UC Davis

IDAV Publications

Title

Progressive, Perceptually Transparent Coder for Very High Quality Images

Permalink

<https://escholarship.org/uc/item/57n4s60n>

Journal

In Applications of Digital Image Processing, 2298

Authors

Algazi, Ralph
Ford, Gary
Najmi, A.
et al.

Publication Date

1994

Peer reviewed

Progressive perceptually transparent coder for very high quality images

V. R. Algazi, G. E. Ford, R. R. Estes Jr., A. El-Fallah, and A. Najmi

CIPIC

Center for Image Processing and Integrated Computing
University of California, Davis

ABSTRACT

In the perceptually transparent coding of images, we use representation and quantization strategies that exploit properties of human perception to obtain an approximate digital image indistinguishable from the original. This image is then encoded in an error free manner. The resulting coders have better performance than error free coding for a comparable quality. Further, by considering changes to images that do not produce perceptible distortion, we identify image characteristics onerous for the encoder, but perceptually unimportant. One such characteristic is the typical noise level, often imperceptible, encountered in still images. Thus, we consider adaptive noise removal to improve coder performance, without perceptible degradation of quality. In this paper, several elements contribute to coding efficiency while preserving image quality: adaptive noise removal, additive decomposition of the image with a high activity remainder, coarse quantization of the remainder, progressive representation of the remainder, using bilinear or directional interpolation methods, and efficient encoding of the sparse remainder.

The overall coding performance improvement due to noise removal and the use of a progressive code is about 18%, as compared to our previous results for perceptually transparent coders. The compression ratio for a set of nine test images is 3.72 for no perceptible loss of quality.

Keywords: Image coding, data compression, high quality imaging, noise removal, perceptual coding.

1 INTRODUCTION

Most image coding techniques are directed to the efficient digital representation of original images of moderate quality. The quality of the original image serves as an implicit measure of the additional tolerable distortion that the coder may introduce. For very high quality and super high definition images, a more suitable approach is to attempt to achieve an imperceptible quality degradation in coding. We refer to such an approach as *Perceptually Transparent Coding* since no change from the original image can be perceived.

We have proposed a framework for perceptually transparent coding of images that first identifies simple image representation and quantization strategies that directly exploit properties of human perception to insure that the image is approximated by a digital image which is indistinguishable from the original.^{2,3} The image is then encoded in an error free manner. This approach has several beneficial characteristics, in addition to resulting in coders with substantially better performance than error free coding for a comparable quality. By considering first the changes that can be made to images that do not produce perceptible distortion, we identify image parameters or characteristics that are onerous for the encoder, but that are perceptually unimportant.

In this paper, we consider adaptive noise removal that will decrease the noise without perceptible degradation of quality. We follow this noise reduction with an additive decomposition of the image which directly exploits

	lena	lynda	bldg	cameraman	SHD
Variance	3.7	1.1	1.5	4.8	2.0
Entropy bits	2.9	2.2	2.4	3.2	2.1
PSNR	42.4	47.7	46.3	41.3	45.1

Table 1: Noise analysis of sample images.

the properties of human perception. Finally, we consider a pyramidal representation for the resulting quantized remainder image. In this progressive representation, higher resolution images are estimated from the lower resolution images using either bilinear interpolation, or a direction interpolation strategy based on image analysis which we have developed. At each stage, only the error images need to be transmitted.

2 ADAPTIVE NOISE REDUCTION

Images commonly used for the evaluation of coding algorithms are quite noisy. To quantify this statement, we have analyzed several images, including an image from the Super High Definition (SHD) image test set provided by Nippon Telegraph and Telephone (NTT).

The analysis was performed by locating in each image, a relatively flat region of sufficient extent to allow statistical analysis. These regions generally exhibit a slow trend or shading, which we removed by subtracting a 3×3 running average. The residual noise was then quantized to integers and analyzed — for all images, we found that the noise is additive with approximately constant variance.

