

Noise Reduction in Video Images Using Coring on QMF Pyramids

by

Arthur J. Kalb

Submitted to the Department of Electrical Engineering and
Computer Science

in partial fulfillment of the requirements for the degree of

Bachelor of Science in Electrical Science and Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

May 1991

© Arthur J. Kalb, MCMXCI.

The author hereby grants to MIT permission to reproduce and
to distribute copies of this thesis document in whole or in part.

Author _____
Department of Electrical Engineering and Computer Science
May 20, 1991

Certified by _____
Edward H. Adelson
Associate Professor
Thesis Supervisor

Accepted by _____
Leonard A. Gould
Chairman, Department Committee on Undergraduate Theses

MASSACHUSETTS INSTITUTE
OF TECHNOLOGY

JUL 11 1991

LIBRARIES

Noise Reduction in Video Images Using Coring on QMF Pyramids

by

Arthur J. Kalb

Submitted to the Department of Electrical Engineering and Computer Science
on May 26, 1991, in partial fulfillment of the
requirements for the degree of
Bachelor of Science in Electrical Science and Engineering

Abstract

This thesis investigates a technique known as "coring" which can be used to reduce noise in two-dimensional video images. Coring passes the high spatial frequency components of a noisy image through a non-linear "coring function." This function operates on each pixel and is a function of pixel intensity. This thesis provides a theoretical justification for the success of coring in noise reduction. Coring functions are generated using the theoretical results and then tested on a set of images. Finally, this thesis develops an algorithm for noise reduction given a statistical characterization of the noise.

Thesis Supervisor: Edward H. Adelson
Title: Associate Professor

Acknowledgments

I would like to thank Ted and Eero for their guidance, support and patience. Thanks also to everyone in the Vision and Modeling lab who helped me out. I would like to thank my friends, especially Juan, Tom, Jason, Steve, Dave, Bob, Phil, Sanjeev, and Ann for helping me keep good spirit even when things looked bleak. I would like to thank my mother and father for their support (both financial and emotional). Most especially, I thank Margaret for her love in good and bad times. Thanks be to God.

Contents

1	Introduction	10
2	Related Work	12
2.1	Image Smoothing	12
2.2	Wiener Filtering	13
3	Pyramid Filtering	14
3.1	Laplacian Pyramids	14
3.2	Quadrature Mirror Filter Pyramids	15
4	An Overview of Coring	21
4.1	The Coring Procedure	21
4.2	Bayer and Powell's Work	22
4.3	Adelson and Ogden's Work	23
4.4	Lee's Work	23
5	Coring Theory	33
5.1	Assumptions	34
5.2	Derivation of the Coring Function	34
6	Results	37
6.1	Testing the Theory	37
6.2	Comparison with Wiener Filtering	39
6.3	Comparison to Bayer and Powell's Coring	40

6.4	Comparison with Adelson and Ogden's Coring	40
6.5	Dependence of Coring on Signal Model	41
6.6	Effect of Different Transforms	41
7	A Semi-Automatic Noise Reduction System	52
7.1	Image Subband Histogram Models	52
7.2	The Noise Reduction System	55
8	Conclusion	58
A	Appendix	60
A.1	QMF Transform Tap Values	60

List of Figures

3-1	The original Lenna image.	16
3-2	The low-pass subband of Lenna.	17
3-3	The vertical subband of Lenna.	17
3-4	The horizontal subband of Lenna.	18
3-5	The diagonal subband of Lenna.	18
3-6	The second level low-pass subband of Lenna.	19
3-7	The second level vertical subband of Lenna.	19
3-8	The second level horizontal subband of Lenna.	20
3-9	The second level diagonal subband of Lenna.	20
4-1	A noisy version of the Lenna image.	24
4-2	Information flow in coring (one level).	25
4-3	The low-pass band of the noisy Lenna image.	25
4-4	The horizontal subband of the noisy Lenna image.	26
4-5	The vertical subband of the noisy Lenna image.	26
4-6	The diagonal subband of the noisy Lenna image.	27
4-7	The coring function for the horizontal subband.	27
4-8	The coring function for the vertical subband.	28
4-9	The coring function for the diagonal subband.	28
4-10	The cored horizontal subband of the noisy Lenna image.	29
4-11	The cored vertical subband of the noisy Lenna image.	29
4-12	The cored diagonal subband of the noisy Lenna image.	30
4-13	A restored version of the noisy Lenna image (one level coring).	31

4-14	The noise-free Lenna image.	32
6-1	The noisy Lenna image. Fixed noise added to Lenna.	43
6-2	The second level diagonal band of the fixed noise image.	44
6-3	The second level diagonal band of the noisy Lenna image.	45
6-4	A histogram for the second level diagonal band of the Lenna image. .	46
6-5	A histogram for the second level diagonal band of the noise image. . .	46
6-6	A histogram for the second level diagonal band of the noisy image. . .	47
6-7	A coring function for the second level diagonal band.	47
6-8	A histogram for the second level cored diagonal band.	48
6-9	A second level diagonal band after coring.	49
6-10	A restored image (one level coring).	50
6-11	A restored image (two level coring).	51
7-1	A restored image using the semi-automatic procedure (two level coring). .	57

List of Tables

6.1	Results of coring on image plus fixed noise over a set of six images. Coring was done to one and two levels.	39
6.2	Results of different noise reduction strategies on a noisy Lenna image.	40
6.3	Results of coring when the coring functions are swapped.	41
6.4	Results of coring as filter length varies	42
7.1	Results of coring when the model error is minimized with respect to the numerator or the denominator of the coring function expression. .	53
7.2	A comparison of various models' SNR Gain.	55
7.3	A comparison of results when using the original image histograms, a two parameter fit model, and the semi-automatic method.	56
A.1	Tap Values used for Construction of QMF Pyramids	61

Chapter 1

Introduction

Noise in video images often detracts from the viewing enjoyment of pictures and frequently can obliterate the pictures entirely. Typically, transmission of pictures introduces noise into the image. The noise that this thesis deals with is that of uncorrelated additive white noise. There are other types of noise, such as pink noise (noise with some spatial correlation), speckle noise or multiplicative noise, quantization errors, and coding/decoding errors. These types of noise will not be dealt with in this thesis.

