An Input-Level Dependent Approach to Color Error Diffusion

Vishal Monga[1], Niranjan Damera-Venkata[2], and Brian L. Evans[1]

[1]Dept. of Electrical and Computer Engineering
The University of Texas at Austin, Austin, TX 78712-1084
vishal@ece.utexas.edu, bevans@ece.utexas.edu

[2]Hewlett-Packard Laboratories
1501 Page Mill Road, Palo Alto, CA 94304-1126
damera@exch.hpl.hp.com

ABSTRACT

Conventional grayscale error diffusion halftoning produces worms and other objectionable artifacts. Tone dependent error diffusion (Li and Allebach) reduces these artifacts by controlling the diffusion of quantization errors based on the input graylevel. Li and Allebach optimize error filter weights and thresholds for each (input) graylevel based on a human visual system model. This paper extends tone dependent error diffusion to color. In color error diffusion, what color to render becomes a major concern in addition to finding optimal dot patterns. We present a visually optimum design approach for input level (tone) dependent error filters (for each color plane). The resulting halftones reduce traditional error diffusion artifacts and achieve greater accuracy in color rendition.

1. INTRODUCTION

Digital halftoning transforms a continuous tone image (grayscale or color) to an image with a reduced number of levels for display (or printing). Examples include converting an 8-bit per pixel grayscale image to a binary image, and a 24-bit color image (with 8 bits per pixel per color) to a 3-bit color image.

In grayscale halftoning by error diffusion, each grayscale pixel is thresholded to white or black, and the quantization error is fed back, filtered (lowpass), and added to the neighboring grayscale pixels.[1] Although the error filter is typically lowpass, the feedback arrangement causes the quantization error to be highpass filtered, i.e. pushed into high frequencies where the human eye is least sensitive. Grayscale error diffusion, however, introduces nonlinear distortion (worms and false textures), linear distortion (sharpening) and additive noise.[2] Many variations and enhancements on error diffusion have been developed to improve halftone quality, which includes using variable thresholds,[3–5] variable filter weights[6] and different scan paths.[7]

Recently, tone dependent error diffusion methods have been developed for grayscale error diffusion.[8,9] These methods use error filters with different coefficients for different graylevels in the input image. The quantizer threshold is also modulated based on the input graylevel.[8] In this paper, we formulate the design of input-level dependent color error diffusion halftoning systems. An independent design for each color plane would ignore the correlation amongst color planes. The color quantization error must ideally be diffused to frequencies and colors to which the eye is least sensitive. Our design procedure trains error filters for each color plane in order to minimize the perceived error between a constant-valued continuous-tone color image and its corresponding halftone pattern. A color human visual system (HVS) model takes into account the correlation among color planes. We employ an HVS model based on a transformation to the Linearized CIELab color space[10] that exploits the spatial frequency sensitivity variation of the luminance and chrominance channels. The efficacy of Linearized CIELab in computing color reproduction errors in halftoning is shown in.[11] The resulting halftones overcome many of the artifacts associated with traditional error diffusion viz. worms and false textures. The accuracy of color rendition is also higher. The color HVS model helps minimize the visibility of the halftone pattern.

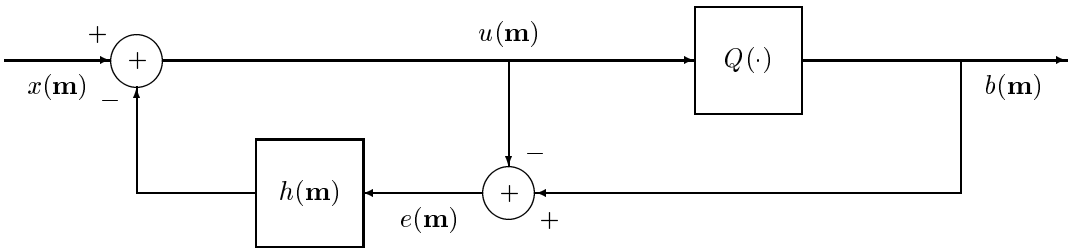

Figure 1. System block diagram for grayscale error diffusion halftoning where $\mathbf{m}$ represents a two-dimensional spatial index (m_1, m_2) and $h(\mathbf{m})$ is the impulse response of a fixed 2-D nonseparable FIR error filter