We show, in Table 1, the results obtained for some commonly used images: lena, lynda, bldg, and cameraman, and one SHD image. The SHD is 4000×4000 pixels, and was acquired with 12 bits of resolution. Only the upper 8 bits were used in this analysis. No gamma correction was done.

Observe that the noise variance and the corresponding entropy are quite high. For the SHD image, the PSNR is 45 dB. For this PSNR, we predict an entropy of 2 bits/pixel for additive Gaussian noise.⁶ This value was measured experimentally and the histogram of the residual noise confirms its Gaussian behavior.

Such noise has a large effect on the performance of coders at high quality levels.⁶ Thus, we apply noise reduction methods which maintain the perceived quality of the image while reducing the noise. We have found that adaptive noise reduction using anisotropic diffusion can substantially reduce the noise while maintaining the structured image details, important in the perception of image quality.^{1,8}

In adaptive noise reduction, an interactive data dependent filtering algorithm is used. It can be shown that filtering with a family of Gaussian filter kernels $G(x, y, t)$ with variance parameter t , i.e.

$$I(x, y, t) = I(x, y) * G(x, y, t), \quad (1)$$

is equivalent to the partial differential diffusion equation

$$I_t = c \nabla^2 I = c(I_{xx} + I_{yy}), \quad (2)$$

where the subscripts denote partial derivatives, and ∇^2 is the Laplacian. In anisotropic diffusion, we allow the conduction coefficient, $C(x, y, t)$, to vary with respect to space and time, so that

$$I_t = C(x, y, t) \nabla^2 I + \nabla C \cdot \nabla I = \nabla \cdot [C(x, y, t) \nabla I], \quad (3)$$

where ∇ represents the gradient operation and $\nabla \cdot$, the divergence. Typically, we take $C = g(\nabla I)$, where g is a nonlinear function to be specified. In this work, we use adaptively scaled mean curvature diffusion (MCD)⁷ by

Figure 1: Differential Quantization

choosing

$$C = g(\nabla I) = \frac{1}{\sqrt{1 + A^2 |\nabla I|^2}}, \quad (4)$$

where A is a scaling parameter. In this case, it can be shown⁷ that the local rate of diffusion is equal to twice the mean curvature of the image surface about each pixel.

This leads to a very effective iterative adaptive noise reduction technique that preserves image structure, characterized by regions of consistently high gradients, and substantially reduces independent, random noise. The Sobel operator is used for gradient estimation.

Using four iterations of the adaptively scaled MCD algorithm, we get more than 10 dB of noise reduction in the flat portions of the image, while introducing no perceptible change in the resulting image — as long as the PSNR of the original image is at least 46 dB. However, the performance of the coder is substantially improved, as discussed below. Note that, for noisy images, adaptive noise removal may actually *improve* the image quality.

3 DIFFERENTIAL QUANTIZATION^{2,3}

There are three properties of human visual perception which can be used to achieve high image quality while reducing the information content, or bit rate. These are the nonlinear perception of luminance according to Weber's law, the very substantial decrease of the contrast sensitivity for spatial-frequencies above 8 cycles/degree, and the visual masking of perturbations or errors by the activity in the image. Most common visual artifacts encountered in images encoded by current techniques occur in the vicinity of high contrast edges, or near the transition between image regions, and are caused by inadequate control of the spatial distribution of errors. Artifacts with a spatial structure, such as contouring, or the end of block effect in transform coders and vector quantization (VQ) are quite perceptible and highly objectionable. Differential quantization circumvents this problem by providing an excellent approximation in the flat portions of an image. The approximation is not as good near edges or in active areas, but in these regions, visual masking, which extends over several minutes of solid angle, allows for substantial errors to occur below the visual threshold of perception. A diagram of the differential quantization approach which exploits these visual properties is shown in Figure 1.