A tradeoff must be made when removing noise. It would be simple to remove all the noise in an image, if one were to also remove all of the desired un-noisy (signal) image as well. This may be an exaggerated example but it does illustrate the basic tradeoff of noise reduction and image fidelity. Some solutions remove the high spatial frequency components of a noisy image. This does eliminate much of the noise, but it also tends to blur the image. It also leaves low frequency noise. Thus, noise reduction comes at the cost of blurring.

Several techniques have been developed for noise reduction. They range in complexity, generally achieving better results with added complexity. They also make different assumptions about the nature of the noise and the signal. Sometimes these assumptions prove to be constrictive for wide practical use. Two schemas, noise smoothing, and Wiener filtering demand some attention. The basic principles, advantages and disadvantages will be discussed in this thesis. It is important to note that both

perform some sort of high frequency elimination, producing blurring.

Coring solves some of the problem inherent in linear filtering techniques. Coring employs a non-linear mapping to reduce the noise. Coring anticipates that the high-frequency component of signal power is somewhat larger than the noise power. Given this assumption, it makes sense that low intensity pixels in a high-passed filtered image will be noise. Whereas high intensity pixels will most likely be signal. It then makes sense to severely attenuate low intensity pixels and leave high intensity pixels unchanged. This is the essence of coring. When coring is performed on high frequency components of a set of directionally filtered images the benefits increase. Since uncorrelated noise power spreads evenly in all directions and edges tend to align themselves directionally, edges need not be attenuated as much in order to achieve the same overall noise reduction. Quadrature Mirror Filter (QMF) Pyramids prove useful for this purpose. Since uncorrelated noise power spreads evenly in all directions and edges tend to align themselves directionally, edges need not be attenuated as much in order to achieve the same overall noise reduction.

Previously, coring functions were chosen intuitively, realizing that low intensities must be attenuated and high intensities should be left unchanged. No theoretical justification for any particular coring function had been developed. This thesis develops a theory for selecting the coring function that generates a cleaned image with mean square error minimal high frequency subbands. The theory is then tested in a set of test cases. Possible improvements in the details of the coring process are also examined. A procedure to clean a noisy image given only a characterization of the noise is also presented.

Chapter 2

Related Work

When reviewing introductory material on image restoration, a reader normally comes across two techniques for noise reduction, image smoothing and Wiener filtering. Image smoothing is a spatial domain technique providing an extremely simple and efficient algorithm to reduce the noise. Unfortunately, smoothing inherently produces blurry results. Another algorithm is Wiener filtering. Among linear filtering techniques, Wiener filtering promises the minimal mean square error in the restored image. Unfortunately, the assumptions necessary for Wiener filtering are rarely satisfied.

2.1 Image Smoothing

Image smoothing was first developed by R.E. Graham in 1962 as a technique to remove uncorrelated additive noise, also termed “snow” [7]. Smoothing essentially performs a low-pass filter operation. Because images tend to exhibit local correlations, much of the signal energy is located in the low frequency spectrum. Thus when white noise (flat frequency energy spectrum) is added, the low frequency signal to noise ratio (SNR) is greater than the high frequency signal to noise ratio [10]. Image smoothing attempts to increase the overall SNR by attenuating the high frequency portion. This generally improves overall SNR at the expense of high frequency signal information such as edges and significant blurring results. The amount of blurring can be reduced depending on the implementation chosen.

2.2 Wiener Filtering

Another technique frequently discussed in literature on noise reduction is Wiener filtering. Wiener filtering was applied to image restoration by Helstrom [7]. Among linear transformations, it promises mean square error minimal restoration provided a number of conditions are satisfied. For a proof of this the reader is referred to [6] and [7]. These conditions demand that the signal and noise be stationary processes. The stationarity assumption is questionable in many applications. The Wiener filter is often expressed as a frequency domain transfer function given by

$$\hat{F}(u, v) = \frac{G(u, v)}{1 + S_n(u, v)/S_f(u, v)}, \quad (2.1)$$

where $\hat{F}(u, v)$ is the two-dimensional Fourier transform of the cleaned image, $G(u, v)$ is the 2-D Fourier transform of the original noisy image, $S_n(u, v)$ is the Fourier transform of the noise correlation image (noise energy spectrum), and $S_f(u, v)$ is the Fourier transform of the original image correlation image (signal energy spectrum). The drawback of Wiener filtering is that it requires a knowledge of the signal and noise energies at each frequency. This is a weighty amount of information, most of which is generally not available.

Chapter 3

Pyramid Filtering

Recent years have seen the development of multi-scale linear image transforms for use in image data compression, motion analysis, texture segmentation and edge detection. These transforms are represented as recursive structures known as pyramids. Pyramids are formed by taking the original image, filtering the high frequency information out and storing it, subsampling the low frequency information by a factor of two and repeating the process. Filtering the high frequency information and storing it, generates an image that has a low entropy. Subsampling the low frequency portion reduces the amount of information to store. Thus, pyramid structures efficiently use storage space and are useful for image compression. Furthermore, multi-scale structures maintain some uniformity of representation over the range of scales, hence the statistics of the images are similar from scale to scale.

3.1 Laplacian Pyramids

Much early work in multi-scale transforms revolved around the Laplacian pyramid [1, 5]. The Laplacian pyramid is generated by first low-pass filtering the original image using a Gaussian filter. The difference between the original and the low-pass image are taken. This creates the first level high frequency portion of the pyramid. The process is repeated recursively on each successive low-pass image (each time having a lower cut-off frequency on the low-pass filter), thereby generating a pyramid structure.

The Laplacian pyramid has a number of benefits. It is a complete description of the original image, that is, there are no losses of information in the storage of the pyramid (excepting quantization errors). The sub-levels of the pyramid are upsampled, filtered and summed back to the original image. The Laplacian pyramid is computationally efficient, requiring few multiplications and additions. The failures of the Laplacian pyramid are its lack of orientation selectivity and its over-completeness (storing of redundant information).

The Laplacian pyramid works fine for application such as image compression. However, it lacks any orientation selectivity in the high-pass portions of the pyramid. This is a drawback for applications such as motion analysis, texture segmentation, edge detection, and coring. As a result, Quadrature Mirror Filter pyramids are becoming favored for image processing applications.

