Journal of Electronic Imaging 1(3), 277–292 (July 1992).

Optimized error diffusion for image display

Bernd W. Kolpatzik
Charles A. Bouman
Purdue University
School of Electrical Engineering
West Lafayette, Indiana 47907-1285

Abstract. *Displaying natural images on an 8-bit computer monitor requires a substantial reduction of physically distinct colors. Simple minimum mean squared error quantization with 8 levels of red and green and 4 levels of blue yields poor image quality. A powerful means to improve the subjective quality of a quantized image is error diffusion. Error diffusion works by shaping the spectrum of the display error. Considering an image in raster ordering, this is done by adding a weighted sum of previous quantization errors to the current pixel before quantization. These weights form an error diffusion filter. We propose a method to find visually optimized error diffusion filters for monochrome and color image display applications. The design is based on the low-pass characteristic of the contrast sensitivity of the human visual system. The filter is chosen so that a cascade of the quantization system and the observer's visual modulation transfer function yields a whitened error spectrum. The resulting images contain mostly high-frequency components of the display error, which are less noticeable to the viewer. This corresponds well with previously published results about the visibility of halftoning patterns. An informal comparison with other error diffusion algorithms shows less artificial contouring and increased image quality.*

1 Introduction

High-quality color images can be displayed on monitors with 24 bits/pixel by assigning 256 different shades (8 bits) of red, green, and blue to each input pixel. Hence, each pixel color is selected out of $2^{24} \approx 16$ million different colors, and the resulting quantization steps are practically invisible to a human observer.

Paper 92-020 received May 18, 1992; revised manuscript received July 16, 1992; accepted for publication July 17, 1992.
1017-9909/92/$2.00. © 1992 SPIE and IS&T.

However, many low-cost computer and display devices can display only 256 different colors (8 bits) at a time, due to hardware constraints. This requires a substantial reduction of physically distinct colors. Using a simple minimum mean squared error (MMSE) quantizer with 8 levels of red and green and 4 levels of blue yields poor image quality.

Applying error diffusion significantly improves the subjective quality of a quantized image. The basic algorithm was first introduced by Floyd and Steinberg[1] for halftoning in the printing process of gray-scale images. It is based on the observation that the human visual sensitivity to display errors is dependent on spatial frequency.

Floyd and Steinberg proposed an algorithm that calculates the quantization error for each pixel and feeds it forward to four unquantized pixels of the input image. As shown in Refs. 2, 3, and 4, this algorithm is equivalent to a feedback system that adds a weighted sum of four previous quantization errors to the current pixel before it is quantized. Since the weighting factors sum to one, it can be shown that the average value of the quantized image is locally equal to the true gray-scale value.

Billotet-Hoffman and Bryngdahl[5] compared the performance of ordered dithering with error diffusion. They concluded that error diffusion, as proposed by Floyd and Steinberg, yields quantized images that are comparable or superior to most ordered dithering techniques, when nonlinear effects such as dot overlap are not significant. However, the image quality still suffers from visible, correlated artifacts. To reduce artificial contouring, Billotet-Hoffman and Bryngdahl proposed a combination of error diffusion and ordered dither, where a dither matrix is used to vary the threshold level of the quantizer.

Ulichney[6] examined the spectral characteristics of the display error for the error diffusion algorithm using a variety

of feedback filters. He proposed an error diffusion filter with randomized weighting coefficients to shape the display error spectrum so that it would have mostly high-frequency content (''blue noise''). He stated that blue noise is less noticeable to a human observer than errors with a white power spectrum. Goertzel and Thompson[7,8] applied error diffusion with a randomized error diffusion filter to monochrome and color images. It was found that randomized coefficients remove deterministic (iterative) patterns in the displayed image. Although in comparison to the approach with a deterministic filter, the image quality is improved, false contouring artifacts remain.

The objective of this paper is to develop an optimized algorithm for a specific human visual model. Most previous methods use information about the human visual system only indirectly or qualitatively.[2,10–13] To achieve this, a modulation transfer function (MTF) for an overall system in a luminance-chrominance space is used. This system model includes the effects of spatial sampling due to the monitor and a model for the human modulation transfer function. We assume that each part of this model can be described by three decoupled system functions for luminance and both chrominance components. The error diffusion algorithm is then independently matched to the three components of the resulting system MTF. This is done by choosing an optimized error diffusion filter for each component. Design of the optimal error diffusion filter is shown to correspond to an optimum 2-D linear prediction problem.

One important underlying assumption of our error diffusion filter design method is that the quantization error has a white power spectrum. However, in general, this assumption is not valid. It is important to notice that the quantization error is defined as the difference between the input and output of the quantizer, while the display error is the difference between the input and output of the entire quantizing system. For error diffusion applied to gray-scale images, it was found that the whiteness assumption about the quantization error may be locally violated. This often leads to false contouring.

Whiteness of the quantization error spectrum can be assured by combining standard error diffusion with dithered quantization (dithered error diffusion, DED). This method is similar to the one described by Billotet-Hoffmann and Bryngdahl[5] but uses a different dither signal. Our method is based on the fact that the quantization error spectrum can be whitened by adding a fixed amount of white noise to the signal before it is quantized. However, this also increases the display error variance. A further refinement of this algorithm is examined, that adds dithering noise locally, based on a nearest neighbor criterion (locally dithered error diffusion, LDED). With this approach, it is possible to remove false contouring without adding excessive noise to the image.

Simulations were carried out using visual models from Sullivan,[14,15] Näsänen,[16] and Mullen,[17] describing the human MTF and reflecting the low-pass characteristic of the human visual system to changes in luminance and chrominance. The models differ in their cutoff frequencies and their spatial symmetry properties.

Considering a modulation transfer function for the overall system and then designing the error diffusion filter accordingly yields an improvement over the conventional methods of error diffusion. The algorithms were applied to gray-scale images as well as to color images, and their performance was compared. The best results were achieved using the LDED algorithm. False contours are broken up, without distorting the image through excessive amounts of dithering noise.

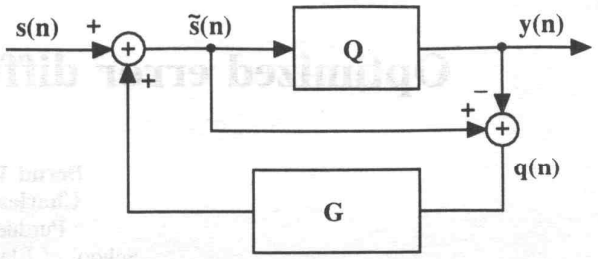


Fig. 1 Block diagram of the basic error diffusion algorithm.

Section 2 describes the error diffusion algorithm for monochrome images and the optimization for a given human MTF. Then we extend our analysis to color images. Section 3 discusses the whitening assumption of the quantization error. Some models of the human MTF are described in Sec. 4, and Sec. 5 contains experimental results.

2 Optimized Error Diffusion

Section 2.1 briefly describes the basic error diffusion algorithm. Section 2.2 develops our approach for designing an optimized error diffusion filter for the luminance component, and Sec. 2.3 extends the algorithm to color images.

2.1 *Basic Error Diffusion*

The basic error diffusion algorithm for monochrome images, as introduced by Floyd and Steinberg,[1] is illustrated in Fig. 1. We will use the analysis given in Refs. 3 and 4 as the basis for our approach. In our simulations, we process the input image in raster order, starting with the pixel in the upper left corner and proceeding line by line to the lower right corner. Let $n=(n_1, n_2)$ denote the pixel location in the image. Referring to Fig. 1, we write $s(n)$ for pixels of the unquantized input image. In general, $s(n)$ will be linearly proportional to image luminance. Since most image data formats incorporate some nonlinear predistortion (e.g., gamma correction), this requires that the data be transformed back to a linear format before processing. This transformation is important since our models for the human visual system and the display assume that the data are proportional to light intensity.

The input to the quantizer, $\tilde{s}(n)$, is computed by adding a sum of weighted, previous quantization errors, $q(n)$, to the current pixel, $s(n)$. This set of weights forms the error diffusion filter G with impulse response $g(n)$. The output image of our quantizing system is $y(n)$. We obtain the display error, $e(n)$, as

$$e(n)=s(n)-y(n) \ .$$

The quantization error $q(n)$ is given by

$$\begin{aligned} q(n) &= \tilde{s}(n)-y(n) \\ &= \sum_{k<n} g(n-k)q(k)+e(n) \ , \end{aligned} \tag{1}$$

Table 1 Coefficients of error diffusion filters. For the filter used by Goertzel, z is a random number, uniformly distributed between 0.25 and 0.75.

	$g_{0,1}$	$g_{1,-1}$	$g_{1,0}$	$g_{1,1}$
Floyd-Steinberg	0.4375	0.1875	0.3125	0.0625
Goertzel	z	0	0	1.0-z
1st order LP (non-scparablc) (luminance)	0.7770	0.009	0.7861	-0.6098
1st order LP (non-separable) (chrominance)	0.8767	0.0359	0.8205	-0.7376

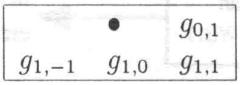


Fig. 2 Location of the filter coefficients.

where $n-k=(n_1-k_1, n_2-k_2)$ and $k<n$ denotes pixels with $(k_1=n_1$ and $k_2<n_2)$ or $k_1<n_1$. We define the two-dimensional Fourier transform of the discrete signal $e(n)$ as in Ref. 18:

$$E(\omega)=\sum_{n\in\mathbb{Z}^2} e(n)\exp[-jn\omega^T] \ ,$$

where $\omega=(\omega_1,\omega_2)$ and $\mathbb{Z}$ is the set of all integers. Then we obtain from Eq. (1) a relationship between the spectrum of the display error $E(\omega)$ and the spectrum of the quantization error $Q(\omega)$,

$$E(\omega)=Q(\omega)[1-G(\omega)] \ , \qquad (2)$$

where $G(\omega)$ is the frequency response of the error diffusion filter. We will refer to $G(\omega)$ and $g(n)$ as the error diffusion filter interchangeably. If we assume that $q(n)$ is spatially uncorrelated, then $Q(\omega)$ is white, and the display error spectrum can be shaped by choosing the frequency response $H(\omega)=1-G(\omega)$. In Sec. 3 we will develop methods to assure that this whiteness assumption holds.