In this scheme, we compute a low frequency approximation to the original image, and then exploit the properties of this approximation by coarsely quantizing the difference between this approximation and the original image. That is, given an original $M \times N$ image, $I(m, n)$, we compute an approximation, $\hat{I}(m, n)$, which can be reconstructed from a smaller, subsampled image, $I_1(j, k)$.^{*} We then coarsely quantize the difference $I_2(m, n) = I(m, n) - \hat{I}(m, n)$ to obtain the quantized remainder image, $\hat{I}_2(m, n)$, which we use as an alternate representation,

^{*}In general, the approximation does not have to be in the form of an image, but typically, it is.

$\{I_1, \hat{I}_2\}$, of the original image.

In this paper, we use a spline based low frequency approximation, computed from an 8×8 subsampled version of the original, which, when coupled with a 45 level non uniform quantizer results in images which are perceptually indistinguishable from the original.

3.1 Spline based approximation

Consider the image to be a surface in 3 dimensional space. We sample the input, $M \times N$ image, $I(m, n)$, on a rectangular grid to obtain the subsampled, $J \times K$ image, $I_1(j, k)$, where $J = \lceil M/S \rceil$, $K = \lceil N/S \rceil$, and S is the subsampling factor. Using bicubic spline patches, we then determine, from the $\{I_1(j, k)\}$, a smooth surface $\hat{I}(x, y)$, that interpolates these sampled values. We obtain, in fact, a set of bicubic spline patches, which are given by:

$$\hat{I}_{jk}(x, y) = \sum_{p, q=0}^3 \gamma_{jkpq} (x-j)^p (y-k)^q \quad (5)$$

for each rectangular region

$$R_{jk} = \{(x, y) : j \leq x \leq j+1, k \leq y \leq k+1\}. \quad (6)$$

The whole image can thus be interpolated with the requirement that $\hat{I}(x, y)$ be continuous at the patch boundaries.

The set of 16 coefficients γ_{jkpq} , which determine the bicubic spline patches, are computed from the subsampled values, $\{I_1(j, k)\}$, as the solution of a set of linear equations. The approximation, $\hat{I}(m, n)$, of the original image is obtained by subsampling $\hat{I}(x, y)$ at $x = m/S$ and $y = n/S$, where $S = 8$ is the subsampling factor that we have chosen to exploit visual masking of errors in the active areas of the image.

Alternatives to the use of bicubic splines are linear splines on a rectangular or quincunx sampling grid. We have found them slightly inferior for our purpose. Note that linear splines have been used in combination with VQ in a fairly low bit rate encoding scheme.⁹ Another alternative is to use an FIR approximation to an ideal low pass filter both prior to sampling and for interpolation. We have found this approach unsatisfactory because of the serious visual artifacts caused by ideal low pass filtering and interpolation of images.¹⁰ Because we reconstruct an approximate image, $\hat{I}(m, n)$, by interpolation of a subsampled version, low pass filtering of the image prior to subsampling has also been considered. Since the remainder image, $I_2(m, n)$, is obtained by taking the difference between the original an approximation of it, the purpose of preprocessing is not to reduce aliasing errors, but principally to reduce the dynamic range of the remainder image and thus, possibly, to reduce its entropy. In some cases, preprocessing leads to a net gain, but for reasons mentioned in Section 4.3, it is not used in this work.

The subsampled array $\{I_1(j, k)\}$, from which we compute the spline approximation, $\hat{I}(m, n)$ is represented with an 8 bits accuracy and is not encoded, i.e. the upper error free encoder of Figure 1 is not used. For an 8×8 subsampling grid, this adds 0.125 bit per pixel to the overall bit rate of the code.

3.2 Non uniform quantization and perceptual transparency

The decomposition discussed in the previous section results in a reduced entropy because the remainder image has a significantly lower variance than the original image. However, the number of quantization levels for the remainder remains high. In fact the dynamic range of the remainder is actually more than 8 bits. However, the principal advantage of the differential quantization scheme comes from our ability to exploit the characteristics of this remainder image.

3.2.1 Luminance and brightness

It is well known that humans do not distinguish 256 shades of gray. The use of 8 bit gray scale images is due to the fact that at low luminance levels, the just noticeable difference (JND) in luminance is approximately 1/256.