3.2 Quadrature Mirror Filter Pyramids

Quadrature Mirror Filter (QMF) pyramids [2, 3, 12, 13, 14] achieve spatial localization while trading off some advantages with the Laplacian pyramid. QMF pyramids for two-dimensional images can be formed using separable low-pass and high-pass kernels. A two-level QMF pyramid of the Lenna image is shown in Figures 3-2 through 3-9. The original Lenna image is shown in Figure 3-1. Note the transformed subbands were not subsampled. On each level of the pyramid, four filtered subbands are generated: the low-pass, the horizontal, the vertical and the diagonal. Each band is named by the orientation of edges within the subband. The low-pass band can be recursively filtered to generate a pyramid structure.

The second level of the pyramid shows similar spatial information to the first. The second level of the pyramid (again not subsampled) is shown in Figures 3-6–3-9. This similarity from scale to scale is a property of the QMF transform. This has repercussions on this thesis' work on coring. Because of this property, all levels of the pyramid exhibit similar intensity distributions (more commonly referred to as "histograms"). Thus it becomes possible to reliably generate a two parameter model



Figure 3-1: The original Lenna image.

of the subband histograms (See Chapter 7).

The orientation selectivity of the QMF pyramids proves beneficial in many applications as mentioned above. This improvement comes with the tradeoff of increased information storage and decreased computational efficiency. Relative to coring, the orientation selectivity accentuates image features relative to noise [11, 4]. This enhances the effectiveness of coring.



Figure 3-2: The low-pass subband of Lenna.




Figure 3-3: The vertical subband of Lenna.




Figure 3-4: The horizontal subband of Lenna.




Figure 3-5: The diagonal subband of Lenna.



Figure 3-6: The second level low-pass subband of Lenna.




Figure 3-7: The second level vertical subband of Lenna.




Figure 3-8: The second level horizontal subband of Lenna.




Figure 3-9: The second level diagonal subband of Lenna.

Chapter 4

An Overview of Coring

The foundations of coring originate with a paper by Bryce E. Bayer and Philip G. Powell [11]. In that paper, the idea of attenuating low intensities values of high-pass images is developed. Indeed, Bayer and Powell indicate the path followed in this thesis, which they had considered intractable. Bayer and Powell's research was followed by a number of unpublished papers on the subject. A paper by Adelson and Ogden [4] extended coring research by implementing the coring algorithm on oriented pyramids. Adelson and Ogden also introduced a continuous coring function, as opposed to Bayer and Powell's discontinuous function. W. Anthony Lee [9] furthered the work of Adelson and Ogden by implementing coring on QMF pyramids. Before delving into previous work, a brief glimpse of the coring procedure is given.

4.1 The Coring Procedure

The coring procedure starts with a noisy image (see Figure 4-1) as its input, and follows the procedure detailed below and shown in Figure 4-2. First the noisy image is broken down into subbands: low-pass (Figure 4-3), horizontal (Figure 4-4), vertical (Figure 4-5), and diagonal (Figure 4-6). At this point a decision is made whether or not to core to two levels. If the decision is made to go a second level, the low-pass subband is then broken down into subbands. This process can be repeated recursively. Typically when using QMF pyramid structures, the subbands are subsampled. This

maintains the orthogonality of the QMF transform. For this research, the subband images were not subsampled (the filters were padded with zeros though, to maintain the octave dividing characteristic of the QMF transform). By not subsampling, the subband is over-complete, meaning it contains redundant information.

After the subband images are generated, the high-pass subbands are operated on by the coring function, pixel by pixel. The coring function maps pixels intensities of the subband to new intensities. A set of coring functions are shown in Figures 4-7 through 4-9 along with their respective cored images in Figures 4-10 through 4-12.

After coring, the images are convolved with the inverting QMF transforms and added back together, leaving the resulting image shown in Figure 4-13. This can be compared with the original noise-free image shown in Figure 4-14.

4.2 Bayer and Powell's Work

Bayer and Powell [11] propose a discontinuous coring function that operates on the high frequency portions of a Laplacian-type pyramid. Their coring function is defined as

$$I_{out} = \begin{cases} 0 & \text{if } I_{in} < T \\ I_{in} & \text{otherwise} \end{cases} \quad (4.1)$$

where I_{out} is the cored output intensity, I_{in} is the input intensity and T is the threshold which varies with the standard deviation of the noise. Bayer and Powell do not give any algorithm for selecting the threshold level, T . Their work suggests the approach, taken in this thesis, of making an estimate of the most probable decomposition of noisy image into noise and signal. They also suggest the use of histogram models for the noise and signal to perform calculations. These two suggestions are the driving force behind this thesis' work on coring.

Bayer and Powell also suggest the use of oriented filtered images. They propose that an oriented pyramid structure would emphasize image features relative to noise, thereby making coring of low intensity noise more effective. When Bayer and Powell wrote their paper, QMF pyramids had not yet been developed. What they proposed

corresponds reasonably well with what is now known as QMF pyramids.

4.3 Adelson and Ogden's Work

Adelson and Ogden [4] extend Bayer and Powell's work to pyramid structures, in particular oriented pyramid structures. Their paper starts with an analysis of coring on a non-oriented Laplacian pyramid. Later in the paper they discuss results with an oriented version of the Burt pyramid. They find that oriented coring does much better than non-oriented coring.

Adelson and Ogden also experimented with filter shapes. In an attempt to avoid edge distortion caused by the threshold coring of Bayer and Powell, they developed a parameterized coring function

$$I_{out} = (1 - e^{-(m|I_{in}|)^p})I_{in}^k \quad (4.2)$$

where I_{out} is the coring output intensity, I_{in} is the input intensity, and m , k and p are parameters. The parameter m controls the amount of coring (the coring width), p controls the attenuation in the coring region, and k controls the amplification outside the coring region. Adelson and Ogden find that this coring function reduced edge distortions created by coring. They state that the shape of the coring function defines the restoration quality as much as the coring width does, yet offer no solutions to the optimal coring shape. This thesis presents a formulation for the optimal coring shape.

4.4 Lee's Work

W. Anthony Lee further advanced the development of coring when he implemented coring on QMF pyramids [9]. He detailed the implementation of the filtering process described in Section 4.1.



Figure 4-1: A noisy version of the Lenna image.