Based on this analysis, the frequency response of the error diffusion filter, $G(\omega)$, is critical in determining the spectrum of the display error. Table 1 lists some common examples of error diffusion filters and Fig. 2 shows the location of the filter coefficients. Floyd and Steinberg proposed a filter, $g(n)$, with four positive, deterministic coefficients. Stucki[22] defined a filter that operates on 12 previous quantization errors. Due to space limitations, it is not listed in the table. The coefficients of most error diffusion filters, including the filters from Floyd and Steinberg, and Stucki,[22] sum to one. Therefore, $G(\omega=0)=1$. It follows from Eq. (2) that the dc component of the display error $E(\omega=0)=0$. This implies that the average gray value of the quantized image, $y(n)$, is locally adjusted to the true gray-scale value of the input image. Ulichney[6] and Goertzel and Thompson[8] examined error diffusion filters with randomized coefficients. In these cases $g(n)$ changes randomly with the location in the image. Table 1 lists the coefficients of a filter, used by Goertzel. The parameter z denotes a random number, uniformly distributed between 0.25 and 0.75 and chosen independently for each pixel. As with the deterministic filters, the sum of the coefficients at any location is also equal to one.

The error diffusion should be designed so that the display error $e(n)$ is the least noticeable to a human observer. Since the human visual system is more sensitive to low-frequency changes in luminance than to high-frequency components, most of the energy of the display error should be shifted to high frequencies. However, forcing the dc component of the display error to zero by setting the sum of the filter coefficients to one is an artificial constraint. Our simulations show that better results can be obtained by tolerating some display error at zero frequency and distributing the error energy over the entire spectral range according to a human visual model. The next section formulates such a minimization criterion.

2.2 *Optimization and Design of the Error Diffusion Filter in Luminance*

To design an optimized error diffusion filter $G(\omega)$ for the luminance component, we consider a model for the overall system, as illustrated in Fig. 3. After the error diffusion algorithm and quantization is applied to the input image, the quantized image is converted by a monitor to a spatially continuous signal. It is filtered with the monitor MTF, $P(\Omega)$, where $\Omega=(\Omega_1,\Omega_2)$ and $\Omega_1,\Omega_2\in(-\infty,\infty)$. This gives rise to aliased terms in the spectrum of the displayed image. The resulting image is seen by a human observer. It passes through the human visual system with MTF, $W(\Omega)$, and becomes the perceived image. Linearity is assumed for each component of our system.

We would like to minimize the energy of the perceived error $E[\varepsilon^2(x)]$ by choosing an optimal error diffusion filter. As illustrated in Fig. 3(a), $\varepsilon(x)$ is the perceived difference between the quantized image $y(n)$ and the original image $s(n)$. Since linearity of the monitor and the human MTF is assumed, it is sufficient to consider the difference signal $s(n)-y(n)$ as the input to our system, as shown in Fig. 3(b). This is the display error $e(n)$ as it was defined in the previous section.

Assuming that the input image can be viewed as a stationary random process, the perceived error $\varepsilon(x)$ is also a stationary random process and the energy of $\varepsilon(x)$ is given by

$$E[\varepsilon^2(x)]=\frac{1}{4\pi^2}\int_{\Omega\in\mathbb{R}^2}\Phi\varepsilon(\Omega)\,d\Omega \ ,$$

where $\Phi_\varepsilon(\Omega)$ is the power spectral density of $\varepsilon(x)$.

In Appendix A, we show that the energy of the perceived error can be expressed in terms of the power spectrum of the display error $\Phi_e(\omega)$ and an equivalent overall system function $\tilde{W}(\omega)$ such that

$$E[\varepsilon^2(x)]=\frac{1}{4\pi^2}\int_{\omega\in\mathbb{P}^2}\Phi_e(\omega)|\tilde{W}(\omega)|^2\,d\omega \ , \qquad (3)$$

where $\mathbb{P}=[-\pi,\pi]$, and the overall system function $\tilde{W}(\omega)$ is given by

$$|\tilde{W}(\omega)|^2=f_s^2\sum_{l\in\mathbb{Z}^2}|P(\omega f_s+2\pi l f_s)W(\omega f_s+2\pi l f_s)|^2 \ ,$$

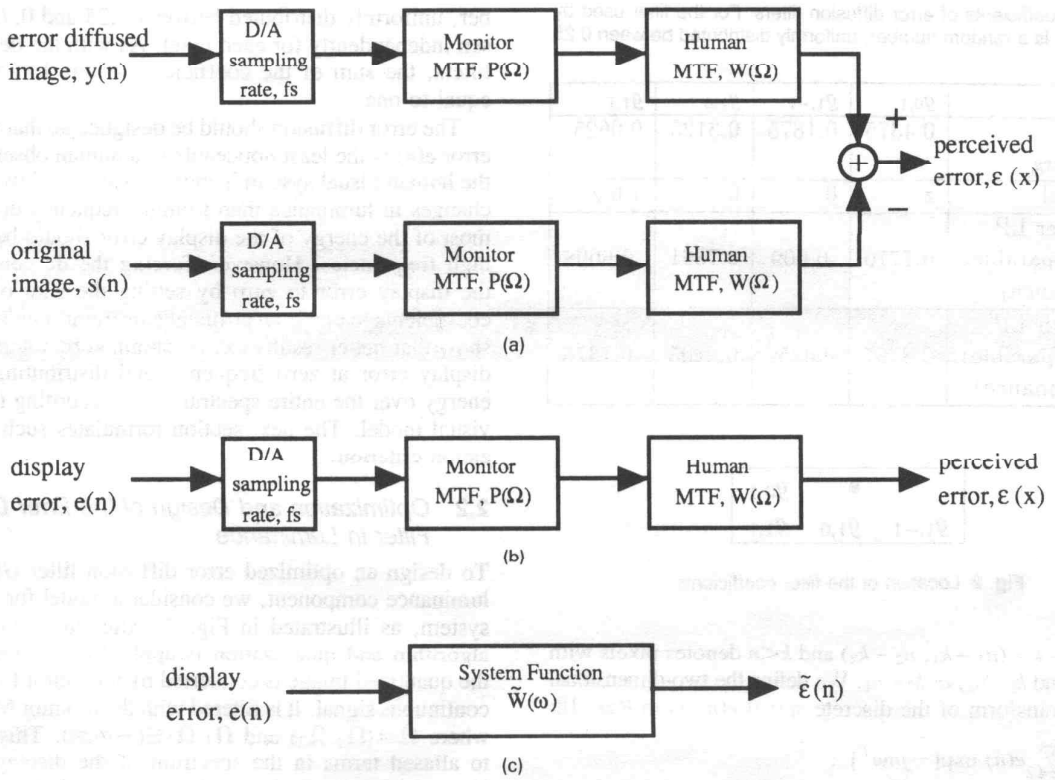


Fig. 3 Block diagram of the monitor and the human visual system.

where $l=(l_1, l_2)$ and f_s is the sampling rate of the monitor in cycles per degree. The sampling rate is obtained by taking the reciprocal of the angle of perception of a single displayed pixel for a fixed viewing distance.

Figure 3(c) illustrates the resulting simplified system model, where $\tilde{\varepsilon}(n)$ is the output of an equivalent discrete system $\tilde{W}(\omega)$. This simplification results from Eq. (3) and the fact that

$$E[\varepsilon^2(x)]=E[\tilde{\varepsilon}^2(n)] \ .$$

Therefore, we can minimize the energy of the perceived error by minimizing the energy of $\tilde{\varepsilon}(n)$. Furthermore, from Eq. (2) it follows that

$$\Phi_e(\omega)=\Phi_q(\omega)|1-G(\omega)|^2 \ .$$

Replacing $\Phi_e(\omega)$ in Eq. (3), we want to minimize

$$E[\tilde{\varepsilon}^2(n)]=\frac{1}{4\pi^2}\int_{\omega\in\mathbb{P}^2}\Phi_q(\omega)|1-G(\omega)|^2|\tilde{W}(\omega)|^2 \, d\omega \ ,$$

subject to the constraint that $G(\omega)$ be a strictly causal filter. Strict causality of $G(\omega)$ is required since it is located in the feedback branch of the error diffusion system.

The optimal choice of $G(\omega)$ has an intuitive interpretation, which is illustrated in Fig. 4. Figure 4(a) shows the quantizing system together with the system $\tilde{W}(\omega)$ and the perceived error $\tilde{\varepsilon}(n)$. By exchanging $\tilde{W}(\omega)$ and $1-G(\omega)$, Fig. 4(b) illustrates that the filter $G(\omega)$ acts as a predictor for the signal $u(n)$. It should be emphasized that the signal $u(n)$ is introduced only to simplify our derivation and does not occur anywhere in the real system.

The problem of minimizing $E[\tilde{\varepsilon}^2(n)]$ is then equivalent to choosing an optimal linear predictor $G(\omega)$ for the stationary random process $u(n)$. This is a classical problem, which may be solved using $\Phi_u(\omega)$, the power spectrum of $u(n)$. If we assume that $q(n)$ is white, then

$$\Phi_u(\omega)=\Phi_q(\omega)|\tilde{W}(\omega)|^2$$
$$=N_0|\tilde{W}(\omega)|^2 \ .$$

Appendix B gives the details of the optimized 2-D linear predictor design for this problem.

2.3 *Extension to Color Images*

When applying error diffusion to color images, the input image is a sequence of vectors with a red (R), green (G), and blue (B) component, such that $s(n)=[r(n),g(n),b(n)]^T$. We will again assume that the three primary components are linearly proportional to light intensity. The extension to color will also require the additional approximation that the quantization errors of the three color components are uncorrelated. This assumption together with an opponent color model for the human visual system will lead to linear prediction equations so that one filter can be designed for each color component. For this work, quantization is performed separately on the RGB components. Other methods of quantization are possible, but separate quantization of RGB components does not lead to gamut mismatch.

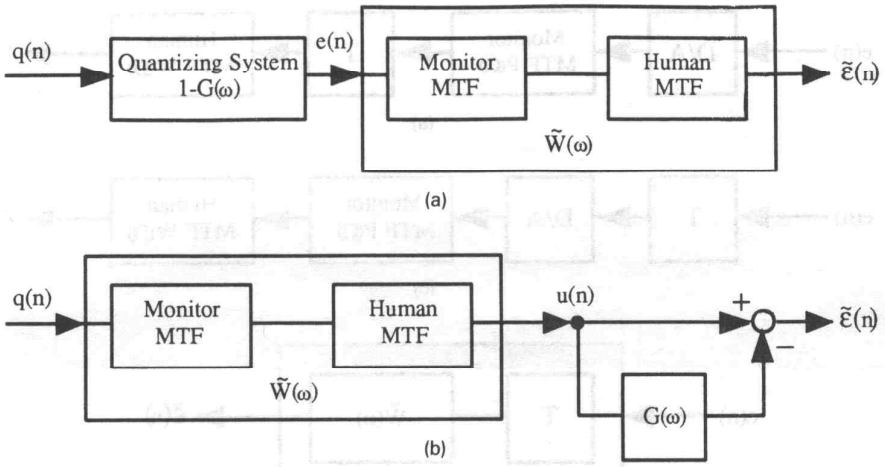


Fig. 4 Block diagram of the overall system.