2. CLASSICAL ERROR DIFFUSION

In grayscale halftoning by error diffusion, each grayscale pixel is thresholded to white or black, and the quantization error is fed back, filtered and added to the neighboring grayscale pixels.[1] The system block diagram shown in Fig. 1, is also known as a noise-shaping feedback coder. In Fig. 1, $x(\mathbf{m})$ denotes the graylevel of the input image at pixel location $\mathbf{m}$, such that $x(\mathbf{m}) \in [0, 255]$. The output halftone pixel is $b(\mathbf{m})$, where $b(\mathbf{m}) \in \{0, 255\}$. Here, 0 is interpreted as the absence of a printer dot and 255 is interpreted as the presence of a printer dot. $Q(\cdot)$ denotes the standard thresholding quantizer function.

The error filter $h(\mathbf{m})$ filters the previous quantization errors $e(\mathbf{m})$:

$$h(\mathbf{m}) * e(\mathbf{m}) = \sum_{\mathbf{k} \in \mathcal{S}} h(\mathbf{k})\, e(\mathbf{m} - \mathbf{k}) \tag{1}$$

Here, $*$ means linear convolution, and the set $\mathcal{S}$ defines the extent of the error filter coefficient mask. The error filter output is fed back and added to the input. Note that $(0, 0) \notin \mathcal{S}$. The mask is causal with respect to the image scan. To ensure that all of the quantization error is diffused, $h(\mathbf{m})$ must satisfy the constraint

$$\sum_{\mathbf{k} \in \mathcal{S}} h(\mathbf{m}) = 1 \tag{2}$$

This ensures that the error filter eliminates quantization noise at DC where the human visual system is most sensitive.[12] The quantizer input $u(\mathbf{m})$ and output $b(\mathbf{m})$ are given by

$$u(\mathbf{m}) = x(\mathbf{m}) - h(\mathbf{m}) * e(\mathbf{m}) \tag{3}$$

$$b(\mathbf{m}) = Q(u(\mathbf{m})) \tag{4}$$

Error diffusion produces high quality halftones because the quantization noise is shaped into the high frequencies where the human eye is least sensitive. Since the halftone dots are of single pixel size, the illusion of continuous-tone is created by varying the dot frequency with graylevel.

3. GRAYSCALE TONE DEPENDENT ERROR DIFFUSION

Tone dependent error diffusion (TDED) methods use error filters $h(\mathbf{m})$ of different sizes and coefficients for different graylevels.[8,9] Optimal error weighting for selected graylevels based on "blue-noise" spectra was introduced in.[13]

Li *et al.* design error filter weights and thresholds so that the resulting halftones approximate dot patterns generated by DBS. DBS[14] produces high quality halftones by searching for the best binary pattern to match a given grayscale image by minimizing a visual distortion criterion. For the TDED algorithm in,[8] the error filter

and threshold matrix, denoted by $h(\mathbf{m})$ and $t(\mathbf{m}; a)$ respectively, are functions of input pixel value $a \in [0, 255]$. The binary output $b(\mathbf{m})$, is determined by

$$b(\mathbf{m}) = \begin{cases} 1, & \textbf{if } u(\mathbf{m}) \geq t(\mathbf{m}; x(\mathbf{m})) \\ 0, & \textbf{otherwise.} \end{cases} \tag{5}$$

The quantization error $e(\mathbf{m})$ and the quantizer input $u(\mathbf{m})$ are then computed as in conventional grayscale error diffusion. The threshold matrix used by Li and Allebach[8] is based on a binary DBS pattern for a constant input of mid-gray.