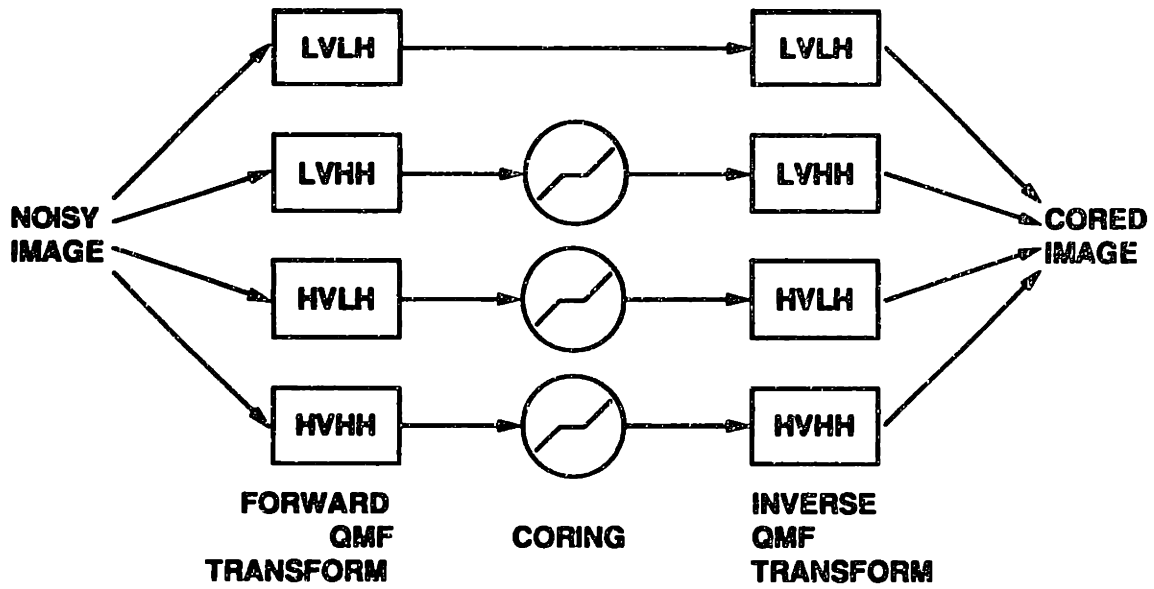


Figure 4-2: Information flow in coring (one level).




Figure 4-3: The low-pass band of the noisy Lenna image.




Figure 4-4: The horizontal subband of the noisy Lena image.




Figure 4-5: The vertical subband of the noisy Lena image.




Figure 4-6: The diagonal subband of the noisy Lenna image.




Figure 4-7: The coring function for the horizontal subband.




Figure 4-8: The coring function for the vertical subband.




Figure 4-9: The coring function for the diagonal subband.




Figure 4-10: The cored horizontal subband of the noisy Lenna image.




Figure 4-11: The cored vertical subband of the noisy Lenna image.




Figure 4-12: The cored diagonal subband of the noisy Lenna image.



Figure 4-13: A restored version of the noisy Lenna image (one level coring).



Figure 4-14: The noise-free Lenna image.

Chapter 5

Coring Theory

Typically, image information is low frequency information and white noise is flat band information. Thus in high pass filtered images, the Signal to Noise Ratio (SNR) is lower than in the original image. It makes sense then to attenuate the high frequency information. This can be done in several ways. Image smoothing low pass filters noisy information. Wiener filtering is a frequency domain technique which makes a linear estimate of the signal image given the noisy image.

Coring is a non-linear technique that has developed in response to many techniques that indiscriminantly eliminate high frequency information in order to remove additive noise. It is possible to throw out only the portion of the noisy information that is believed to be noise. This is the principle behind coring. While coring has been in use for about five years now, a theoretical justification for its success has never been proffered. The following material develops a formulization for generating coring functions. These coring functions produce mean square error minimal images in each high-pass subband, provided the assumptions of this formulation are satisfied. When the restored subband images are reconstructed, the resulting image is significantly cleaned of noise.

5.1 Assumptions

Few assumptions are placed on the theory to be developed in this chapter. The first is that the noise process is assumed to be an additive process. The noise is composed of an array of random variables, which represent pixel intensities. The random variables are given by a zero mean process and are independent from the signal. The second assumption is that the transform be lossless (or nearly lossless). There are other considerations while choosing a specific transform though. While the reason for coring's success depends on the reduced SNR in the high-pass band relative to the low-pass band, coring does better as the SNR in the high-pass band increases. Hence, the transform which maximizes the high-pass SNR should be selected for coring. This is also the reason why oriented transforms, such as the QMF, do better than non-oriented transforms.

5.2 Derivation of the Coring Function

Using the suggestion of Bayer and Powell [11], a Bayesian estimate of the optimal coring function in each subband is developed. Assume that in a subband, the noisy image \mathbf{Y} is composed of the sum of the noise \mathbf{N} and the original image signal \mathbf{X} . Mathematically,

$$\mathbf{Y} = \mathbf{X} + \mathbf{N}$$

where \mathbf{Y} , \mathbf{X} and \mathbf{N} are arrays of random variables. Since \mathbf{Y} , \mathbf{X} and \mathbf{N} are random variables at each pixel, they can be described by probability distribution functions (pdf's) $P_y(y)$, $P_x(x)$ and $P_n(n)$. Because the noise process is additive,

$$P_y(y) = P_x(x) * P_n(y)$$

where the $*$ represents convolution. Writing out the convolution,

$$P_y(y) = \int_{-\infty}^{\infty} P_x(x)P_n(y - x) dx.$$

Assuming an additive noise process, the equation

$$P_{y|x}(x, y) = P_n(n) = P_n(y - x)$$

relates the pdf of the conditional probability of y given x to the pdf for the noise.