Further, any deviation in the mapping from the quantized signal to the brightness of the display will result in perceptible contouring in the low frequency subareas of the image if fewer quantization levels are used.

The number of gray levels in our scheme is always greater than 256 because of the additive contributions of both the (continuous) spline approximation and the (quantized) remainder. Further, the spline approximation is best in the low frequency subareas of the images. Therefore, coarser quantization of the remainder now becomes feasible.

3.2.2 Visual masking

It is known that errors in images are substantially less visible in active portions of the image.¹² The phenomenon of visual masking by image activity is generally difficult to exploit in image coding since it requires some analysis of the image. Here, since the remainder is the difference between a smooth approximation and the original, large values in the remainder correspond to the most active portions of the image, for which visual masking will be significant. It has been determined that visual masking, at a viewing distance of 6 times picture height, will occur a distance of up to six to seven pixels from a sharp transition. This suggests the use of a 8×8 subsampling grid, so that the maximum distance from the two dimensional grid of subsamples is less than six pixels.

3.2.3 Non uniform quantization

Based on the above considerations, we devised a non uniform quantization scheme for the remainder. Such a scheme provides fine quantization for the low frequency subareas, where the remainder is small, and allows for coarse quantization in the active areas in the image.

We designed a minimum mean square error quantizer which provides a desired non uniform characteristic based on the first order probability density function of the remainder. The number of quantization levels was progressively decreased until we reached the threshold of perception for the quantization error.

We find that, for all images, quantization of the remainder to approximately 5 bits, or 32 levels, is sufficient to ensure perceptual transparency. Further work, based on these statistically based quantizers, results in a *universal* non uniform quantizer which assures perceptual transparency for *all* images. This non uniform quantizer has 31 levels in the range of -115 to $+115$,² and uses a uniform quantizer with step size 20 for larger errors. The theoretical maximum number of levels needed by this quantizer is then 45.

4 ERROR FREE ENCODING OF THE QUANTIZED REMAINDER

In the previous section, we have justified the use of differential quantization using a non uniform quantizer as a means for generating perceptually transparent codes, but we have not dealt with the issue of compressing the resulting, quantized remainder image. That is the topic of this section.

In the proposed scheme, the non uniform quantization is as coarse as possible while still maintaining perceptual transparency. Therefore, to avoid additional and generally uncontrollable image degradation that may result by further quantization, all subsequent coding of the quantized remainder is error free.

The quantized remainder image is just an 45 level gray scale image, and thus we can encode it with any lossless encoding method available for gray scale images. As luck, or fate, would have it, there are not many choices. Our first attempts used DPCM and, therefore, it will be presented first. In this work, we consider the use of a progressive strategy based on hierarchical pyramids and an interpolation strategy which has not been presented before. Keep in mind, that for the remainder of this section, we will be only be concerned with the error free coding of the remainder image.

4.1 DPCM

Note that, in our situation, the prediction error is not guaranteed to fall on one of the allowed quantization levels and, even if the prediction is quantized similarly, the difference between the two non uniformly quantized values does not lie on an allowed quantization level. In short, an error free encoding strategy cannot be developed in such a fashion. We can however, enumerate the quantization levels, take differences between these indices, and develop an error free DPCM encoding strategy. Further, the sums in the encoder and decoder can be taken modulo 45 (the number of levels in the non uniform quantizer) without affecting the error free nature of the code. This allows us to represent the difference between indices with 45 levels instead of 89, and leads to a significant coding gain.

The DPCM predictor used in our experiments is shown in Figure 2, while an overview of the entire spline/DPCM encoding scheme is shown in Figure 3, where we have chosen not to perform preprocessing due, partially, to the fact we are encoding an image which has already been filtered using MCD to remove excess noise.

Figure 2: DPCM kernel used to encode the quantized remainder image.