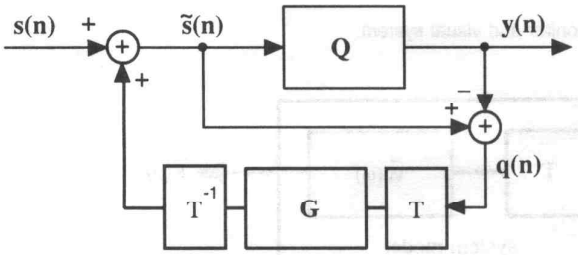


Fig. 5 Block diagram of the error diffusion algorithm applied to color images.

A more natural color space to evaluate perception of the human visual system is spanned by a luminance-chrominance coordinate system. Experimental results[17,19] are available that describe the contrast sensitivity function of the human visual system to variations in luminance, as well as for variations in two chrominance components. For our experiments we use a color coordinate system specified by the primaries in Mullen's paper.[17] We write "Y" for luminance, "rg" for the red-green component, and "yb" for the yellow-blue component. Furthermore, we selected the D_{6500} standard white point as the chromaticity of the luminance component.

Figure 5 illustrates our color error diffusion algorithm, where the error diffusion is performed in the Y,rg,yb space, while the quantization is done in RGB coordinates. We define a transformation matrix $\mathbf{T}$ to convert from the RGB space to Y,rg,yb coordinates. Matrix $\mathbf{T}$ can be decomposed into the product of two matrices $\mathbf{A}$ and $\mathbf{B}$, such that

$$\mathbf{T}=\mathbf{B}^{-1}\mathbf{A} \ ,$$

where $\mathbf{A}$ is a device-dependent matrix, converting from RGB coordinates to standard X,Y,Z coordinates,[20] and $\mathbf{B}$ transforms Y,rg,yb coordinates to the X,Y,Z space. Matrix $\mathbf{A}$ is defined as

$$\mathbf{A}=[\tilde{r},\tilde{g},\tilde{b}] \ ,$$

where $\tilde{r}$, $\tilde{g}$, and $\tilde{b}$ are column vectors containing the X,Y,Z coordinates of the three primaries of the display device. These vectors have to be determined for each display. The transformation matrix $\mathbf{B}$ can be calculated for the wavelengths of the color stimulus described in Ref. 17. We obtained $\mathbf{B}$ as

$$\mathbf{B}=\begin{pmatrix} 0.0670 & -0.1116 & -0.0830 \\ 0.0707 & 0 & 0 \\ 0.0767 & 0.0049 & -1.0000 \end{pmatrix} .$$

The quantization error $q(n)$ is transformed to Y,rg,yb coordinates. Each component of the transformed quantization error is passed through a separate error diffusion filter. Then the filtered error is transformed back to RGB coordinates and added to the next unquantized input pixel.

The analysis of the system is analogous to the case of luminance error diffusion, where signal variables now denote vectors of three color coordinates. The system functions $\mathbf{G}(\omega)$ and $\tilde{\mathbf{W}}(\omega)$ are now matrices of functions of spatial frequency. Including the color transformations $\mathbf{T}$ and $\mathbf{T}^{-1}$ of the feedback loop into the calculation, we obtain a relationship for the display error spectrum in terms of the quantization error spectrum,

$$E(\omega)=[\mathbf{I}-\mathbf{T}^{-1}\mathbf{G}(\omega)\mathbf{T}]Q(\omega) \ ,$$

where $\mathbf{I}$ is the identity matrix and $\mathbf{T}^{-1}$ denotes the inverse transformation from the luminance-chrominance to RGB coordinates.

A model for the monitor and the human visual system is shown in Fig. 6(a). The components of the system are analogous to the monochrome model. A transformation $\mathbf{T}$ is inserted, since $\mathbf{W}(\Omega)$ is given for luminance-chrominance coordinates. We assume that human contrast sensitivity to variations in luminance and chrominance is decoupled, and that $\mathbf{W}(\Omega)$ is therefore diagonal. If we also assume that $P(\Omega)$ and the digital-to-analog conversion are the same for each of R, G, and B, we may exchange the order of the system blocks and obtain the model illustrated in Fig. 6(b).

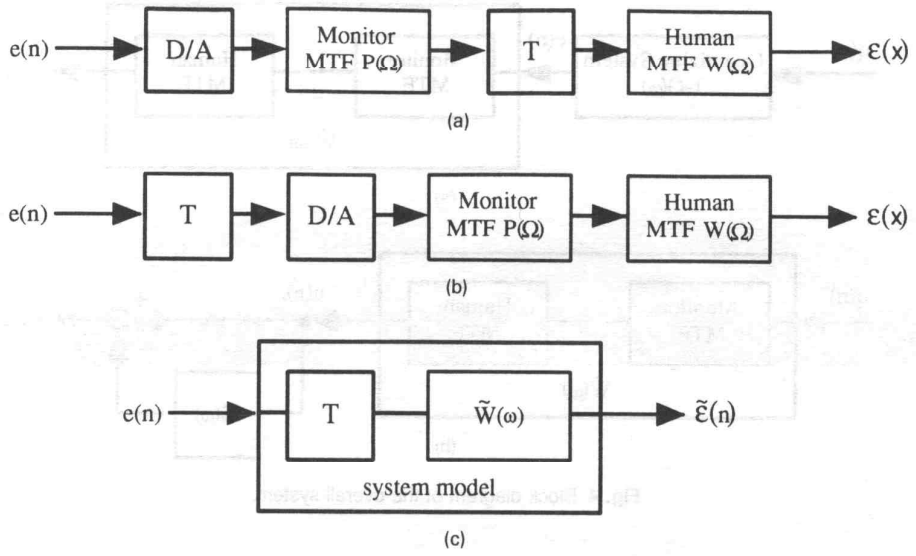


Fig. 6 Block diagram of the monitor and visual system.

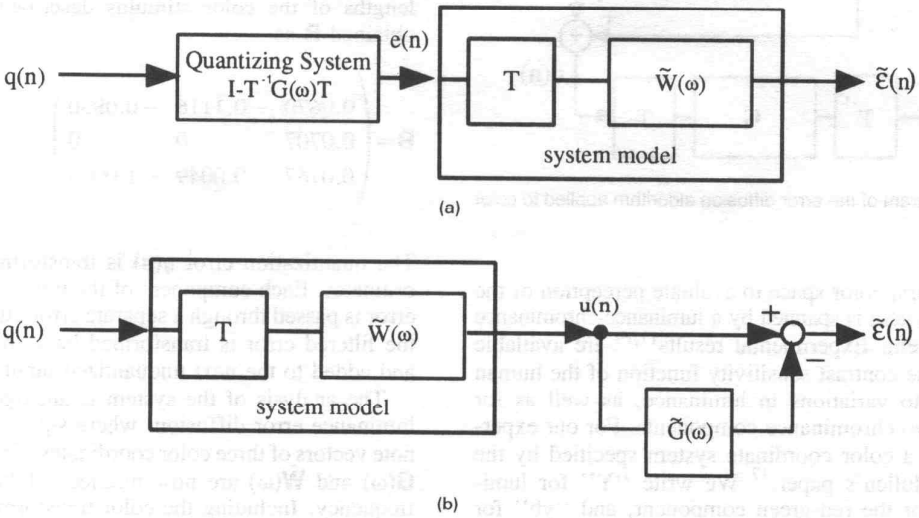


Fig. 7 Block diagram of the overall color system.

As shown earlier for the monochrome model, this system has a discrete equivalent, given in Fig. 6(c), where

$$\tilde{\mathbf{W}}(\omega)\mathbf{T}=\begin{bmatrix}\tilde{W}_Y(\omega) & 0 & 0\\ 0 & \tilde{W}_{rg}(\omega) & 0\\ 0 & 0 & \tilde{W}_{by}(\omega)\end{bmatrix}\mathbf{T}\ .$$

Figure 7(a) shows the overall discrete color system. As illustrated in Fig. 7(b), the spectrum of the perceived error is related to the spectrum of the quantization error by

$$\begin{aligned}\tilde{\varepsilon}(\omega)&=\tilde{\mathbf{W}}(\omega)\mathbf{T}[\mathbf{I}-\mathbf{T}^{-1}\mathbf{G}(\omega)\mathbf{T}]Q(\omega)\\&=\tilde{\mathbf{W}}(\omega)[\mathbf{I}-\mathbf{G}(\omega)]\mathbf{T}Q(\omega)\\&=[\mathbf{I}-\tilde{\mathbf{G}}(\omega)]\tilde{\mathbf{W}}(\omega)\mathbf{T}Q(\omega)\ ,\end{aligned}\tag{4}$$

where

$$\tilde{\mathbf{G}}(\omega)=\tilde{\mathbf{W}}(\omega)\mathbf{G}(\omega)\tilde{\mathbf{W}}^{-1}(\omega)\ .\tag{5}$$

Notice that $\tilde{\mathbf{G}}(\omega)$ should be an optimal predictor for the signal $u(n)$. Therefore, the objective is to minimize the energy of $\tilde{\varepsilon}(n)$ by choosing the filter $\tilde{\mathbf{G}}(\omega)$ and then converting to $\mathbf{G}(\omega)$, using Eq. (5).

Let $\tilde{\mathbf{R}}(m)$ be the autocorrelation matrix of the transformed quantization error $\tilde{q}(n)=\mathbf{T}q(n)$. Then we will assume that

$$\begin{aligned}\tilde{\mathbf{R}}(m)&=E[\tilde{q}(n+m)\tilde{q}^T(n)]\\&=\tilde{\mathbf{R}}_0\delta(m)\ ,\end{aligned}\tag{6}$$

where $\tilde{\mathbf{R}}_0$ is assumed to be a diagonal matrix and $\delta(m)$ is the Dirac delta function, given by

$$\delta(m)=\begin{cases}1 & \text{if } m=0\\ 0 & \text{otherwise}\end{cases}.$$

Again, the spatial independence of $\tilde{q}(n)$ is ensured by the methods described in Sec. 3. In general, the structure of $\tilde{\mathbf{R}}_0$ is difficult to determine when nonuniform quantization is used. If $\tilde{\mathbf{R}}_0$ is not diagonal and the quantization error between components is correlated, we could calculate the optimal $\mathbf{G}(\omega)$ according to a minimum mean square error criterion for the perceived error. However, this is a fairly complex problem, which involves the design of six error diffusion filters to account for the cross-coupling between the quantization error components. In practice, $\tilde{\mathbf{R}}(m)$ is usually unknown, but our experiments indicate that the cross-correlation between the three error components is small, particularly when the quantization error is spatially white. Therefore we will assume that $\tilde{\mathbf{R}}_0(m)$ is diagonal.