Bayer and Powell's idea is to develop a Bayesian estimate for x given y . Bayesian analysis attempts to estimate the original intensity of a pixel given its noisy value. The first step is to apply Bayes' Law,

$$P_x(x)P_{y|x}(x, y) = P_y(y)P_{x|y}(x, y),$$

which relates the conditional probabilities of uncorrelated random events. Rewriting for $P_{x|y}(x, y)$, the probability of x happening if y happens,

$$P_{x|y}(x, y) = \frac{P_x(x)P_{y|x}(x, y)}{P_y(y)}.$$

Substituting in for $P_y(y)$ and $P_{y|x}(x, y)$ from above,

$$P_{x|y}(x, y) = \frac{P_x(x)P_n(y - x)}{P_x(y) * P_n(y)}.$$

The next step is to find the expected value of x given y using the relation

$$\hat{x}(y) = \int_{-\infty}^{\infty} x P_{x|y}(x, y) dx$$

which makes a minimal mean square error estimate for x . Substituting in for $P_{x|y}(x, y)$ one arrives at the coring function,

$$\hat{x}(y) = \frac{\int_{-\infty}^{\infty} x P_x(x) P_n(y - x) dx}{\int_{-\infty}^{\infty} P_x(x) P_n(y - x) dx}, \quad (5.1)$$

where $\hat{x}(y)$ is a function of y . It yields the expected intensity value of x given a particular y .

The coring function is computed for each subband of a transformed noisy image

and is then applied to the each noisy subband image. The cored subbands are then used in the reconstruction process.

Chapter 6

Results

Using the theory of Chapter 5, a test experiment was conducted to check the validity of the theory. A fixed noise image was added to six test images. The images were transformed using a nine tap QMF transform, then cored, and a clean image was reconstructed. Signal to Noise Ratio data is then collected. A comparative study was also done between Wiener filtering, Bayer-Powell coring, Adelson-Ogden coring and this research. Next, the dependence of coring performance on signal characterization was examined. Finally, the effect of filter kernel length on coring performance was examined.

6.1 Testing the Theory

In order to test the theory, a fixed noise source was added to a set of test images, which were subsequently cored. The noise was chosen to be Gaussian noise of variance $\sigma^2 = 64$. This level of noise was sufficient to test noise removal. It provided enough graininess to the noisy image to be visually bothersome, while at the same time did not obliterate edge information.

The test images were chosen to be "Lenna", "Einstein", "Kate", "Paolina", "Mandrill" and "Bench." They were selected for their variety in textures as well as the general quality of the image.

The following procedure was developed to test coring. It follows closely a proce-

ture developed by Adelson and Ogden[4]. The outline is as follows:

1. Add fixed noise image to the test image. Compute the Signal to Noise Ratio (SNR) of the noisy image using the formula

$$SNR(dB) = 10 \log \left(\frac{\sigma_{signal}^2}{\sigma_{noise}^2} \right),$$

where σ_{signal}^2 and σ_{noise}^2 are the variances of the signal and noise respectively.

2. Transform image, noise, and noisy image into subbands (not subsampled however). Use a two level pyramid structure.
3. Compute histograms for the noise image and original image subbands — these are the effective pdf's for use in the coring formula.
4. For each high-pass subband, compute the coring function from Equation 5.1.
5. For each high-pass subband, apply the coring function to the noisy image subband.
6. Reconstruct two images: one that is cored on two levels, one that is cored on one level.
7. Measure the mean square error between the cored images and the original image. Compute the SNR gain.

To assist in visualization of the process, various images are shown in the following pages. Figure 6-1 shows the Lenna image with the fixed noise added. Figure 6-2 shows the diagonal subband of the noise image and Figure 6-3 shows the diagonal subband of the noisy image. Figure 6-4, Figure 6-5, and Figure 6-6 show the original image diagonal subband histogram, the noise image diagonal subband histogram and the noisy image diagonal subband histogram respectively. Figure 6-7 shows the coring function calculated from the image and noise histograms of the diagonal subband and Figure 6-9 shows the cored diagonal subband and Figure 6-8 shows the histogram for the cored diagonal subband. This histogram can be compared with Figure 6-4 and

Image	SNR (pre-core)	Cored SNR (1 level)	Cored SNR (2 levels)	SNR Gain (1 level)	SNR Gain (2 levels)
Lenna	37.508	47.895	49.767	10.387	12.259
Einstein	30.985	38.722	40.149	7.737	9.164
Kate	36.247	45.377	46.248	9.130	10.001
Paolina	31.871	41.379	43.549	9.508	11.678
Mandril	31.043	34.951	35.195	3.908	4.152
Bench	42.429	46.954	47.379	4.525	4.950

Table 6.1: Results of coring on image plus fixed noise over a set of six images. Coring was done to one and two levels.

Figure 6-6, the coring functions for the original and noisy diagonal subbands, to see the effects of coring. Figure 6-10 and Figure 6-11 show the resultant cored images after one and two level respectively. The results of this study are shown in Table 6.1. The results show that the coring functions do a satisfactory job. The results also demonstrate that multi-level coring does a better job than single-level coring, which was previously known.

6.2 Comparison with Wiener Filtering

Among linear filters, Wiener filtering is the mean square error minimal linear filter. In this section, the Bayesian estimated coring function is compared with an approximation of the Wiener filter. An approximation of the Wiener filter is made in each subband instead of using a transfer function that is continuous with respect to frequency. The signal and noise energies in each band are measured and the appropriate linear transformations are calculated using Equation 2.1.

This comparison is done only on the Lenna image to one and two levels. The results demonstrate that the non-linear coring function does considerably better in terms SNR gain. The numerical results of this experiment are shown in Table 6.2.

Technique	SNR (pre-core)	Cored SNR (1 level)	Cored SNR (2 levels)	SNR Gain (1 level)	SNR Gain (2 levels)
Bayesian	37.508	47.895	49.767	10.387	12.259
Adelson- Ogden	37.508	48.297	50.593	10.789	13.085
Bayer- Powell	37.508	45.880	47.632	8.372	10.124
Wiener	37.508	42.442	42.089	4.934	4.581

Table 6.2: Results of different noise reduction strategies on a noisy Lenna image.

6.3 Comparison to Bayer and Powell's Coring

The Bayesian estimated coring functions are also compared with the Bayer-Powell coring functions given by Equation 4.1. The threshold level T is set to $2\sigma_{noise}$ in each band. Coring is done to one and two levels. The Bayer-Powell coring functions generate reasonable results, yet the Bayesian estimated coring functions have approximately $2dB$ more SNR gain when using the Lenna image. The exact SNR gains are shown in Table 6.2.