$$t(\mathbf{m}; a) = \begin{cases} t_u(a) & \text{if } p[\mathbf{m}, 128] = 0 \\ t_l(a) & \text{otherwise.} \end{cases} \tag{6}$$

where $t_u(a)$ and $t_l(a)$ are tone dependent parameters satisfying $t_u(a) \geq t_l(a)$. The function $p[\mathbf{m}, 128]$ is a halftone pattern generated by DBS that represents a constant patch of mid-gray. By substituting (6) into (5), the thresholding process can be represented by

$$b(\mathbf{m}) = \begin{cases} 1 & \text{if } u(\mathbf{m}) \geq t_u(x(\mathbf{m})) \\ 0 & \text{if } u(\mathbf{m}) < t_l(x(\mathbf{m})) \\ p[\mathbf{m}, 0.5] & \text{otherwise.} \end{cases} \tag{7}$$

For the error filter design, the authors chose the magnitude of the DFT of the DBS pattern as an objective spectrum for the halftone pattern for input graylevel values in the midtones (21-234). For the highlight and shadow regions (graylevel values in 0-20 and 235-255) the objective spectrum is the DFT of the graylevel patch. Let $B^{DBS}(k, l)$ and $B^{TDED}(k, l)$ denote the DFT of the DBS and the tone dependent error diffusion patterns, respectively. The goal is then to search for the tone dependent parameter vector $v = (t_u(a), t_l(a), h(\mathbf{m}; a))$ that minimizes

$$J = \sum_k \sum_l (|B^{TDED}(k, l)| - |B^{DBS}(k, l)|)^2 \tag{8}$$

subject to the constraints

$$t_u(a) + t_l(a) = 1 \tag{9}$$

$$t_u(a) \geq t_l(a) \tag{10}$$

$$\sum_{\mathbf{k} \in \mathcal{S}} h(\mathbf{k}; a) = 1 \tag{11}$$

$$h(\mathbf{k}; a) \geq 0 \quad \forall \quad \mathbf{k} \in \mathcal{S} \tag{12}$$

A serpentine scan is used in the design. Note that $B^{DBS}(k, l)$ would be replaced by $X(k, l)$ (the DFT of the input graylevel patch) for the highlight and shadow regions. The difference $|B^{TDED}(k, l)| - |X(k, l)|$ is then weighted (convolved) by a human visual system frequency response[14] to compute the perceived error. The algorithm to search for the optimum tone dependent parameter vector v_{opt} is described in.[8]

Independent design/application of grayscale halftoning methods to each color plane fails to exploit the visual response to color noise. The error filters must hence be designed to shape quantization errors to frequencies and colors, to which the HVS is least sensitive.

4. PERCEPTUAL MODEL

This section describes the model for calculating the perceived halftone in Linearized CIELab color space using the frequency responses of a channel-separable HVS.

4.1. Linearized Uniform Color Space

The linearized CIELab color space is obtained by linearizing the CIELab space about the D65 white point[10] in the following manner:

$$Yy = 116\,\frac{Y}{Y_n} - 16 \tag{13}$$

$$C_x = 500\left[\frac{X}{X_n} - \frac{Y}{Y_n}\right] \tag{14}$$

$$C_z = 200\left[\frac{Y}{Y_n} - \frac{Z}{Z_n}\right] \tag{15}$$

The Y_y component is proportional to the luminance and the C_x and C_z components are similar to the R-G and B-Y opponent color chrominance components on which Mullen's data[15] is based. The original transformation to the CIELab from CIEXYZ is a non-linear one.[16] The nonlinearity in the transformation from CIELab distorts the spatially averaged tone of the images, which yields halftones that have incorrect average values.[10] The linearized color space overcomes this, and has the added benefit that it decouples the effect of incremental changes in (Y_y, C_x, C_z) at the white point on (L, a, b) values:

$$\nabla_{(Yy,Cx,Cz)}(L^*, a^*, b^*)|_{D_{65}} = \frac{1}{3}\mathbf{I} \tag{16}$$