Making the assumption of Eq. (6) and that $\tilde{\mathbf{R}}_0$ is diagonal we have that

$$\mathbf{\Phi}_u(\omega)=\begin{bmatrix}\Phi_{u11}(\omega) & 0 & 0\\ 0 & \Phi_{u22}(\omega) & 0\\ 0 & 0 & \Phi_{u33}(\omega)\end{bmatrix}, \tag{7}$$

where

$$\Phi_{uii}(\omega)=\tilde{\mathbf{W}}_{ii}(\omega)\mathbf{R}_{0ii}.$$

By the form of Eq. (7), the three components of $u(n)$ are uncorrelated and the optimal predictor for $u(n)$ is therefore of the form

$$\tilde{\mathbf{G}}(\omega)=\mathbf{G}(\omega)$$

$$=\begin{bmatrix}G_{11}(\omega) & 0 & 0\\ 0 & G_{22}(\omega) & 0\\ 0 & 0 & G_{33}(\omega)\end{bmatrix},$$

where $G_{ii}(\omega)$ is chosen to be the optimal predictor for the signal $u_i(n)$ with power spectrum $\Phi_{uii}(\omega)$. Notice, that $\tilde{\mathbf{W}}$ and $\tilde{\mathbf{W}}^{-1}$ cancel in Eq. (5) since $\tilde{\mathbf{W}}(\omega)$ and $\tilde{\mathbf{G}}(\omega)$ are diagonal, and therefore $\tilde{\mathbf{G}}(\omega)=\mathbf{G}(\omega)$.

3 Whitening the Quantization Error Spectrum

3.1 *Dithered Error Diffusion*

The design of the error diffusion filter assumed that the quantization error $q(n)$ has a white power spectrum. However, simulations showed that this assumption is not always valid. Based on our simulations, two observations about the power spectrum of $q(n)$ were made. On the one hand, the whiteness of $q(n)$ depends on the error diffusion filter. On the other hand, the simulations indicated that the quantization error is not white in regions of the image where the input gray-scale value is close to one of the quantization levels. The quantization error resulting from regions of such input pixels is not large enough to alter the output value of the quantizer. This results in bands of constant gray level [as can be seen in Fig. 11(a)].

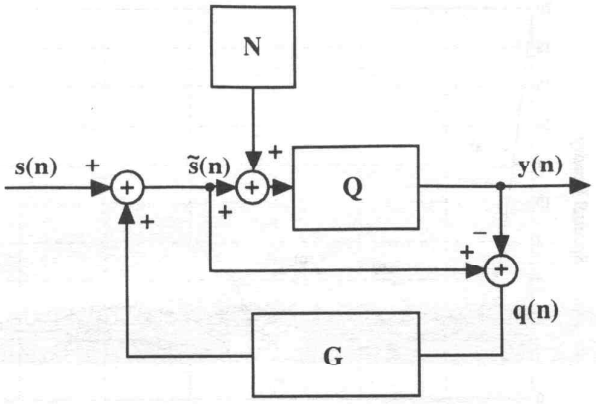


Fig. 8 Block diagram of the dithered error diffusion algorithm.

In Appendix C, we show that a possible way to obtain a white quantization error spectrum is to vary randomly the thresholds $t(n)$ of the quantizer Q. Moreover, it is shown that the expected value of the output image, $E[Q[\tilde{s}(n)]]$, is equal to the expected value of the quantizer input, $E[\tilde{s}(n)]$, and that the signal is uncorrelated with $q(n)$. Let α_k for $k=0, \ldots, M-1$ be the M nonuniformly spaced quantization levels, and assume that the signal falls in the range $\alpha_0 \leq s(n) \leq \alpha_{M-1}$. The dithered quantization of $s(n)$ is given by

$$y(n)=\tilde{Q}[s(n)]=\begin{cases}\alpha_{k-1} & \text{if } \alpha_{k-1} \leq s(n) < t(n)\\ \alpha_k & \text{if } t(n) \leq s(n) < \alpha_k\end{cases},$$

where $t(n)$ is a sequence of independent random thresholds uniformly distributed on (α_{k-1}, α_k).

We also show in Appendix C that for a uniform quantizer the variation of the thresholds is equivalent to adding noise to the signal before it is quantized. This is done by defining a noise sequence $N(n)$ of independent random variables. Each N is uniformly distributed on the interval $(-\Delta q/2; \Delta q/2)$. As for the nonuniform quantizer, the expected value of the output image, $E[Q[\tilde{s}(n)]]$, is equal to the expected value of the quantizer input, $E[\tilde{s}(n)]$, and the signal is uncorrelated with the noise. Figure 8 illustrates the resulting quantization system. We will refer to this method as *dithered error diffusion* (DED).

Figure 9 shows estimates of the power spectrum of the quantization error obtained for a natural image using the Floyd-Steinberg filter and an optimized error diffusion filter. The estimates of the power spectrum are calculated by averaging periodograms over the image. They are plotted as a function of f_1, while f_2 is set to zero. We used a uniform 3-bit quantizer and limited the input signal to the range of the quantizer to avoid saturation effects. A comparison of the spectra, resulting from the basic error diffusion method and the DED algorithm, clearly shows the whitening effect of the DED algorithm when the Floyd-Steinberg filter is used. We also notice that the optimized error diffusion filter seems to produce a white quantization error, without the addition of dithering noise. However, experimental results in Sec. 5 will show that the quantization error can be locally nonwhite. Therefore, dithering noise has to be added to

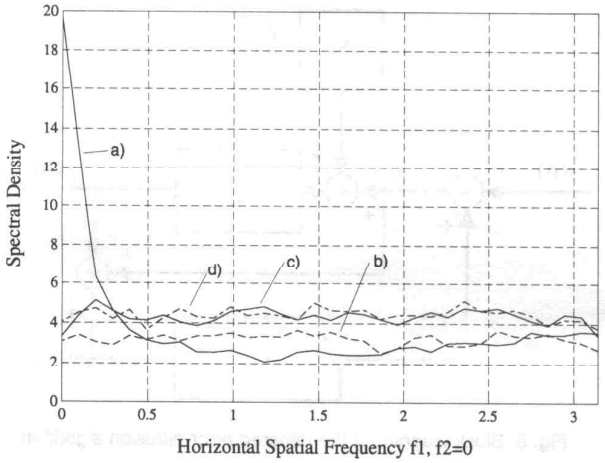


Fig. 9 Estimates of the power spectrum of the quantization error as a function of $f_1 \in (0,\pi)$ and $f_2 = 0$: (a) basic error diffusion, Floyd-Steinberg filter; (b) optimized error diffusion, linear predictor; (c) DED, Floyd-Steinberg filter; and (d) DED, linear predictor.

ensure whiteness of the quantization error. The disadvantage of the DED algorithm is an overall increase of graininess and noise in the output image.

3.2 *Locally Dithered Error Diffusion*

To produce results similar to DED but with less noise, we selectively add noise to the image based on a nearest neighbor criterion. We will call this method *locally dithered error diffusion* (LDED). Figure 10 illustrates the algorithm. Our simulations indicate that areas with smooth, slowly varying gray-scale values in the input image often show artificial textural contouring after they are quantized. This may be due to the fact that most pixels in a neighborhood of the input image are represented by the same quantization level and that the diffused error is too small to change the quantization level of successive pixels. As a result, the quantization error tends to have the same sign in such areas. Therefore, it seems to be desirable to whiten the quantization error by adding some noise in critical regions of the image. As a criterion, the algorithm compares the quantization levels of two previously quantized pixels, located to the left and above the current input pixel, such that

```
if (y(n1 − 1,n2) = y(n1,n2 − 1))
   dither quantization
else
   do not dither quantization .
```

Whenever the same quantization level was assigned to these two pixels, dithered quantization was used. This modification reduces the variance of the quantization error compared to dithered error diffusion, but also whitens the quantization error in critical regions of the image.

4 Models for the Human Visual System

4.1 *Luminance Contrast Sensitivity*

In this section, we briefly describe two of the various human MTF models found in the literature, and then develop the model used in our experiments. Most models for the human

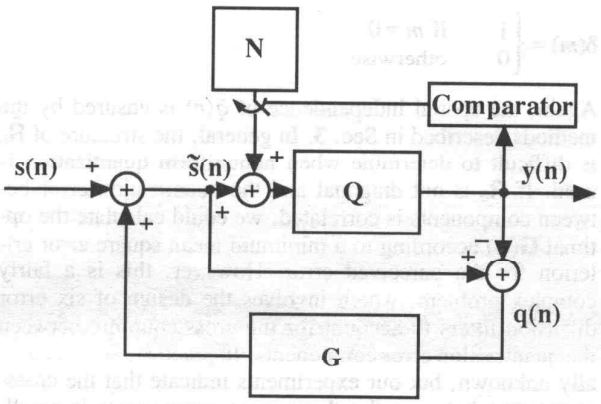


Fig. 10 Block diagram of the locally dithered error diffusion algorithm.

MTF have a low-pass characteristic. It should be emphasized that our proposed method for designing an optimal error diffusion filter for a given visual model is applicable to any given modulation transfer function describing the specific display and viewing conditions.

Sullivan and colleagues[14,15] quote a proprietary model from Daly that describes the human MTF by a weighted exponential function of the frequency vector. The model is also a function of the viewing angle and decreases faster for diagonal frequencies to account for the reduced sensitivity of the human visual system toward luminance changes in diagonal directions.

Näsänen[16] describes a circular symmetric model based on an exponential function. This model features parameters, which depend on the average luminance of the display. The slope of the exponential decreases with increasing average luminance.