4.2 Interpolation based pyramids

We now consider a progressive representation of the quantized remainder, and the reconstruction — by interpolation — of the highest resolution image from lower resolution subimages. The basic scheme is shown in Figure 4, where we show only 2 stages of the pyramid and haven’t included the error free encoding that occurs between the encoder and decoder. In the experiments conducted, we used anywhere from 1–6 stages of interpolation and error encoding.

The interpolation process is a key to the performance of such a technique. Driven by the encouraging results obtained by Najmi,¹¹ where noise filtering followed by a hierarchical pyramid led to good compression ratios at high, but not perceptually transparent, quality levels,[†] we set out to determine how such a technique would work in an error free setting. Such a technique *may* do significantly better than DPCM since non causal interpolative prediction contexts are more informative than causal DPCM predictors, and our analysis based directional interpolation algorithm results in improved predictions.

Some of our directional interpolation work has been presented previously,⁴ and was further refined in.¹¹ The main observation made in the latter is that a blended interpolation strategy works better than the threshold based strategy that was used previously. In fact, directional interpolation, as we had previously described it, is just one (important) component of our new strategy (see Figure 5). For this work, we extended the directional interpolation technique to include quincunx pyramids, interpolate using a piecewise linear profile approximation along the low frequency direction at each pixel, and use a new, more robust method for averaging angles and determining the confidence of these estimates. Brevity restricts us from presenting the details here, and limits our discussion to some of the more important aspects.

Typically, subsampling and interpolation are performed on a rectangular grid, in which case, the smallest symmetric subsampling yields a 4 : 1 reduction in the number of pixels in each stage of the pyramid. As an alternative, quincunx sampling only reduces the number of pixels by a factor of two at each stage, and we can cascade two stages of quincunx interpolation to get the equivalent of one stage of rectangular interpolation. In additive decomposition based coding, such as this, the quincunx technique has many advantages. Briefly, each pixel is encoded with respect to its four nearest neighbors at the *current* resolution, and the special cases

[†]His technique compressed images to approximately one half the size that JPEG did, at the same quality level.

Figure 3: Spline/DPCM based transparent coding.

Figure 4: Two stage pyramidal code: encoder (top) and decoder (bottom).

Figure 5: Overview of blended, directional interpolation strategy (top) and directional interpolation (bottom).

associated with rectangular interpolation reduce to a single, simple case.[†] Further, the error image is decomposed into two subimages and more pixels will be known by the encoder and can be used to help predict future pixels. In the belief that quincunx pyramids might do significantly better than rectangular ones, we expand on this last comment.

We compare a 1 stage rectangular pyramid with a 2 stage quincunx pyramid, both of which, for simplicity of discussion, are interpolated using bilinear interpolation. Figure 6a shows the rectangular case, i.e. the white high resolution pixels are interpolated from the black low resolution pixels. The remainder is then encoded and

Figure 6: Rectangular versus quincunx pyramids. On the left we have a single stage of the rectangular pyramid, on the right two stages of a quincunx pyramid.

transmitted. In Figure 6b, we perform one stage of quincunx interpolation, again obtaining bilinear estimates for the white pixels. These estimates are *identical* to those obtained in the rectangular case for these pixels, thus the statistics of the resulting errors are identical. Now, in the quincunx technique, we add the first error image to these predicted values to obtain the *original* pixels as shown in black in Figure 6c. We now have the four nearest neighbors at this highest resolution available to interpolate the remaining pixels, whereas, in the rectangular case, there are only half as many pixels upon which to base this interpolation.

In directional interpolation (top of Figure 5), we perform a local analysis of the image to determine the low frequency direction about each high resolution pixel. This analysis starts with a planar fit to each 2×2 neighborhood in the image, from which the gradient angle estimate and confidence are computed. In the next step, angles are estimated for all high resolution pixels using confidence weighted combinations of neighboring estimates. A larger neighborhood is then examined about each low confidence estimate, from which we try to infer its structural orientation. These angle estimates are then fed into a directional interpolation algorithm which exploits them by interpolating along the low frequency direction using a linear approximation of the edge profile. Finally, the directional interpolates are blended with bilinear interpolates, according to our confidence in the directional estimates.