4.2. Human Visual Frequency Response

Nasanen and Sullivan[17] chose an exponential function to model the luminance frequency response

$$W_{(Y_y)}(\tilde{\rho}) = K(L)e^{-\alpha(L)\tilde{\rho}} \tag{17}$$

where L is the average luminance of display, $\tilde{\rho}$ is the radial spatial frequency, $K(L) = a\,L^b$ and $\alpha(L) = \frac{1}{c\ln(L)+d}$. The frequency variable $\tilde{\rho}$ is defined[10] as a weighted magnitude of the frequency vector $\mathbf{u} = (u, v)^T$, where the weighting depends on the angular spatial frequency ϕ.[17] Thus,

$$\tilde{\rho} = \frac{\rho}{s\,(\phi)} \tag{18}$$

where $\rho = \sqrt{u^2 + v^2}$ and

$$s\,(\phi) = \frac{1-\omega}{2}\cos(4\phi) + \frac{1+\omega}{2} \tag{19}$$

The symmetry parameter ω is 0.7, and $\phi = \arctan\left(\frac{v}{u}\right)$. The weighting function $s\,(\phi)$ effectively reduces the contrast sensitivity to spatial frequency components at odd multiples of 45^o. The contrast sensitivity of the human viewer to spatial variations in chrominance falls off faster as a function of increasing spatial frequency than does the response to spatial variations in luminance.[18] Our chrominance model reflects this[19]:

$$W_{(C_x,C_z)}(\rho) = Ae^{-\alpha\rho} \tag{20}$$

Both the luminance and chrominance response are lowpass in nature but only the luminance response is reduced at odd multiples of 45^o. This will place more luminance error across the diagonals in the frequency domain where the eye is less sensitive. Using this chrominance response as opposed to identical responses for both luminance and chrominance will allow more low frequency chromatic error, which will not be perceived by the human viewer.

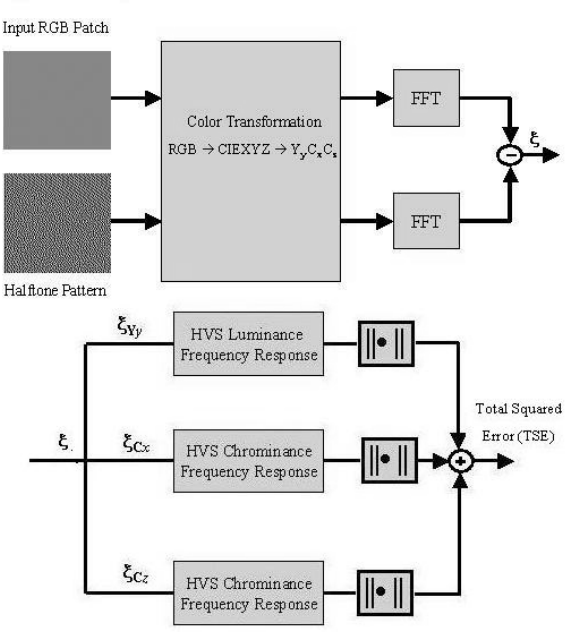

Figure 2. Block Diagram for Calculating Perceptual Error Metric

5. TONE DEPENDENT COLOR ERROR DIFFUSION

5.1. Perceptual Error Metric

We train error filters to minimize a visually weighted squared error between the magnitude spectra of a "constant" input color image and its halftone pattern. Let $x_{(R,G,B)}(\mathbf{m})$ and $b_{(R,G,B)}(\mathbf{m})$ denote the constant valued continuous tone and halftone images respectively. The calculation of the perceptual error metric is illustrated in Fig. 2. $x_{(Y_y,C_x,C_z)}(\mathbf{m})$ and $b_{(Y_y,C_x,C_z)}(\mathbf{m})$ are obtained by transforming $x_{(R,G,B)}(\mathbf{m})$ and $b_{(R,G,B)}(\mathbf{m})$ to the $Y_yC_xC_z$ space. The difference in their spectra $\xi(k,l)$ is then computed as $\xi(k,l) = \mathbf{X}_{(Y_y,C_x,C_z)}(k,l) - \mathbf{B}_{(Y_y,C_x,C_z)}(k,l)$ where

