\mu_x^2 + \mu_y^2 + C_1}$ compares the mean of the signals, $c(\mathbf{x}, \mathbf{y}) = \frac{2\sigma_x\sigma_y + C_2}{\sigma_x^2 + \sigma_y^2 + C_2}$ compares the variance of the signals, and $s(\mathbf{x}, \mathbf{y}) = \frac{\sigma_{xy} + C_3}{\sigma_x\sigma_y + C_3}$ measures the correlation of the signals. The quantities μ_x, μ_y are the sample means of \mathbf{x} and \mathbf{y} respectively, σ_x^2, σ_y^2 are the sample variances of \mathbf{x} and \mathbf{y} respectively, and σ_{xy} is the sample cross-covariance between \mathbf{x} and \mathbf{y} . The constants C_1, C_2, C_3 are used to stabilize the metric for the case where the means and variances become very small. The parameters $\alpha > 0$, $\beta > 0$, and $\gamma > 0$, are used to adjust the relative importance of the three components. We use the following simplified form of the SSIM Index in our work (with $\alpha = \beta = \gamma = 1$, and $C_3 = C_2/2$):

$$SSIM(\mathbf{x}, \mathbf{y}) = \left(\frac{2\mu_x\mu_y + C_1}{\mu_x^2 + \mu_y^2 + C_1} \right) \left(\frac{2\sigma_{xy} + C_2}{\sigma_x^2 + \sigma_y^2 + C_2} \right). \quad (2)$$

In image quality assessment, pixel values of local image patches from the reference and distorted image constitute \mathbf{x} and \mathbf{y} respectively. The term $l(\mathbf{x}, \mathbf{y})$ compares the luminance, $c(\mathbf{x}, \mathbf{y})$ compares the contrast, and $s(\mathbf{x}, \mathbf{y})$ compares the structure of the local image patches. The average of the SSIM values across the image (also called mean SSIM or MSSIM) gives the final quality measure. The key idea behind the SSIM Index is to recognize that natural images are highly structured, and that the measure of structural correlation (between the reference and the distorted image) is very important in deciding the overall visual quality. Further, the SSIM Index measures quality locally and is able to capture local dissimilarities better, unlike global quality measures such as MSE (and hence PSNR). Though (2) has a form that is more complicated than MSE, it remains analytically tractable.

3. EXPRESSING SSIM INDEX IN WAVELET DOMAIN

To express the space domain SSIM Index in terms of wavelet coefficients, the space domain means, variances, and cross-covariance terms must be expressed in terms of the wavelet coefficients. Of the several classes of wavelet transforms [10], only orthogonal wavelets are energy preserving. This property allows for the space domain

variance and covariance terms to be expressed in terms of the wavelet coefficients in a straightforward manner. The analysis in this paper considers only orthogonal wavelet basis and holds good for any orthogonal wavelet basis.

3.1. Mean Calculation

The calculation of the mean from the wavelet coefficients depends on the image size $N \times N$ and the number of levels of decomposition. The approximation subband (low-low (LL) subband) of the resulting wavelet subband contains all the information required to calculate the mean of the space domain signal. A known scaling factor k is applied to the mean of the LL subband to find the mean. Let \mathbf{x} denote an image patch of size $N \times N$ and \mathbf{X} denote the L level wavelet transform of the patch (also of size $N \times N$).

$$\mu_{\mathbf{x}} = k^L \mu_{\mathbf{X},LL}, \quad (3)$$

where k is the scaling factor, and $\mu_{\mathbf{X},LL}$ is the mean of LL subband of \mathbf{X} . For e.g., if three levels of decomposition were applied to an 8×8 patch, $\mu_x = (k)^3 X(0,0)$ (since $\mu_{\mathbf{X},LL} = X(0,0)$).