4.3 Error free coding

We have yet to discuss the actual encoding strategy used to encode the error images. Given the sometimes large percentage of zeros in these images, Huffman codes can be inefficient. Further, these zeros exhibit coherency, since the errors are localized to the active areas in the image. To exploit this coherency and make Huffman coding more efficient, we encode the error images using a hybrid strategy which first encodes the position of the non zero pixels using a binary image encoder (the QM-code) with a seven pixel predictor, after which we encode the the non zero pixels with a Huffman code.⁵ This strategy results in an average 4% bit rate reduction, but for the simpler images can be as high as 28%. In addition, the 4096 samples from the image transmitted by the differential quantization stage propagate through the the entire pyramid and, thus, the lowest resolution image in our pyramid consists entirely of zeros, and need not be encoded.

[†]Indeed, two stages of the quincunx pyramid run faster than a single stage of the rectangular pyramid.

Figure 7: Test images. In order, left to right and top to bottom: baboon, bldg, daisys, flowers, lena, lynda, smile, wheel, and wine.

5 RESULTS

In this section, we summarize the results we have obtained with our perceptually transparent coding scheme, with and without adaptive noise removal. For each case, we have encoded the remainder with DPCM (as described in Section 4.1) and four alternate versions of our pyramidal strategy, as described in Section 4.2. The test set consists of the nine 512×512 , 8 bit images shown in Figure 7. The results obtained are shown in Tables 2 and 3.

The column labeled DPCM gives the results when the remainder image is quantized non uniformly and encoded in an error free manner by DPCM. As a reference, error free JPEG yields 4.11 bit/pixel for the bldg image, 3.46 bit/pixel for lynda and 6.4 bit/pixel for baboon. For the hierarchical technique, we consider 3 stage rectangular and 6 stage quincunx pyramids using both bilinear and directional interpolation. Our results for the original image are presented in Table 2. Note that for rectangular subsampling, the directional interpolation is not always better than bilinear. For quincunx subsampling directional interpolation is consistently slightly better, but by only about 1%. The performance improvement of the progressive scheme as compared to a one stage DPCM scheme is 4%.

The reason that the quincunx technique does better than the rectangular technique is twofold: first, the neighborhood used to compute angle estimates is symmetric about each pixel to be interpolated, and secondly, the neighborhood used to estimate the edge profile is symmetric about each high resolution pixel and is more localized resulting in better approximations — although we use a 2×2 neighborhood to estimate angles, we sometimes use a 3×2 pixel neighborhood our rectangular based technique to estimate the edge profile; significant curvature of image features occurs at this scale and our approximation of the low frequency contour by its tangent fails. Note that with a good directional interpolation strategy, coarser quantization could be used while

Image	Technique				
	DPCM	3 stage rectangular		6 stage quincunx	
		bilinear	directional	bilinear	directional
baboon	4.344	4.378	4.380	4.282	4.262
bldg	2.715	2.851	2.855	2.725	2.627
daisys	3.426	3.462	3.447	3.357	3.309
flowers	1.981	2.064	2.057	1.973	1.952
lena	3.104	3.116	3.101	3.029	2.991
lynda	1.907	1.870	1.868	1.806	1.795
smile	1.624	1.524	1.524	1.450	1.442
wheel	2.081	2.180	2.166	2.110	2.057
wine	2.275	2.258	2.230	2.175	2.131

Table 2: Transparent coding on the original images.

Image	Technique				
	DPCM	3 stage rectangular		6 stage quincunx	
		bilinear	directional	bilinear	directional
baboon	3.911	3.950	3.944	3.857	3.804
bldg	2.581	2.636	2.640	2.547	2.426
daisys	2.960	3.005	2.988	2.909	2.839
flowers	1.949	1.997	1.991	1.925	1.881
lena	2.299	2.285	2.265	2.220	2.144
lynda	1.674	1.563	1.560	1.512	1.476
smile	1.445	1.129	1.127	1.090	1.077
wheel	1.914	1.753	1.739	1.738	1.662
wine	2.203	2.163	2.133	2.090	2.030

Table 3: Transparent coding on the noise reduced images.

still maintaining perceptual transparency[§].