For our simulations, a model was used that combines the models from Sullivan et al. and Näsänen. Our model is based on an exponential function as used by Näsänen[16]:

$$W(\tilde{f}) = K(L)\,\exp[-\alpha(L)\tilde{f}]\ ,$$

where L is the average luminance in cd/m^2, $\tilde{f}$ is the spatial frequency in cycles per degree,

$$K(L) = aL^b\ ,$$

$$\alpha(L) = \frac{1}{c\,\ln(L) + d}\ ,$$

and $a = 131.6$, $b = 0.3188$, $c = 0.525$, and $d = 3.91$. Deviating from Näsänen's model, we define $\tilde{f}$ to be the weighted magnitude of the frequency vector $f = (f_1, f_2)$, where the weighting has an angular dependence as used by Sullivan, Ray, and Miller[14] of

$$\tilde{f} = \frac{\|f\|}{s(\Theta)}\ ,$$

where $\|f\| = (f_1^2 + f_2^2)^{1/2}$ and

$$s(\Theta) = \frac{1-w}{2}\cos(4\Theta) + \frac{1+w}{2}\ .$$

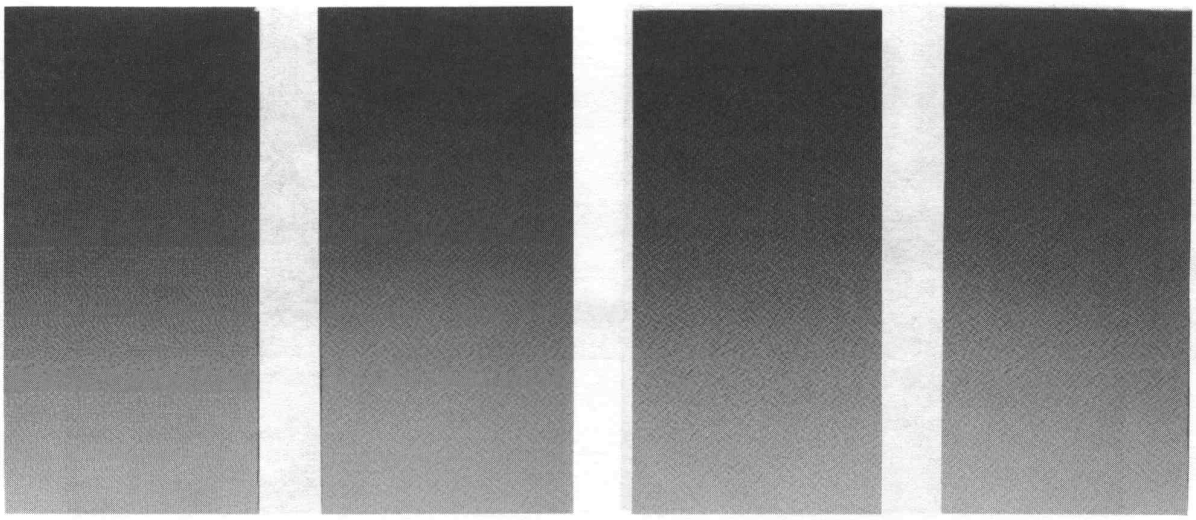


Fig. 11 Error diffusion applied to a gray-scale ramp (zoomed by factor of 2): (a) basic error diffusion, Floyd-Steinberg filter; (b) basic error diffusion, first-order linear predictor; (c) DED, first-order linear predictor; and (d) LDED, first-order linear predictor.

The angle Θ is defined as

$$\Theta = \arctan\left(\frac{f_1}{f_2}\right)$$

and the symmetry parameter $w = 0.7$. Throughout our simulations, we assumed a monitor MTF $P(\Omega) = 1$. An investigation of the spatial characteristics of gray-scale CRT pixels was recently presented by Naiman and Makous.[21]

4.2 *Chrominance Contrast Sensitivity*

We base our model for human sensitivity toward variations in chrominance on the experimental results obtained by Mullen.[17] Evaluating the presented data, the chrominance contrast sensitivity function can be approximated by an exponential function,

$$W(f) = A \exp(-\alpha \tilde{f}) \ .$$

We determined $\alpha = 0.419$, $A = 100$ for both chrominance components, and assumed circular symmetry for our model such that $\tilde{f} = (f_1^2 + f_2^2)^{1/2}$. As with the luminance model, we assume $P(\Omega) = 1$ for both chrominance components.

5 Experimental Results

5.1 *Error Diffusion in Luminance*

Our experiments were carried out on a Sun SPARCstation 1 with a calibrated color monitor. Simulations were performed with 512×512 pixel natural images and a gray-scale ramp. A monochrome version of a natural color image was obtained by calculating the luminance component Y according to the Y,cr,cb standard[20] from RGB image data. The monochrome image was then scaled and rounded to integer values between 0 and 255. Uniform and nonuniform quantizers were examined. We chose to use a nonuniform quantizer where steps increase toward higher luminance levels since this corresponds to the fact expressed in Weber's law that human contrast sensitivity follows a logarithmic function of luminance. Let $\alpha_{Y,k}$ denote the luminance quantization levels, where a value of 0 corresponds to zero intensity, 255 denotes maximal intensity, and intermediate values are proportional to the energy of the corresponding image pixel. The luminance quantization levels were assigned such that

$$\alpha_{Y,k} = 255(k/7)^{1.8} \ ,$$

where $k = 0, \ldots, 7$.

Our output images were displayed on an 8-bit color monitor. We calibrated the monitor using measured values of gamma and offset. The error-diffused image was transformed according to the nonlinearity of the display. A normal viewing distance was assumed and measured to be 40 cm, which is equivalent to approximately 3.5 times the image height. The quantized images shown in this paper are zoomed by a factor of 2 and should therefore be viewed at a distance of approximately 7 times the image height. Calculating the angle of perception for a single pixel yields $\Delta\phi = 0.0317$ deg/pixel. This determines the sampling rate, $f_s = (\Delta\phi)^{-1} = 31.5$ cycles/deg. The average luminance of our display was measured to be approximately $L = 11.0$ cd/m^2.

The coefficients for the first-order linear prediction filters for luminance and chrominance are listed in Table 1. Notice that the coefficient $g_{1,1}$ of both prediction filters is negative and that the coefficients do not sum to one. Since the optimized filters shape the display error spectrum according to a human visual model, some display error energy is allowed at zero frequency. Notice also that the coefficient $g_{1,-1}$ is very close to zero and can be neglected to reduce the computational complexity of the algorithm.

Figures 11(a) and 11(b) show a quantized gray-scale ramp, comparing the performance of the basic error diffusion algorithm, using the Floyd-Steinberg filter, and the first-order

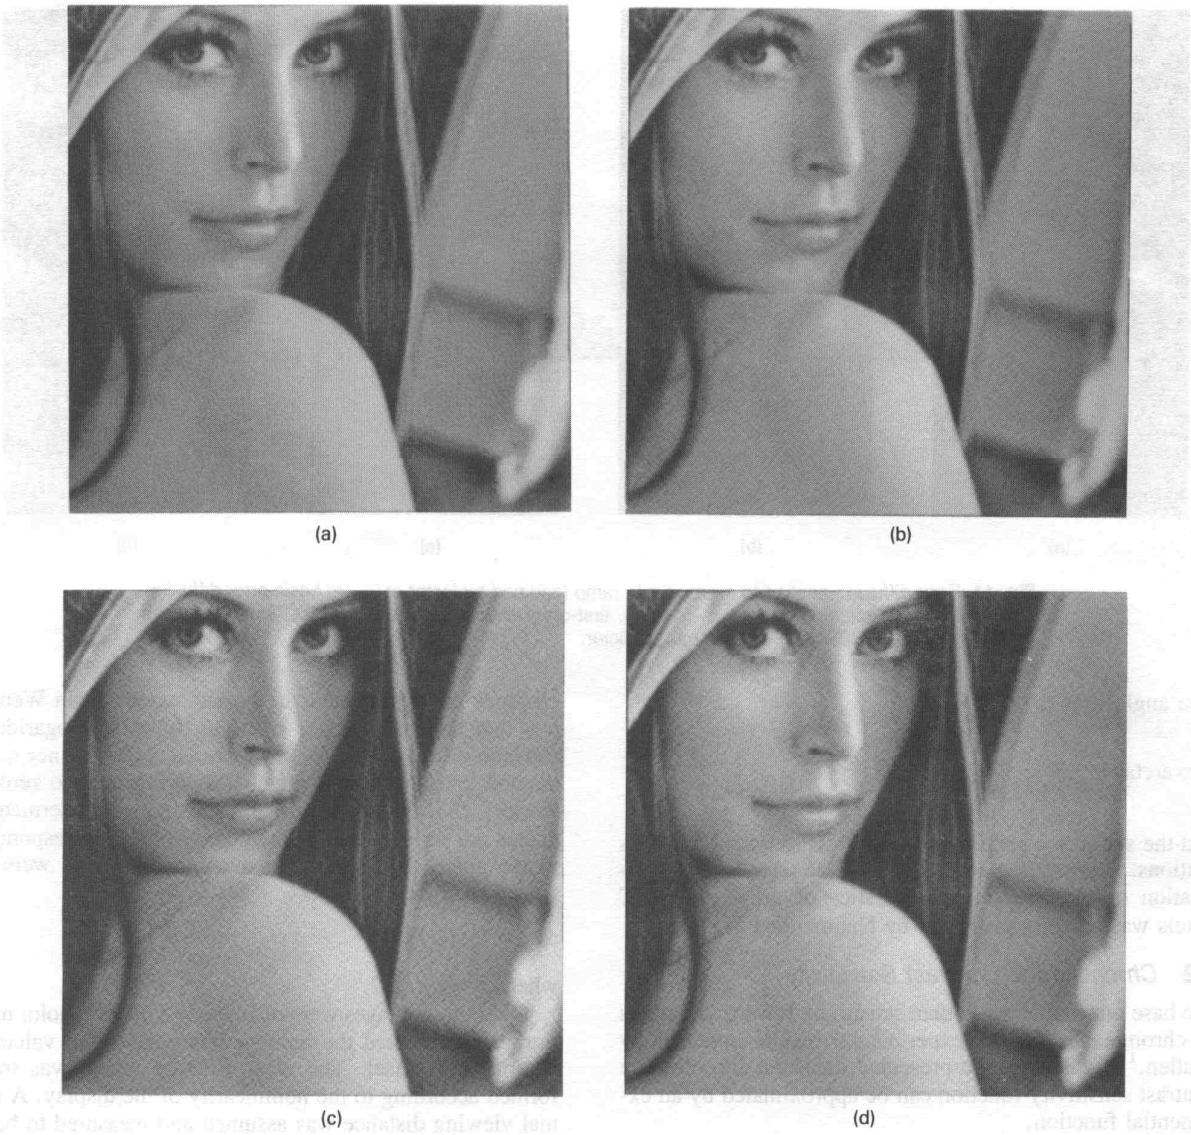


Fig. 12 Error diffusion applied to a natural image (zoomed by factor of 2): (a) basic error diffusion, Floyd-Steinberg filter; (b) basic error diffusion, first-order linear predictor; (c) DED, first-order linear predictor; and (d) LDED, first-order linear predictor.

linear predictor for luminance, according to the human visual model described in Sec. 4. When using the Floyd-Steinberg filter, the quantized image shows diagonal patterns and artificial contours. Applying a first-order linear predictor with the basic error diffusion shows fewer patterns.

Our simulation indicated that for basic error diffusion averaged over the entire image the quantization error spectrum is whiter when using a linear prediction filter than for the alternative filters (see Fig. 9). This may result from limit cycles and saturation effects, which occur more frequently for filters with coefficients that sum to one. For all tested filters, the quantization error is not equally white in all regions of the image. Whenever the input gray-scale value is close to a quantization level, the whiteness assumption seems to be locally violated. In these regions, the contribution of the quantization errors of adjacent pixels to the current pixel is negligible. Thus, the quantizer assigns the same output gray level to almost all pixels in that region. In Figs. 11(a) and 11(b) this causes artifacts that are visible as lines in the gray-scale ramp.

Figure 11(c) shows the resulting image for the DED algorithm using the first-order linear predictor. The resulting image does not suffer from lines and patterns as the previous two examples. However, the overall perception of the image is more noisy.