3.2. Variance and Covariance Calculation

The calculation of variance and covariance makes use of the property that orthogonal wavelet bases obey the Parseval's theorem. Let \mathbf{x}, \mathbf{y} , represent image patches of size $N \times N$ and \mathbf{X}, \mathbf{Y} be their respective orthogonal transforms. From Parseval's theorem, it follows that

$$\sigma_{\mathbf{x}}^2 = \frac{1}{N^2} \sum_{i=0}^{N-1} \sum_{j=0}^{N-1} X_{i,j}^2 - (k^L \mu_{\mathbf{X},LL})^2, \quad (4)$$

$$\sigma_{\mathbf{xy}} = \frac{1}{N^2} \sum_{i=0}^{N-1} \sum_{j=0}^{N-1} X_{i,j} Y_{i,j} - (k^L \mu_{\mathbf{Y},LL})(k^L \mu_{\mathbf{X},LL}). \quad (5)$$

Replacing the space domain mean, variance, and covariance terms in (2) with the expressions in (3), (4), (5) gives

$$SSIM(\mathbf{x}, \mathbf{y}) = \left(\frac{2(k^L \mu_{\mathbf{X},LL})(k^L \mu_{\mathbf{Y},LL}) + C_1}{(k^L \mu_{\mathbf{X},LL})^2 + (k^L \mu_{\mathbf{Y},LL})^2 + C_1} \right) \left(\frac{2 \frac{1}{N^2} \sum_{i=0}^{N-1} \sum_{j=0}^{N-1} X_{i,j} Y_{i,j} - (k^{2L} \mu_{\mathbf{Y},LL} \mu_{\mathbf{X},LL}) + C_2}{\frac{1}{N^2} \sum_{i=0}^{N-1} \sum_{j=0}^{N-1} X_{i,j}^2 + Y_{i,j}^2 - k^{2L} (\mu_{\mathbf{X},LL}^2 + \mu_{\mathbf{Y},LL}^2) + C_2} \right). \quad (6)$$

4. SSIM-BASED SOFT THRESHOLDING

The soft thresholding operator in (7) with threshold λ is applied to the wavelet coefficients of the noisy image (denoted by y). The thresholded wavelet coefficients are inverted to get the space domain denoised image. The approximation subband is not thresholded.

$$g(y) = \text{sgn}(y)(|y| - \lambda)_+. \quad (7)$$

The wavelet domain representation of the SSIM Index in (6) allows for the formulation of the SSIM-optimal soft thresholding problem. In this paper, it is assumed that one threshold per subband is used. The optimization problem comprises the design of the thresholds so that the SSIM Index between the reference and the soft thresholded output is maximized.

4.1. Problem Formulation

Let \mathbf{x} be the reference image patch of size $N \times N$, \mathbf{n} be zero mean Gaussian noise, and $\mathbf{y} = \mathbf{x} + \mathbf{n}$ be the noisy observation of \mathbf{x} . Let \mathbf{X}, \mathbf{Y} represent an L level orthogonal wavelet transform of \mathbf{x}, \mathbf{y} respectively (all of size $N \times N$). An L level orthogonal transform consists of $3L$ subbands, and hence $3L$ thresholds. Let $\Lambda = [\lambda_1, \lambda_2, \dots, \lambda_{3L}]$ denote the vector of thresholds applied to each of the subbands. Let $\hat{\mathbf{X}}$ (a function of \mathbf{Y}, Λ) be the soft thresholded output, and let $\hat{\mathbf{x}}$ be the space domain version of $\hat{\mathbf{X}}$.

It is assumed that the noise variance is known at the receiver. Since only the observation \mathbf{y} is known, a direct evaluation of the SSIM Index between \mathbf{x} and $\hat{\mathbf{x}}$ is not possible. The following observations combined with the assumption that wavelet coefficients are Gaussian distributed are used to evaluate the SSIM Index between the reference and soft thresholded image patches.

Since the noise is zero mean, the mean of the reference and the thresholded estimate are identical (since the approximation subband is not thresholded). This makes the mean term in the SSIM Index equal to identity.