$$\mathbf{X}_{(Y_y,C_x,C_z)}(k,l) = FFT(x_{(Y_y,C_x,C_z)}(\mathbf{m})) \tag{21}$$

$$\mathbf{B}_{(Y_y,C_x,C_z)}(k,l) = FFT(b_{(Y_y,C_x,C_z)}(\mathbf{m})) \tag{22}$$

where FFT is the fast Fourier transform.

HVS filters in Section 4 are applied to the luminance and chrominance components of the error image in the spatial frequency domain. This corresponds to a multiplication of the filter and error image spectra $\mathbf{P}(k,l) = \xi(k,l)\mathbf{H}_{HVS}(k,l)$. Here, $\mathbf{H}_{HVS}(k,l)$ denotes the FFT of the human visual spatial filter. Note, $\mathbf{P}(k,l)$, $\mathbf{H}_{HVS}(k,l)$ and $\xi(k,l)$ are vector-valued

$$\xi(k,l) = (\xi_{Y_y}(k,l), \xi_{C_x}(k,l), \xi_{C_z}(k,l)) \tag{23}$$

$$\mathbf{H}_{HVS}(k,l) = (H_{Y_y}(k,l), H_{C_x}(k,l), H_{C_x}(k,l)) \tag{24}$$

$$\mathbf{P}(k,l) = (P_{Y_y}(k,l), P_{C_x}(k,l), P_{C_z}(k,l)) \tag{25}$$

We define the perceived error metric as the total squared error (TSE) given by

$$TSE = \sum_k \sum_l |P_{Y_y}(k,l)|^2 + |P_{C_x}(k,l)|^2 + |P_{C_z}(k,l)|^2 \tag{26}$$

5.2. Formulation of the Design Problem

The design problem is then to obtain error filters for each color plane that minimize the TSE defined in eqn. (26), subject to the constraints that all quantization error to be diffused

$$\sum_{\mathbf{k}\in\mathcal{S}} h_m(\mathbf{k};a) = 1, h_m(\mathbf{k};a) \geq 0 \quad \forall \quad \mathbf{k} \in \mathcal{S} \tag{27}$$

where the subscript m takes on values R, G and B and hence the constraints are imposed on error filters in each of the 3 color planes. The error filter coefficients are a function of the input level $a \in [0, 255]$. The design objective is to obtain error filter weights for each (R, G, B) vector in the input. For 24 bit color images, this would amount to a total of 256^3 input combinations. We consider input values along the diagonal line of the color cube i.e. $(R, G, B) = ((0, 0, 0), (1, 1, 1)....(255, 255, 255))$. This choice is made as the HVS is very sensitive to colors around the neutral axis.[20] This results in 256 error filters for each color plane. Note that the TSE in general, is not a convex function. Hence, a global minimum cannot be guaranteed. The space of solutions (error filter weights) however comprises a convex set. The algorithm to search for the error filter weights is described in.[8] The design is based on a Floyd-Steinberg[1] support for the error filter.

6. RESULTS

Figs. 3 and 4 shows color halftone images generated by (1) Floyd-Steinberg error diffusion, (2) grayscale TDED[8] applied separately to the R, G and B color planes, and (3) the proposed color TDED method. The color TDED halftone in Fig. 4(d) does not suffer from directional artifacts such as diagonal worms. Worms can be seen in the Floyd-Steinberg halftone in the yellow and blue extremes of the ramp. False textures in the Floyd-Steinberg halftone are prominent in the middle of the yellow region (a third of the ramp length from the left) and in the center of the ramp (where yellow turns into blue). These are nearly absent in Fig. 4(d). The choice of color to render is also better for the color TDED halftone. In Fig. 4(b), white dots are rendered in the blue region. These are replaced by a mixture of magenta, cyan and black dots in the color TDED halftone which are less visible. By virtue of the Li and Allebach's design,[8] traditional error diffusion artifacts viz. worms and false textures are almost completely removed in Fig. 4(c).