In Table 3, we show the effect of adaptive noise reduction. For the DPCM scheme, the noise reduction leads to a 11% decrease in bit rate, averaged over all images. The best progressive coding scheme produces a further improvement of 9%. Combining adaptive noise reduction with a progressive code leads to a decrease of 18% in the average bit rate. The average compression rate for such a scheme and for all images is now 3.72 with no perceptible loss of image quality.

6 DISCUSSION

This paper makes two major contributions. The first one is conceptual. It is possible to process and represent images so as to improve their compressibility without loss of image quality. In particular, adaptive noise reduction leads to a substantial increase in compressibility with no visible change in the image. The second contribution is that a progressive code, when used in conjunction with image analysis may also lead to a slight improvement in performance, by contrast to the slight loss of performance generally associated with progressive coding schemes.

[§]This is the focus of on going research.

Note that our approach to perceptually transparent coding, which controls image quality by first introducing imperceptible changes in the image, now requires increased attention to efficient error free coding schemes for quantized gray scale and color images.

7 ACKNOWLEDGMENTS

This research was supported in part by the UC MICRO program, Pacific Bell, Lockheed and Hewlett Packard.

8 REFERENCES

- [1] V. R. Algazi, G. E. Ford, M. Mow, and A. Najmi. Design of subband coders for high quality images, based on perceptual criteria. In *Proceeding of the SPIE Symposium on Applications of Digital Image Processing XVI*, volume 2028, pages 40–6, July 1993.
- [2] V. R. Algazi, E. Maurincommé, and G. E. Ford. Error free and transparent coding of images using approximations by splines. In *Proceedings of the SPIE Symposium on Image Processing Algorithms and Techniques II*, volume 1452, pages 364–70, Boston, MA, November 1991.
- [3] V. R. Algazi, T. R. Reed, G. E. Ford, E. Maurincommé, I. Hussain, and R. Potharlanka. Perceptually transparent coding of still images. *IEICE*, E75-B(5):340–8, March 1992.
- [4] V. Ralph Algazi, Gary E. Ford, and Ravindra Potharlanka. Directional interpolation of images based on visual properties and rank order filtering. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 3005–8, May 1991.
- [5] V. Ralph Algazi, Phillip L. Kelly, and Robert R. Estes. Compression of binary facsimile images by preprocessing and color shrinking. *IEEE Transactions on Communications*, 38(9):1592–8, September 1990.
- [6] V. Ralph Algazi, T. R. Reed, Gary E. Ford, and R. R. Estes. Image analysis for adaptive noise reduction in super high definition image coding. In *Proceedings of the SPIE Symposium on Visual Communications and Image Processing*, volume 1818, Boston, MA, November 1992.
- [7] A. I. El-Fallah and G. E. Ford. Mean curvature evolution and surface area scaling in image filtering. Submitted to the *IEEE Transactions on Image Processing*, June 1994.
- [8] Adel El-Fallah and Gary E. Ford. Image filtering by gradient inverse inhomogeneous diffusion. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, volume V, pages V73–6, April 1993.
- [9] H. M. Hang and B. J. Haskell. Interpolative vector quantization of color images. *IEEE Transactions on Communications*, 36(4):465–70, April 1988.
- [10] T. A. Hentea and V. R. Algazi. Perceptual models and filtering of high-contrast achromatic images. *IEEE Transactions on Systems, Man and Cybernetics*, SMC-14(2):230–45, 1984.
- [11] Azfar Najmi. Adaptive noise filtering with high quality interpolation-based image coding. Master's thesis, University of California, Davis, 1994.
- [12] A. Netravali and B. Haskell. *Digital Pictures*. Plenum Press, 1988. Chapter 4.