Since the noise is additive, the source variance can be estimated to be the difference between the variance of the observation σ_y^2 and the noise variance σ_n^2 as

$$\begin{aligned} \sigma_x^2 &\approx \sigma_y^2 - \sigma_n^2 \\ &= \frac{1}{N^2} \sum_{i=0}^{N-1} \sum_{j=0}^{N-1} Y_{i,j}^2 - (k^L \mu_{\mathbf{Y},LL})^2 - \sigma_n^2. \end{aligned} \quad (8)$$

The expression for the SSIM Index is rewritten as

$$SSIM(\mathbf{x}, \hat{\mathbf{x}}) = \left(\frac{2 \frac{1}{N^2} \sum_{i=0}^{N-1} \sum_{j=0}^{N-1} X_{i,j} \hat{X}_{i,j} - (k^L \mu_{\mathbf{Y},LL})^2 + C_2}{\frac{1}{N^2} \sum_{i=0}^{N-1} \sum_{j=0}^{N-1} Y_{i,j}^2 + \hat{X}_{i,j}^2 - 2(k^L \mu_{\mathbf{Y},LL})^2 - \sigma_n^2 + C_2} \right). \quad (9)$$

To evaluate the above expression, the summation is split across the approximation subband and the rest of the subbands as

$$SSIM(\mathbf{x}, \hat{\mathbf{x}}) = \frac{N(\mathbf{X}, \hat{\mathbf{X}})}{D(\mathbf{Y}, \hat{\mathbf{X}})}, \quad (10)$$

where $N(\mathbf{X}, \hat{\mathbf{X}}) = 2 \frac{1}{N^2} (\sum_{i,j \in LL} X_{i,j} \hat{X}_{i,j} + \sum_{i,j \notin LL} X_{i,j} \hat{X}_{i,j}) - (k^L \mu_{\mathbf{Y},LL})^2 + C_2$, $D(\mathbf{Y}, \hat{\mathbf{X}}) = \frac{1}{N^2} \sum_{i,j} Y_{i,j}^2 + \frac{1}{N^2} (\sum_{i,j \in LL} \hat{X}_{i,j}^2 + \sum_{i,j \notin LL} \hat{X}_{i,j}^2) - 2(k^L \mu_{\mathbf{Y},LL})^2 - \sigma_n^2 + C_2$. Thresholding is not applied to the approximation band, which means $\hat{X}_{LL} = Y_{LL}$. Due to the additive nature of the noise, the first term in the above summation can be approximated by

$$\frac{1}{N^2} \sum_{i,j \in LL} X_{i,j} \hat{X}_{i,j} \approx \frac{N_{LL}^2}{N^2} (\sigma_{\mathbf{Y},LL}^2 - \sigma_n^2 + \mu_{\mathbf{Y},LL}^2), \quad (11)$$

where N_{LL}^2 is the number of wavelet coefficients in the approximation band, $\sigma_{\mathbf{Y},LL}^2, \mu_{\mathbf{Y},LL}$ are the variance and mean of the approximation band of the noisy observation \mathbf{y} . Similarly $D(\mathbf{Y}, \hat{\mathbf{X}})$ can be simplified as,

$$\frac{1}{N^2} \sum_{i,j \in LL} \hat{X}_{i,j}^2 \approx \frac{N_{LL}^2}{N^2} (\sigma_{\mathbf{Y},LL}^2 + \mu_{\mathbf{Y},LL}^2), \quad (12)$$

where the same notation as above is used. Substituting these into the SSIM expression in (6),

$$SSIM(\mathbf{x}, \hat{\mathbf{x}}) = \frac{N_1(\mathbf{X}, \hat{\mathbf{X}})}{D_1(\mathbf{Y}, \hat{\mathbf{X}})}, \quad (13)$$