6.4 Comparison with Adelson and Ogden's Coring

In order to get a sense of the extent to which the coring formulization developed in this thesis works, a comparison was made with Adelson and Ogden's work. Using the parameterized coring function

$$I_{out} = (1 - e^{-(m|I_{in}|)^p})I_{in}^k,$$

the same noisy images were cored (to one and two levels). The parameterization for the function was selected by minimizing the mean square error of the restored image as compared with the original. The same coring function was used on all subbands of

Image	Coring Function	SNR	Cored SNR (one level)	Cored SNR (two level)	SNR Gain (one-level)	SNR Gain (two-level)
Lenna	Lenna	37.508	47.895	49.767	10.387	12.259
Lenna	Einstein	37.508	47.063	48.850	9.555	11.342
Einstein	Einstein	30.985	38.722	40.149	7.737	9.164
Einstein	Lenna	30.985	38.587	39.385	7.602	8.400

Table 6.3: Results of coring when the coring functions are swapped.

the pyramid. Because of the availability of the original image, the coring done with the parameterized coring function should yield a somewhat cleaner image than the Bayesian estimated coring function which only uses histogram information. This was found to be the case. This study did show that the coring done using the theory, does almost as well as a system that takes advantage of much more information. The SNR improvements for each technique are listed in Table 6.2.

6.5 Dependence of Coring on Signal Model

In order to determine the dependence of coring effectiveness on the signal statistics the following experiment was undertaken. Coring functions were generated separately for the Lenna and the Einstein images. Instead of applying the Lenna coring functions to Lenna, they were applied to Einstein, and vice-versa. The purpose of this experiment was to see how dependent the success of coring is on the model of the signal histograms (in this case, the actual histograms were taken to be the models). The results showed a diminished success, yet they indicate that reasonably accurate approximations to the signal histograms will work well. The results are shown in Table 6.3.

6.6 Effect of Different Transforms

Another issue that needed to be addressed is the selection of the best transform. A study of 5,7,9 and 11 tap QMF filters was conducted to examine if one filter performs better than the others. Using the six test images and a fixed noise source of variance

Image	Level	5 Tap MSE	7 Tap MSE	9 Tap MSE	11 Tap MSE
Lenna	1	22.234	22.018	22.566	21.803
Lenna	2	19.254	18.706	18.830	18.081
Einstein	1	27.639	27.734	29.117	28.506
Einstein	2	24.548	24.448	25.857	24.648
Kate	1	24.596	24.491	25.211	24.548
Kate	2	23.265	22.892	22.968	22.447
Paolina	1	23.068	23.097	24.081	23.820
Paolina	2	19.360	19.101	19.339	19.184
Mandrill	1	42.552	42.464	42.037	42.194
Mandrill	2	42.021	41.773	41.021	41.137
Bench	1	39.370	39.096	40.002	39.252
Bench	2	38.410	37.696	39.202	37.504

Table 6.4: Results of coring as filter length varies

$\sigma^2 = 64$, the number of taps in the QMF transforms were varied. The coring functions were calculated for each subband image generated and the coring was performed on one and two levels. The cored subbands were then reconstructed. The results of this study show little difference between filter selection and noise reduction. Table 6.4 shows the mean square error of the cored images as filter lengths are varied.




Figure 6-1: The noisy Lenna image. Fixed noise added to Lenna.




Figure 6-2: The second level diagonal band of the fixed noise image.




Figure 6-3: The second level diagonal band of the noisy Lena image.




Figure 6-4: A histogram for the second level diagonal band of the Lenna image.




Figure 6-5: A histogram for the second level diagonal band of the noise image.

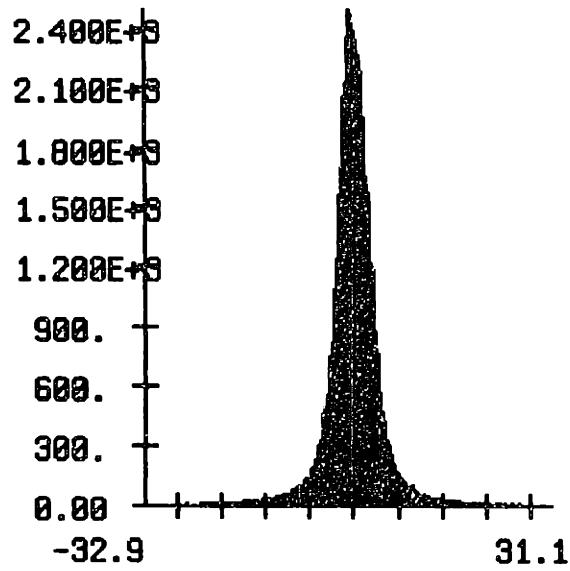


Figure 6-6: A histogram for the second level diagonal band of the noisy image.




Figure 6-7: A coring function for the second level diagonal band.




Figure 6-8: A histogram for the second level cored diagonal band.




Figure 6-9: A second level diagonal band after coring.



Figure 6-10: A restored image (one level coring).



Figure 6-11: A restored image (two level coring).

Chapter 7

A Semi-Automatic Noise Reduction System

Once the theory was satisfactorily demonstrated, image subband histogram models were constructed and tested for success. Eventually, a two-parameter model was selected for its simplicity and effective results in coring. Using this model, a semi-automatic procedure for noise reduction was developed. This procedure assumes the noise is Gaussian and it assumes that the standard deviation is given.

7.1 Image Subband Histogram Models

If enough subband histograms are viewed, it becomes apparent that they all have a characteristic shape that looks something like a laplacian pdf. The idea was developed that if it were possible to construct a model of the clean image subband histograms given only the noisy image, an automatic procedure for removing noise could be developed. This procedure would utilize the theory developed in this thesis to generate a nearly optimal noise reduced image. Although an extremely good two parameter model was developed, it appears that the probabilistic nature of the problem prevents fully automatic estimation of the model parameters, thereby preventing automatic noise reduction. However, if a characterization of the noise is developed, it becomes relatively simple to calculate the model parameters and do a semi-automatic noise

Image	SNR (pre-core)	Cored SNR (denominator)	Cored SNR (numerator)	SNR Gain (denominator)	SNR Gain (numerator)
Lenna	37.508	48.125	50.019	10.616	12.510
Einstein	30.985	38.760	40.217	7.775	9.233
Kate	36.247	45.967	46.352	9.721	10.105
Paolina	31.871	42.249	43.253	10.378	11.382
Mandrill	31.043	35.162	35.125	4.120	4.083
Bench	42.429	47.562	47.115	5.133	4.686

Table 7.1: Results of coring when the model error is minimized with respect to the numerator or the denominator of the coring function expression.

reduction.