Figure 11(d) shows the resulting image for the LDED algorithm. This algorithm also removes the false contours, but the LDED algorithm is superior to the DED algorithm since it reduces the amount of noise in the image.

Figure 12 shows the same algorithms applied to a natural image. Critical regions of the image such as the shoulder and the cheek appear to be smoother for the LDED algorithm

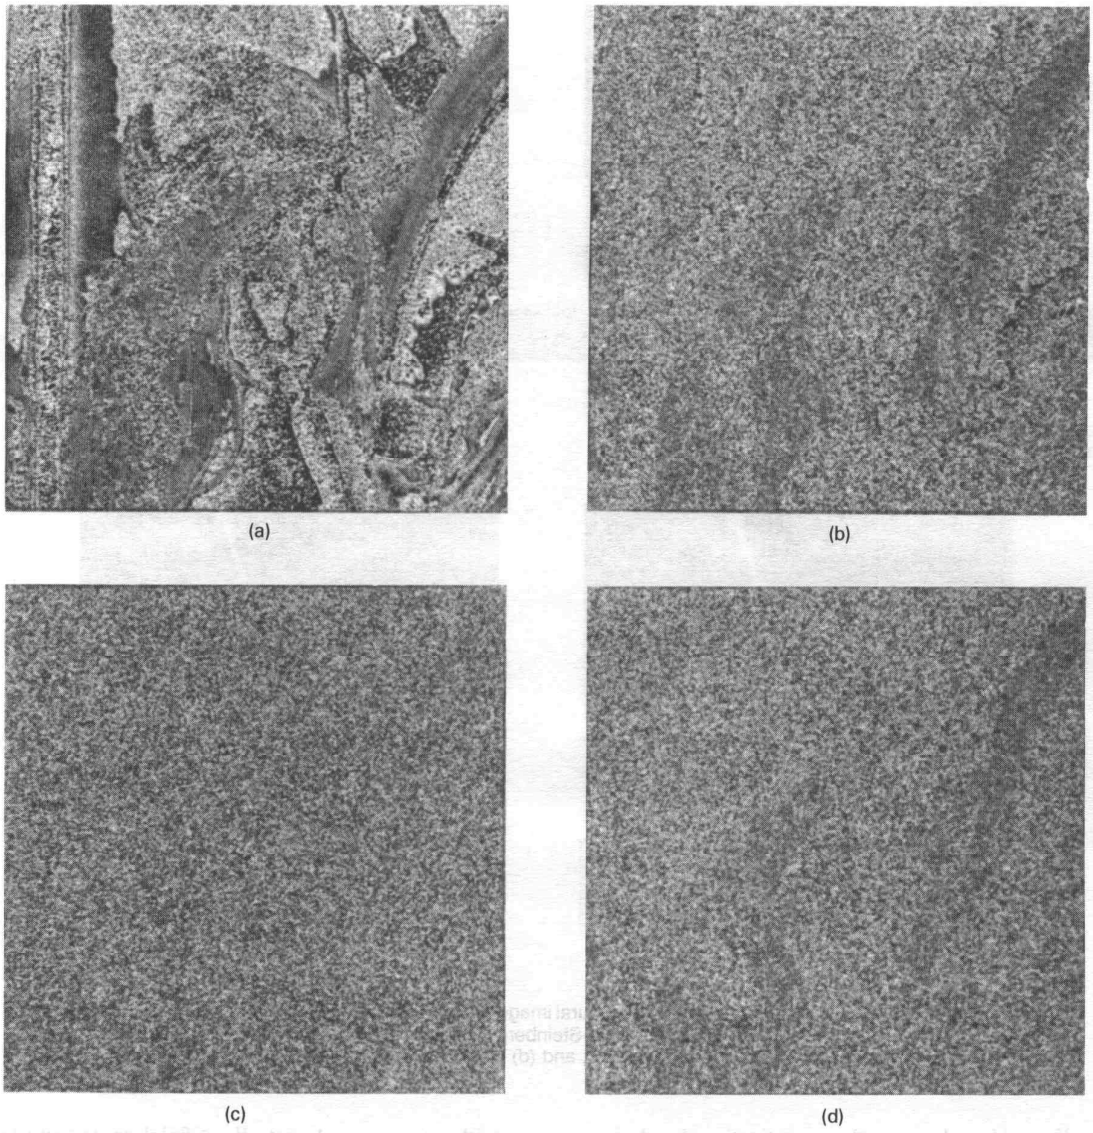


Fig. 13 Quantization error images for a natural image (full image): (a) basic error diffusion, Floyd-Steinberg filter; (b) basic error diffusion, first-order linear predictor; (c) DED, first-order linear predictor; and (d) LDED, first-order linear predictor.

than for basic error diffusion, using either the Floyd Steinberg filter or the optimal linear predictor.

Knox[9] proposed to display the quantization error image to examine the correlation between the input image and the quantization error. Figure 13(a) shows the quantization error image, for a natural image, processed with basic error diffusion, using the Floyd-Steinberg filter. Note that Fig. 13 shows the full error images. It can be seen that the quantization error is highly correlated with the input image. Figure 13(b) shows the quantization error image for the same natural image, processed with our optimized error diffusion algorithm. Notice that the quantization error becomes less correlated to the input image simply by changing the coefficients of the error diffusion filter. The whitening effect of adding dithering noise can be seen in Figs. 13(c) and 13(d), which show the quantization error images for DED and LDED. Adding dithering noise globally to the image produces a quantization error that is almost uncorrelated. The LDED algorithm improves the whiteness of the quantization error without adding an excessive amount of noise.

5.2 *Color Error Diffusion*

Our images were quantized in RGB coordinates, using 256 colors, where we assigned 3 bits (8 levels) to red, 3 bits to green, and 2 bits (4 levels) to blue. Our input data are linear in luminance. Again, uniform and nonuniform quantizers were examined, where a nonuniform quantizer was chosen with increasing steps toward higher luminance levels. Let $\alpha_{r,k}$, $\alpha_{g,k}$, and $\alpha_{b,k}$ denote the R, G, and B quantization levels, respectively, where a value of 0 corresponds to zero intensity, 255 denotes maximal intensity, and intermediate values are proportional to the energy of the corresponding

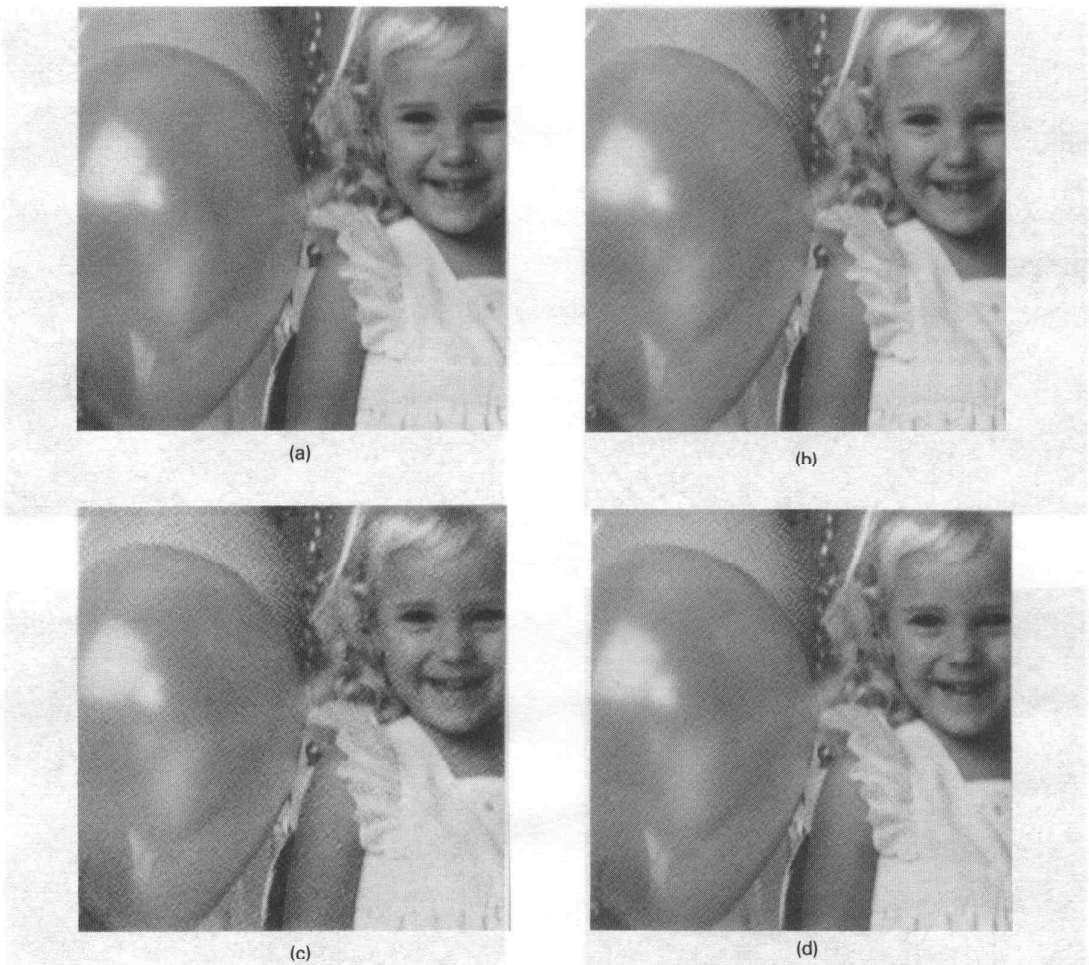


Fig. 14 Color error diffusion applied to the natural image "Balloon" (monochrome reproduction, zoomed by factor of 2): (a) basic error diffusion, Floyd-Steinberg filter; (b) basic error diffusion, first-order linear predictor; (c) DED, first-order linear predictor; and (d) LDED, first-order linear predictor.

primary. For red and green the quantization levels were assigned such that

$$\alpha_{r,k}=\alpha_{g,k}=255(k/7)^{1.8} ,$$

where $k=0, \ldots, 7$. For the blue primary we used

$$\alpha_{b,k}=255(k/3)^{1.8} ,$$

where $k=0, \ldots, 3$. The monitor was calibrated using offsets and gamma for each color gun. The color of the monitor was also calibrated by measuring the elements of matrix **A** in Sec. 2.3. Display and viewing conditions were the same as for the monochrome case.

Figure 14 shows monochrome reproductions of our experimental results using four different error diffusion methods applied to a color image. For all cases we used a color transformation to a luminance-chrominance space as described in Sec. 2.3. Notice that using the Floyd-Steinberg filter on all three components of the quantization error cancels the effect of the transformation **T** and its inverse. As in the case of luminance error diffusion, the optimized algorithm removes almost all artificial contouring in smooth areas of the image. Furthermore, it eliminates the deterministic patterns typically associated with the Floyd-Steinberg filter.