where $N_1(\mathbf{X}, \hat{\mathbf{X}}) = 2\frac{1}{N^2}(N_{LL}^2(\sigma_{\mathbf{Y},LL}^2 - \sigma_n^2 + \mu_{\mathbf{Y},LL}^2) + \sum_{i,j \notin LL} X_{i,j} \hat{X}_{i,j}) - (k^L \mu_{\mathbf{Y},LL})^2 + C_2$, $D_1(\mathbf{Y}, \hat{\mathbf{X}}) = \frac{1}{N^2} \sum_{i,j} Y_{i,j}^2 + \frac{1}{N^2}(N_{LL}^2(\sigma_{\mathbf{Y},LL}^2 + \mu_{\mathbf{Y},LL}^2) + \sum_{i,j \notin LL} \hat{X}_{i,j}^2) - 2(k^L \mu_{\mathbf{Y},LL})^2 - \sigma_n^2 + C_2$. The \hat{X} values from the remaining subbands are a function of the thresholds Λ . Further, the above expression cannot be simplified any further based on the information available to the denoiser. To estimate the remaining terms in the summation in (13), a Gaussian model for the source statistics of the wavelet coefficients is assumed. It is also known that these subbands have zero mean. With this assumption, the empirical values for the correlation and variance in (13) are replaced by their statistical equivalents.

In the following, expressions for the correlation between the reference and the thresholded estimate, and the variance of the thresholded estimate are derived under the Gaussian assumption for the source. From (7), the first and second order statistics of \hat{X} are

$$\mu_{\hat{X}} = \mathbf{E}[\hat{X}] = \mathbf{E}[g(Y)] = 0, \quad (14)$$

$$\begin{aligned} \sigma_{\hat{X}}^2 &= f(\sigma_Y, \lambda) = \mathbf{E}[(\hat{X} - \mu_{\hat{X}})^2] = \mathbf{E}[(g(Y))^2] \\ &= (\sigma_Y^2 + \lambda^2) \left[1 - \operatorname{erf} \left[\frac{\lambda}{\sqrt{2}\sigma_Y} \right] \right] - \sqrt{\frac{2}{\pi}} \sigma_Y \lambda \exp \left[\frac{-\lambda^2}{2\sigma_Y^2} \right], \end{aligned} \quad (15)$$

where $\operatorname{erf}(x) = \frac{2}{\sqrt{\pi}} \int_{t=0}^x e^{-t^2} dt$ is the error function. The derivation is omitted for brevity.

To calculate the covariance term, the MSE result from (10) in [11] is used. Also, since the source and the noise are assumed to have zero mean, the covariance and correlation terms are identical. The covariance between the source and the estimate is

$$\begin{aligned} \sigma_{X\hat{X}} &= h(\sigma_X, \sigma_Y, \lambda) = \sigma_X^2 + \sigma_{\hat{X}}^2 - MSE(X, \hat{X}) \\ &= \sigma_X^2 \left[1 - \operatorname{erf} \left[\frac{\lambda}{\sqrt{2}\sigma_Y} \right] \right]. \end{aligned} \quad (16)$$

Using (11), (12), (15), (16), the optimization problem becomes

$$\begin{aligned} \Lambda^* &= \operatorname{argmax}_{\Lambda \in \mathbf{R}_+^{3L}} SSIM(\mathbf{x}, \hat{\mathbf{x}}) \\ &= \operatorname{argmax}_{\Lambda \in \mathbf{R}_+^{3L}} \frac{N_g(\mathbf{x}, \hat{\mathbf{x}})}{D_g(\mathbf{x}, \hat{\mathbf{x}})}, \end{aligned} \quad (17)$$