Before discussing the various attempts at forming a model for the subband pdf's, it is necessary to elucidate an important point about how to select a model. Typically when fitting a model to a pdf, it would be intuitive to perform a mean square error fit to the pdf. However, in the context of the coring problem this does not provide optimal results. Early on in the model fitting process, disappointing results were achieved relative the Adelson-Ogden coring function. At that time, model pdf's were being fit to the histogram pdf's of the original image and then coring functions were generated using these models. Minimizing the MSE difference of the model pdf with the original pdf essentially minimizes the error in the denominator of the coring function expression. It turns out, that it is more critical to minimize the error in the numerator portion of the expression. Table 7.1 gives noise reduction results when the model is fit to the histogram and when the model is fit to the histogram multiplied by the intensity. The table includes only the data for a two level coring process. The data for one level coring manifest the same result. The data demonstrates that fitting the numerator expression does better than fitting the denominator. This is equivalent to fitting a model to $xP_x(x)$ where x is the intensity and $P_x(x)$ is the histogram for the subband image. A mathematical justification for this discovery has not yet been derived.

Three coring function parameterizations were developed. The first model is a two

parameters model of the form

$$P_{X_1}(X) = \frac{1}{2\tau^2\Gamma(\frac{1}{p})}e^{-|X/\tau|^p}$$

where Γ is the gamma function. This is the model that was eventually selected for its simplicity and effectiveness. Another model that was attempted was a sum of two exponentials of the form above. A weighting constant A is necessary, resulting in a five parameter model

$$P_{X_2}(X) = \frac{A}{2\tau_1^2\Gamma(\frac{1}{p_1})}e^{-|X/\tau_1|^{p_1}} + \frac{(1-A)}{2\tau_2^2\Gamma(\frac{1}{p_2})}e^{-|X/\tau_2|^{p_2}}$$

The third model is of the form

$$P_{X_3}(X) = C + Ae^{-|X/\tau|^p}$$

where A is a normalization factor and not a free parameter of the model. All models were fit to minimize the numerator of the coring function expression.

The set of six test images used earlier was used for this test. Coring was done to one and two levels. The models above were fit to the original image histograms (normalized) and a smooth approximation to a Gaussian of variance $\sigma^2 = 64$ was constructed. Using these models of the subband images and noise, the coring functions were computed. All models performed roughly equivalently in a mean square error sense, all of them achieving noise reduction on par with the original image statistics (see Table 7.2). The third model however tended to give visually unpleasing results, with thin lines of uncored noise. The other two models performed equally well relative to visual inspection. Because the first model involved only two parameters, and the second five, the first model was chosen to represent the subband histograms. This model also lends itself to simple calculations of its central moments. This property is beneficial for a semi-automatic noise reduction system.

Image	SNR	SNR Gain (Model 1)	SNR Gain (Model 2)	SNR Gain (Model 3)
Lenna	37.508	12.510	12.412	11.984
Einstein	30.985	9.233	9.227	9.226
Kate	36.247	10.105	10.158	10.183
Paolina	31.871	11.382	11.674	11.462
Mandril	31.043	4.083	4.146	3.369
Bench	42.429	4.686	5.058	5.140

Table 7.2: A comparison of various models' SNR Gain.

7.2 The Noise Reduction System

With the image model chosen in the previous section, it becomes possible to calculate, given a statistical characterization of the noise, the model parameters for each subband histogram. Examining the contributions to the central moments of the noisy image, it is possible to back-calculate the central moments of the image histogram. Assuming the noise and the signal are uncorrelated,

$$\sigma_{noisy}^2 = \sigma_{noise}^2 + \sigma_{signal}^2$$

and

$$\mu_{noisy}^4 = \mu_{noise}^4 + \mu_{signal}^4$$

where σ^2 is the variance and μ^4 is the fourth central moment. For the two-parameter model chosen in the previous section, the relationships σ_{signal} , μ_{signal} , τ , and p are straightforward.

$$\sigma_{signal}^2 = \tau^2 \frac{\Gamma(\frac{3}{p})}{\Gamma(\frac{1}{p})}$$

and

$$\mu_{signal}^4 = \tau^4 \frac{\Gamma(\frac{5}{p})}{\Gamma(\frac{1}{p})}$$

where Γ is the gamma function. These equations can be solved for τ and p , given the noisy image and a characterization of the central moments of the noise. A recursive

Image	SNR	SNR Gain (Histograms)	SNR Gain (Model)	SNR Gain (Semi-automatic)
Lenna	37.508	12.259	12.510	12.244
Einstein	30.985	9.164	9.233	9.274
Kate	36.247	10.001	10.105	10.508
Paolina	31.871	11.678	11.382	11.845
Mandrill	31.043	4.152	4.083	4.371
Bench	42.429	4.950	4.686	4.992

Table 7.3: A comparison of results when using the original image histograms, a two parameter fit model, and the semi-automatic method.

method such as Newton's method can be used.

Code was written in OBVIUS [8], the Object Based Vision and Image Understanding System, to calculate the model parameters, calculate the coring functions, core the noisy image and reconstruct it. The code achieved results on par with those achieved by using the original image statistics. The results are shown in Table 7.3. A sample result using the Lenna image is shown in Figure 7-1.



Figure 7-1: A restored image using the semi-automatic procedure (two level coring).

Chapter 8

Conclusion

The technique of coring to achieve noise reduction has recently developed. It has been found in previous research that coring with oriented transforms improve noise reduction with relatively less loss of edge information. This thesis uses QMF transforms as an oriented transform. Most importantly, this thesis proposes a mathematical formulation for the construction of a coring function. The formula is verified to assure that it works. Studies are undertaken to test the dependence of successful coring on accuracy in modeling of the subband histograms. It is demonstrated that the dependence is slight, and that it would be worthwhile to construct a model of the histograms and see how well it works. Three models are then attempted and they are found to produce nearly equal results in mean square error. A three parameter model produces slightly less pleasing visual results. Because of its simplicity, a two parameter model is selected. A technique is then developed to semi-automatically clean a noisy image. This technique depends on having a characterization of the noise.