Figure 15 shows the Y component of the quantization error images for a natural color image. As in the case of luminance error diffusion, the correlation between the input image and the quantization error is greatly reduced, when using an optimized error diffusion filter. However, Figs. 15(c) and 15(d) show that the contours of the original image are still noticeable, even when dithering noise is added to the quantizer input. This results from saturation of the quantizer, mainly for the blue component, since only four output levels are assigned to this component and the white quantization error assumption is often violated.

6 Conclusion

Starting with an analysis of the basic error diffusion algorithm proposed by Floyd and Steinberg, we developed an optimization criterion for the design of an error diffusion filter for monochrome images, based on a model for the

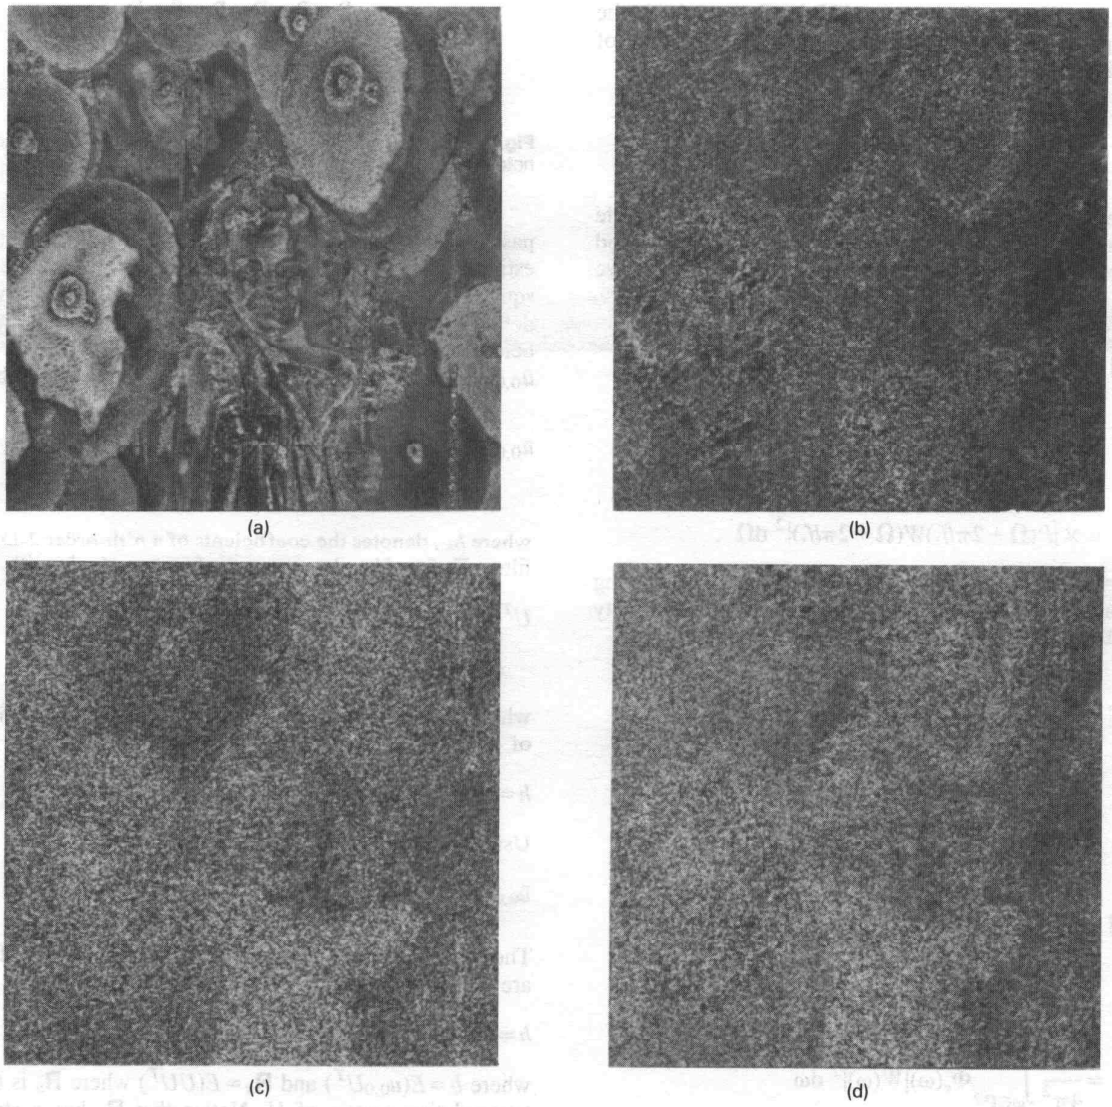


Fig. 15 Luminance quantization error for "Balloon" (full image): (a) basic error diffusion, Floyd-Steinberg filter; (b) basic error diffusion, first-order linear predictor; (c) DED, first-order linear predictor; and (d) LDED, first-order linear predictor.

human visual system. We showed that the analysis can also be extended to color images. Additive white noise is used to ensure a white quantization error spectrum, which is an important underlying assumption for the design of the error diffusion filter. The resulting images show less artificial contouring and a uniform texture. In addition, the noise in slowly varying regions of the image apears to be more uncorrelated and less noticeable.

7 Appendix A

In this section we will show that $E[\varepsilon^2(x)] = E[\tilde{\varepsilon}^2(n)]$, where

$$E[\tilde{\varepsilon}^2(n)] = \frac{1}{4\pi^2} \int_{\omega \in \mathbb{P}^2} \Phi_e(\omega) |\tilde{W}(\omega)|^2 \, d\omega$$

and $\mathbb{P} = [-\pi, \pi]$, while $E[\varepsilon^2(x)]$ is given by

$$E[\varepsilon^2(x)] = \frac{1}{4\pi^2} \int_{\Omega \in \mathbb{R}^2} \Phi\varepsilon(\Omega) \, d\Omega \ , \tag{8}$$

where $\Omega = (\Omega_1, \Omega_2)$ is a frequency vector with $\Omega_1, \Omega_2 \in (-\infty, \infty)$. The digital image is displayed and converted to a spatially continuous signal. It is filtered by the monitor MTF $P(\Omega)$. Ideally we would like $P(\Omega)$ to be a perfect low pass with

$$P(\Omega) = \begin{cases} 1 & \text{if } \Omega \in \mathbb{D}^2 \\ 0 & \text{otherwise} \ , \end{cases}$$

where $\mathbb{D} = [-\pi f_s, \pi f_s]$, and f_s is the sampling frequency of the monitor. Since the monitor is not ideal, we also have to consider effects due to the side lobes of $P(\Omega)$. Accounting

also for the effect of the human MTF $W(\Omega)$, we obtain the power spectrum of the perceived error $\Phi_\epsilon(\Omega)$ in terms of the display error power spectrum $\Phi_e(\Omega/f_s)$:

$$\Phi_\epsilon(\Omega)=\Phi_e\left(\frac{\Omega}{f_s}\right)|P(\Omega)W(\Omega)|^2 \ .$$

We note that $\Phi_e(\Omega/f_s)$ is the power spectrum of a discrete signal and is therefore a 2-D periodic function with period 2π. Substituting in Eq. (8) and expanding the integral we obtain

$$\begin{aligned}E[\varepsilon^2(x)]&=\frac{1}{4\pi^2}\int_{\Omega\in\mathbb{R}^2}\Phi_e\left(\frac{\Omega}{f_s}\right)|P(\Omega)W(\Omega)|^2\,d\Omega\\&=\frac{1}{4\pi^2}\sum_{l\in\mathbb{Z}^2}\int_{\Omega\in\mathbb{D}^2}\Phi_e\left(\frac{\Omega}{f_s}+2\pi l\right)\\&\quad\times|P(\Omega+2\pi lf_s)W(\Omega+2\pi lf_s)|^2\,d\Omega \ ,\end{aligned}$$

where $l=(l_1, l_2)$ and $\mathbb{Z}$ is the set of all integers. Exchanging the order of the sum and the integral and using the periodicity of $\Phi_e(\Omega/f_s)$, we have

$$\begin{aligned}E[\varepsilon^2(x)]&=\frac{1}{4\pi^2}\int_{\Omega\in\mathbb{D}^2}\Phi_e\left(\frac{\Omega}{f_s}\right)\\&\quad\times\sum_{l\in\mathbb{Z}^2}|P(\Omega+2\pi lf_s)W(\Omega+2\pi lf_s)|^2\,d\Omega \ .\end{aligned}$$

Finally, we make the substitution $\omega=\Omega/f_s$ to obtain

$$\begin{aligned}E[\varepsilon^2(x)]&=\frac{1}{4\pi^2}\int_{\omega\in\mathbb{P}^2}\Phi_e(\omega)f_s^2\\&\quad\times\sum_{l\in\mathbb{Z}^2}|P(\omega f_s+2\pi lf_s)W(\omega f_s+2\pi lf_s)|^2\,d\omega\\&=\frac{1}{4\pi^2}\int_{\omega\in\mathbb{P}^2}\Phi_e(\omega)|\tilde{W}(\omega)|^2\,d\omega\\&=E[\tilde{\varepsilon}^2(n)] \ ,\end{aligned}$$

where $\tilde{W}(\omega)$ is any discrete space filter, which satisfies

$$|\tilde{W}(\omega)|^2=f_s^2\sum_{l\in\mathbb{Z}^2}|P(\omega f_s+2\pi lf_s)W(\omega f_s+2\pi lf_s)|^2 \ .$$

8 Appendix B

In this section we develop linear predictor equations for the 2-D random process $u(n)$. The signal $u(n)$ is the output of the linear system $\tilde{W}(\omega)$. Since the input of the system is a stationary white noise signal $q(n)$, with power spectral density N_0, the output $u(n)$ is also stationary. The power spectral density of $u(n)$ is given by

$$\Phi_u(\omega)=N_0|\tilde{W}(\omega)|^2 \ . \tag{9}$$

Assuming raster order when processing the input image, each present pixel divides the image into a past and future halfplane as illustrated in Fig. 16. Based on pixels from the past (causal) halfplane of the signal, we want to obtain an estimate of the present pixel that is optimal in the mean square sense. Let $u_{0,0}$ denote the present pixel. Then we define $u_{k,l}$ to be the pixel l steps to the right and k steps below the present pixel. An estimate of the present pixel, $\tilde{u}_{0,0}$, in terms of weighted previous pixels is given by

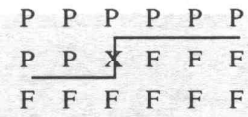


Fig. 16 P = past and F = future halfplanes of the image and X denotes the present pixel.