where $N_g(\mathbf{X}, \hat{\mathbf{X}}) = 2\frac{1}{N^2}(N_{LL}^2(\sigma_{\mathbf{Y},LL}^2 - \sigma_n^2 + \mu_{\mathbf{Y},LL}^2) + \sum_{i=1}^{3L} N_i h(\sigma_{X_i}, \sigma_{Y_i}, \lambda_i)) - (k^L \mu_{\mathbf{Y},LL})^2 + C_2$, $D_g(\mathbf{Y}, \hat{\mathbf{X}}) = \frac{1}{N^2} \sum_{i,j} Y_{i,j}^2 + \frac{1}{N^2}(N_{LL}^2(\sigma_{\mathbf{Y},LL}^2 + \mu_{\mathbf{Y},LL}^2) + \sum_{i=1}^{3L} N_i f(\sigma_{Y_i}, \lambda_i)) - 2(k^L \mu_{\mathbf{Y},LL})^2 + C_2$. The functions $h()$, $f()$ are from (16), (15) respectively, N_i is the number of wavelet coefficients in subband i . The source subband variance $\sigma_{X_i}^2$ is estimated using the relation $\sigma_{X_i}^2 \approx \sigma_{Y_i}^2 - \sigma_n^2$.

4.2. Solution

The objective function is nonlinear in the design parameters Λ , and returns a scalar value for the vector input. The only constraint on the solution is that Λ be non-negative. Of the several solutions available to solve such nonlinear optimization problems [12], the quasi-Newton method provides a good tradeoff between complexity and performance. Specifically, the Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm [12] is used to find the local optimum. An implementation of this algorithm from Matlab's optimization toolbox (called `fminunc`) is used in the solution. There is currently no guarantee, however, whether the solution found is global.

5. RESULTS

The computed (locally optimal) solution is now applied to denoise images that have been distorted with additive white Gaussian noise. The steps involved in the implementation of the denoising algorithm are outlined below. It is assumed that the noise variance is known at the receiver.

- Divide noisy image into non-overlapping blocks of size 32×32
- Apply L level orthogonal wavelet transform to each block
- For each wavelet transformed block, compute the statistics of the subbands using the simplifications in Section 4.1, (15), (16)
- Solve the optimization problem in (17) using the BFGS algorithm to find a locally optimal Λ^* . The algorithm is initialized using the MSE-optimal soft thresholding solution from Chang et al. [11]
- Soft threshold the noisy wavelet coefficients using Λ^*
- Apply inverse wavelet transform on a block by block basis, for all the blocks in the image

The performance of this locally SSIM-optimal algorithm is compared to the MSE-optimal solution by Chang et al. [11]. The Chang et al., soft thresholding solution has been shown to be a very powerful denoising method. Their solution has been shown to be nearly MSE-optimal for several popularly used models for wavelet coefficients including Gaussian, Laplacian, and Generalized Gaussian sources. Further, their solution is space varying and adapts based on the local subband statistics.

The denoising results are presented in Fig. 1. A comparison of Figs. 1(c) and 1(d) reveals the differences between the MSE-optimal and the proposed SSIM-based algorithm. First, the MSE and SSIM values of the images denoised using the algorithms in question are consistent – MSE-optimal solution has lower MSE and SSIM Index, and vice-versa for the SSIM-based solution (even though the proposed algorithm is locally SSIM-optimal). More importantly, the perceptual quality of the SSIM-based solution is higher than the MSE-optimal solution. This claim is made based on the following observations: better overall contrast, retention of more image detail especially in the whiskers and hair region, and finally a higher SSIM value. Similar improvement (subtle but important) was seen in several other test images from the 'Austin and Vicinity' database [13]. These results are summarized in Table 1.

6. CONCLUSIONS AND FUTURE WORK

In this paper, we presented a novel algorithm for SSIM-based soft thresholding and applied it to denoise images in the wavelet domain.

Image	σ_n	L	Noisy		MSE-optimal		SSIM-based	
			MSE	SSIM	MSE	SSIM	MSE	SSIM
Img0039	40	3	1600	0.4016	512	0.5154	586	0.5350
Img0073	30	3	901	0.5545	386	0.6293	408	0.6551
Mandrill	50	3	2492	0.2766	509	0.4835	577	0.4954
Img0043	40	3	1593	0.4458	557	0.5733	600	0.5815

Table 1. Denoising results for a set of images from the ‘Austin and Vicinity’ database.