The research done in this paper opens the doors to many other paths of investigation. It would be worthwhile to analyze precisely what properties are desirable for the transform to be used. Clearly an oriented transform is best, yet no mathematical formulation of this empirical result has been developed. It would be worthwhile to try other QMF structures such as hexagonal QMFs[13, 12, 3, 14]

Another area to further study is the possibility to perform a blind noise reduction. Although theoretically possible using the models developed in this paper, the practical

issues of dealing with probabilities impedes this. There are two possible alternatives to the method of matching moments. A noisy image can be subdivided into small patches and the variance for each patch can be measured. Take the minimum variance out of the set of patches. Assuming some part of the original image has a smooth texture, this should yield a reasonable approximation of the noise variance. Another possibility is to examine the high frequency spectrum. Typically image energy is distributed in a $1/f$ fashion. White noise has a flat frequency spectrum. Thus, at high frequency it can be assumed that all the energy is noise energy, and this energy should be proportional to the variance of the noise. It may be possible to use this idea to guess the noise statistics. Being able to automatically clean noise out of a noisy image is a desirable goal.

A third and final area of research to pursue is the possibility of extending the formulation given for coring to such issues as blurring. Initial attempts at so doing have been made, but have not yet been successful. There is an important theoretical consideration for de-blurring that differs from noise reduction. The success of coring depends on the extent to which the noisy image subbands are uncorrelated. However, in blurring they will be strongly correlated. Thus the simple analysis done in this thesis may or may not be applicable.

This thesis presents an application of Bayesian analysis to the specific problem of noise reduction. A Bayesian estimate was made for coring functions, which performed well in noise reduction. Based on this success of this research, it seems likely that other problems could be addressed using a Bayesian approach.

Appendix A

Appendix

A.1 QMF Transform Tap Values

QMF transform tap values are symmetric around the center tap. Included in this appendix are the low-pass and high-pass kernels.

Tap Values					
Number of Taps					
Kernel	5	7	9	11	13
Low-Pass				3.96839E-4	-0.01456
			0.01995	0.01726	0.02165
		-0.005251	-0.04271	-0.03947	-0.09800
	-0.05381	-0.051776	-0.05224	-0.05178	-0.05783
	0.25000	0.255251	0.29271	0.28907	0.42995
	0.60762	0.603553	0.56458	0.56904	0.77371
	0.2500	0.255251	0.29271	0.28907	0.42995
	-0.05381	-0.051776	-0.05224	-0.05178	-0.05783
		-0.005251	-0.04271	-0.03947	-0.09800
			0.01995	0.01726	0.03905
High-Pass				3.96839E-4	-0.01456
			0.01995	-0.01726	-0.02165
		-0.005251	0.04271	-0.03947	0.03905
	-0.05381	0.051776	-0.05224	0.05178	0.09800
	-0.25000	0.255251	-0.29271	0.28907	-0.05783
	0.60762	-0.603553	0.56458	-0.56904	-0.42995
	-0.25000	0.255251	-0.29271	0.28907	0.77371
	-0.05381	0.051776	-0.05224	0.05178	-0.42995
		-0.005251	0.04271	-0.03947	-0.05783
			0.01995	-0.01726	0.09800
			3.96839E-4	-0.02165	
				-0.01456	

Table A.1: Tap Values used for Construction of QMF Pyramids

Bibliography

- [1] E. H. Adelson, C. H. Anderson, J. R. Bergen, P. J. Burt, and J. M. Ogden. Pyramid methods in image processing. *RCA Engineer*, 29(6):33–41, November/December 1984.
- [2] Edward H. Adelson, Eero Simoncelli, and Rajesh Hingorani. Orthogonal pyramid transforms for image coding. In *Proceedings of SPIE*, volume 845, pages 50–58, Cambridge, MA, October 1987.
- [3] Edward H. Adelson, Eero P. Simoncelli, and William T. Freeman. Pyramids and multiscale representations. In *Proc. European Conference on Visual Perception*, Paris, August 1990.
- [4] E.H. Adelson and J.M. Ogden. Computer simulations of oriented multiple spatial frequency band coring. Technical Report Company Private Technical Report PRRL-85-TR-013, RCA David Sarnoff Research Center, 1985.
- [5] P. J. Burt and Edward H. Adelson. The Laplacian pyramid as a compact image code. *IEEE Trans. Communications*, COM-31(4):532–540, April 1983.
- [6] Rafael C. Gonzalez. *Digital Image Processing*. Applied Mathematics and Computation. Addison-Wesley Publishing Company, Inc., Reading, MA, 1979.
- [7] Ernest L. Hall. *Computer Image Processing and Recognition*. Computer Science and Applied Mathematics. Academic Press, New York, NY, 1979.

- [8] David J. Heeger, Eero P. Simoncelli, and Michael Sokolov. Obvius: Object-based vision and image understanding system. Vision Science Technical Report 121, MIT Media Laboratory, April 1989.
- [9] W. Anthony Lee. Coring on oriented bandpassed images, January 1988. Senior Independent Project.
- [10] Jae S. Lim. *Two-Dimensional Signal and Image Processing*. Signal Processing. Prentice-Hall, Inc., Englewood Cliffs, NJ, 1990.
- [11] P.G. Powell and B.E. Bayer. A method for the digital enhancement of unsharp, grainy photographic images. In *Intl. Conf. on Electronic Image Processing, IEE no. 214*, page 179, 1982.
- [12] E Simoncelli and E H Adelson. Multi-scale image transforms. Technical Report 101, Vision Science Group, Media Lab, 1988.
- [13] Eero P. Simoncelli and Edward H. Adelson. Non-separable extensions of quadrature mirror filters to multiple dimensions. *Proceedings of the IEEE: Special Issue on Multi-dimensional Signal Processing*, April 1990.
- [14] Eero P. Simoncelli and Edward H. Adelson. Subband transforms. In John W. Woods, editor, *Subband Image Coding*, chapter 4. Kluwer Academic Publishers, Norwell, MA, 1990.