$$\tilde{u}_{0,0}=\sum_{l=1}^{p}h_{0,l}u_{0,-l}+\sum_{k=1}^{p}\sum_{l=-p}^{p}h_{k,l}u_{-\kappa,-l} \ ,$$

where $h_{k,l}$ denotes the coefficients of a p'th-order 2-D causal filter. Define U to be a vector of image pixels with

$$\begin{aligned}U^T=[&u_{0,-1},\ldots,u_{0,-p},u_{-1,p},\ldots,u_{-1,-p},\ldots,\\&u_{-p,p},\ldots,u_{-p,-p}] \ ,\end{aligned}$$

where T denotes the transpose of the vector. The dimension of U is $2(p^2+p)$. Then let

$$h=[h_{0,1},\ldots,h_{0,p},h_{1,-p},\ldots,h_{1,p},\ldots,h_{p,-p},\ldots,h_{p,p}] \ .$$

Using these definitions we can write for $\tilde{u}_{0,0}$

$$\tilde{u}_{0,0}=hU \ .$$

The coefficients of h that form the MMSE estimator for $\tilde{u}_{0,0}$ are given by

$$h=\mathbf{R}_u^{-1}b \ .$$

where $b=E(u_{0,0}U^T)$ and $\mathbf{R}_u=E(UU^T)$ where $\mathbf{R}_u$ is the autocorrelation matrix of U. Notice that $\mathbf{R}_u$ has a structure known as *block Toeplitz*. Let the point spread functions of the monitor and the human visual system be denoted by $p(x)$ and $w(x)$. Then the coefficients of b and $\mathbf{R}_u$ can be obtained by calculating and sampling the autocorrelation function $r_u(x)$ of the convolution of $p(x)$ and $w(x)$,

$$\begin{aligned}r_u(x)&=N_0p(x)*p(-x)*w(x)*w(-x)\\&=N_0\mathcal{F}^{-1}[|P(\Omega)W(\Omega)|^2] \ ,\end{aligned}$$

and $\mathcal{F}^{-1}$ denotes the inverse Fourier transform. We can approximate $r_u(x)$ by

$$r_u(x)\approx\Delta f^2\sum_{k\in\mathbb{Z}^2}|P(k\Delta f)W(k\Delta f)|^2\exp[j2\pi k\Delta fx] \ ,$$

where Δf is sufficiently small. To obtain the coefficients of $\mathbf{R}_u$, $r_u(x)$ is sampled at multiples of the spacing between adjacent pixels, Δx, with

$$E(u_nu_m)=r_u[(n-m)\Delta x] \ .$$

9 Appendix C

We show that dithered quantization of a signal $s(n)$ will generate quantization error $q(n)$, which is zero mean, white, and uncorrelated with the signal $s(n)$.

Let α_k for $k=0, \ldots, M-1$ be the M nonuniformly spaced quantization levels, and assume that the signal falls in the range of $\alpha_0 \leq s(n) \leq \alpha_{M-1}$. The dithered quantization of $s(n)$ is given by

$$y(n)=\tilde{Q}[s(n)]=\begin{cases}\alpha_{k-1} & \text{if } \alpha_{k-1} \leq s(n) < t(n) \\ \alpha_k & \text{if } t(n) \leq s(n) < \alpha_k\end{cases},$$

where $t(n)$ is a sequence of independent random thresholds uniformly distributed on $[\alpha_{k-1}, \alpha_k]$. Notice that in the case of uniform quantization with spacing $\Delta\alpha$, this corresponds to the simple operation

$$y(n)=Q[s(n)+N(n)] \ ,$$

where $Q[\cdot]$ is normal minimum distance quantization and $N(n)$ is a sequence of independent noise samples with uniform distribution on $[-\Delta\alpha/2, \Delta\alpha/2]$.

The quantization error is given by $q(n)=s(n)-y(n)$. We first compute the conditional mean of the quantization error given the signal:

$$\begin{aligned}E[q(n)|s(n)] &= E[s(n)-y(n)|s(n)] \\ &= s(n)-E[\tilde{Q}[s(n)]|s(n)] \\ &= s(n)-\alpha_{k-1}\frac{\alpha_k-s(n)}{\alpha_k-\alpha_{k-1}}+\alpha_k\frac{s(n)-\alpha_{k-1}}{\alpha_k-\alpha_{k-1}} \\ &= s(n)-s(n) \\ &= 0 \ .\end{aligned}$$

From this we may show that $q(n)$ is mean zero:

$$E[q(n)]=E[E[q(n)|s(n)]]=0$$

and uncorrelated for $n \neq m$:

$$\begin{aligned}E[q(n)q(m)] &= E[E[q(n)q(m)|s(n)s(m)]] \\ &= E[E[q(n)|s(n)]E[q(m)|s(m)]] \\ &= 0 \ .\end{aligned}$$

The variance of $q(n)$ is given by

$$\begin{aligned}E[[q(n)]^2] &= E[E][q(n)]^2|s(n)] \\ &= E\left[[s(n)-\alpha_{k-1}]^2\frac{\alpha_k-s(n)}{\alpha_k-\alpha_{k-1}}\right. \\ &\quad \left.+[\alpha_k-s(n)]^2\frac{s(n)-\alpha_{k-1}}{\alpha_k-\alpha_{k-1}}\right] \\ &= E[[s(n)-\alpha_{k-1}][\alpha_k-s(n)]] \\ &= \text{constant}=N_0 \ .\end{aligned}$$

Therefore, $q(n)$ has autocorrelation

$$E[q(n)q(m)]=\begin{cases}N_0 & \text{if } m=n \\ 0 & \text{otherwise}\end{cases}.$$

Finally, $q(n)$ and $s(m)$ are uncorrelated for all n and m:

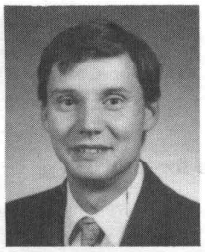
$$\begin{aligned}E[q(n)s(m)] &= E[E[q(n)s(m)|s(m)]] \\ &= E[s(m)E[q(n)|s(m)]] \\ &= E[s(m)0]=0\end{aligned}$$

Acknowledgment

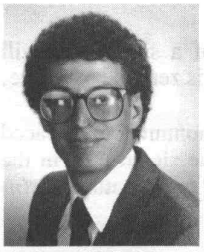
The author would like to thank the Eastman Kodak Company for the use of their image "Balloon" for the experiments. This work was supported by an NEC Faculty Fellowship.

References

1. R. W. Floyd and L. Steinberg, "An adaptive algorithm for spatial greyscale," *Proc. SID* **17**(2), 75–77 (1976).
2. E. Barnard, "Optimal error diffusion for computer-generated holograms," *J. Opt. Soc. Am. A* **5**(11), 1803–1817 (1988).
3. C. A. Bouman and M. T. Orchard, "Color image display with a limited palette size," *Proc. SPIE* **1199,** 522–533 (1990).
4. M. T. Orchard and C. A. Bouman, "Color quantization of images," *IEEE Trans. Signal Proc.* **39**(12), 2677–2690 (1991).
5. C. Billotet-Hoffmann and O. Bryngdahl, "On error diffusion technique for electronic halftoning," *Proc. SID* **24**(3), 253–258 (1983).
6. R. A. Ulichney, "Dithering with blue noise," *Proc. IEEE* **76**(1), 56–79 (1988).
7. G. Goertzel and G. R. Thompson, "Digital halftoning on the IBM 4250 printer," *IBM J. Res. Develop.* **31**(1), 2–15 (1987).
8. G. Goertzel and G. R. Thompson, " 'Halftoning' techniques for displaying images with a limited color palette," *Proc. EI West 1990,* Pasadena, CA, pp. 102–108 (1990).
9. K. T. Knox, "Error image in error diffusion," *Proc. SPIE* **1657,** 268–279 (1992).
10. J. Sullivan and R. Miller, "A new algorithm for image halftoning using a human visual model," *Proc. SPSE 43rd Annual Meeting,* Rochester, NY, pp. 145–148 (1990).
11. R. Miller and J. Sullivan, "Color halftoning using error diffusion and a human visual system model," *Proc. SPSE 43rd Annual Meeting,* Rochester, NY, pp. 149–152 (1990).
12. B. Girod, H. Almer, L. Bengtsson, B. Christensson, and P. Weiss, "A subjective evaluation of noise-shaping quantization for adaptive intra /interframe DPCM coding of color television signals," *IEEE Trans. Comm.* **36**(3), 332–346 (Mar. 1988).
13. T. Pappas and D. L. Neuhoff, "Least-squares model-based halftoning," *Proc. SPIE* **1666,** 165–176 (1992).
14. J. Sullivan, L. Ray, and R. Miller, "Design of minimum visual modulation halftone patterns," *IEEE Trans. Syst., Man, Cybernetics,* **21**(1) 33–38 (1991).
15. J. Sullivan, R. Miller, and G. Pios, "Image halftoning using visual error diffusion," submitted for publication.
16. R. Näsänen, "Visibility of halftone dot textures," *IEEE Trans. Syst., Man, Cybernetics,* **SMC-14**(6), 920–924 (1984).
17. K. T. Mullen, "The contrast sensitivity of human color vision to red-green and blue-yellow chromatic gratings," *J. Physiol.* **359,** 381–400 (1985).
18. J. G. Proakis and D. G. Manolakis, *Introduction to Digital Signal Processing,* Macmillan Publishing Company, New York (1988).
19. S. A. Rajala, H. J. Trussel, and B. Krishnakumar, "Visual sensitivity to color-varying stimuli," *Proc. SPIE* **1666,** 375–386 (1992).
20. A. N. Netravali and B. G. Haskell, *Digital Pictures,* Plenum Press, New York (1988).
21. A. C. Naiman and W. Makous, "Spatial nonlinearities of gray-scale CRT pixels," *Proc. SPIE* **1666,** 41–56 (1992).
22. P. Stucki, "MECCA-A multiple-error correcting computation algorithm for bilevel image hardcopy reproduction," *Res. Rep. RZ1060,* IBM Res. Lab., Zurich, Switzerland (1981).



Bernd W. Kolpatzik received a Dipl.-Ing. in electrical engineering from the Ruhr-Universität Bochum, Germany, in 1989. Since 1990 he has been a PhD student in the School of Electrical Engineering at Purdue University. His research interests are in color coding and halftoning applications.



Charles A. Bouman received a BSEE degree from the University of Pennsylvania in 1981, and a MS degree in electrical engineering from the University of California at Berkeley in 1982. From 1982 to 1985, he was a staff member in the Analog Device Technology Group at the Massachusetts Institute of Technology, Lincoln Laboratory. There his work included the development of algorithms and architectures, which used wide bandwidth analog signal processing devices. In 1987 and 1989, he received MA and PhD degrees in electrical engineering from Princeton University. In 1989, he joined the faculty of the School of Electrical Engineering at Purdue University as an assistant professor. Dr. Bouman's research interests include statistical image modeling, multiscale processing, and the application of human visual system models to the display, printing, and coding of images.