Fig. 1. 1(a) Reference image Mandrill. 1(b) Noisy image with $\sigma_n = 50$, MSE = 2492.49, SSIM Index = 0.2766. 1(c) Image denoised with the MSE-optimal algorithm, MSE = 509.16, SSIM Index = 0.4835. 1(d) Image denoised with the SSIM-based algorithm, MSE = 577.54, SSIM Index = 0.4954.

The SSIM Index is first expressed in terms of wavelet transform coefficients, following by the formulation of the optimization problem. A locally optimal solution is found using a quasi-Newton approach. The performance of the proposed denoising solution is shown to be better than the MSE-optimal soft thresholding solution in terms of visual quality (as confirmed by higher SSIM Index values).

The proposed solution is a first step towards SSIM-optimal soft thresholding. There are several improvements that can be made to the proposed solution. A globally optimal solution, though involved, should give better denoising results. The Gaussian model for wavelet statistics can be replaced with more accurate models such as generalized Gaussian or Gaussian scale mixture (GSM). On the implementation side, non-overlapping blocks can be replaced with overlapping image blocks to reduce blocking artifacts. We are working on incorporating all these improvements into the proposed solution.

7. REFERENCES

[1] Kostadin Dabov, Alessandro Foi, Vladimir Katkovnik, and Karen Egiazarian, “Image Denoising by Sparse 3-D Transform-Domain Collaborative Filtering,” *IEEE Trans. Image Processing*, vol. 16, pp. 2080–2095, 2007.

[2] Aleksandra Pizurica and Wilfried Philips, “Estimating the

Probability of the Presence of a Signal of Interest in Multiresolution Single- and Multiband Image Denoising,” *IEEE Trans. Image Processing*, vol. 15, no. 3, pp. 654–665, Mar. 2006.

[3] Javier Portilla, V. Strela, M. J. Wainwright, and E. P. Simoncelli, “Image Denoising using Scale Mixtures of Gaussians in the Wavelet Domain,” *IEEE Trans. Image Processing*, vol. 52, no. 11, pp. 1338–1351, Nov. 2003.

[4] Buades A., Coll B., and Morel J. M., “A review of image denoising algorithms, with a new one,” *SIAM Multiscale Model. Simul.*, vol. 4, pp. 490–530, 2005.

[5] Bernd Girod, “What’s Wrong with Mean-squared Error?,” in *Digital Images and Human Vision*, Andrew B. Watson, Ed., pp. 207–220. Cambridge, MA: MIT Press, 1993.

[6] Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli, “Image Quality Assessment: From Error Visibility to Structural Similarity,” *IEEE Trans. Image Processing*, vol. 13, no. 4, pp. 600–612, Apr. 2004.

[7] Sumohana S. Channappayya, Alan C. Bovik, and Robert W. Heath Jr., “A Linear Estimator Optimized for the Structural Similarity Index and its Application to Image Denoising,” in *Proc. of IEEE Intl. Conference on Image Processing*, 2006, vol. 12.

[8] Zhou Wang, Q. Li, and X. Shang, “Perceptual Image Coding Based on a Maximum of Minimal Structural Similarity Criterion,” in *IEEE International Conference on Image Processing*, Sept. 2007.

[9] David L. Donoho, “Nonlinear solution of Linear Inverse Problems by Wavelet-Vaguelette Decomposition,” *Applied and Computational Harmonic Analysis*, vol. 2, no. 2, pp. 101–126, 1995.

[10] Martin Vetterli and Jelena Kovacevic, *Wavelets and Subband Coding*, Prentice Hall, 1995.

[11] S. Grace Chang, Bin Yu, and Martin Vetterli, “Adaptive wavelet tresholding for image denoising and compression,” *IEEE Trans. Image Processing*, vol. 9, pp. 1532–1546, 2000.

[12] Dmitri P. Bertsekas, *Nonlinear Programming*, Athena Scientific, Belmont, Massachusetts, 1995.

[13] Visual Delights Inc., “Austin and Vicinity,” CD, 2003.