The halftone textures in Fig. 4(c) are also homogeneous. However, the color rendition is similar to that of Floyd-Steinberg error diffusion in Fig. 4(b). This is expected because the separable design for each color plane does not necessarily shape the color noise to frequencies of least visual sensitivity. Detail of the halftones in Fig. 4(b), (c) and (d) are shown in Fig. 3(a), (b) and (c). Note the significant improvement in the reduction of color halftone noise in Fig. 3(c) over Figs. 3(a) and (b).

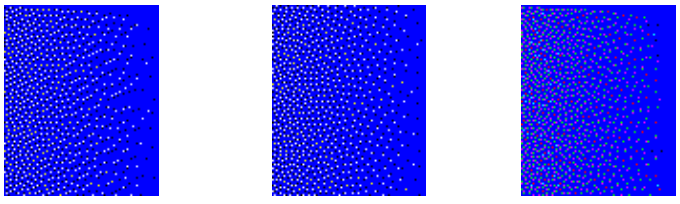

(a) Floyd-Steinberg (b) Grayscale TDED (c) Color TDED

Figure 3. Halftones of a portion of the blue section of the color ramp in Fig. 4(a). Grayscale TDED is applied separably to each color plane.

(a) Original Color Ramp Image

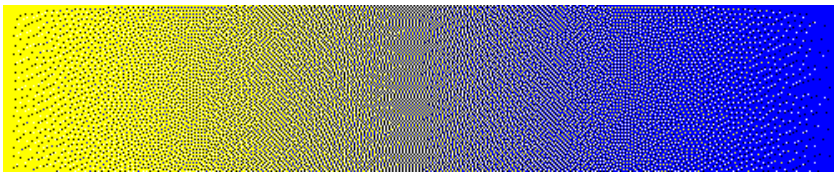

(b) Floyd-Steinberg error diffusion

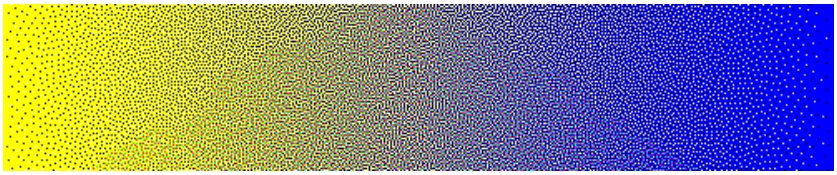

(c) Separable application of grayscale TDED[8]

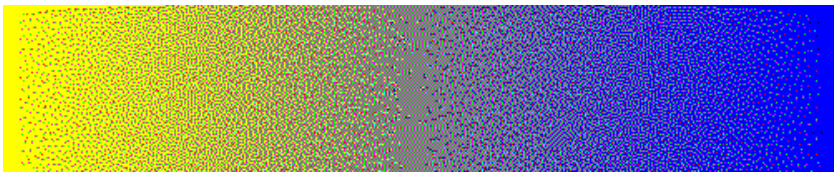

(d) Color TDED with serpentine scan

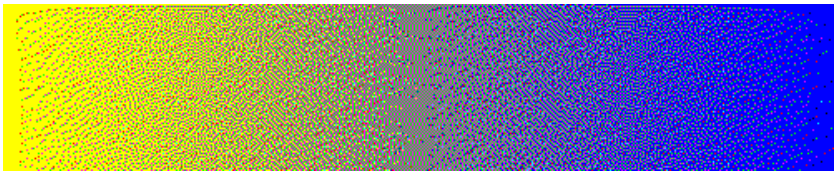

(e) Color TDED with traditional raster scan

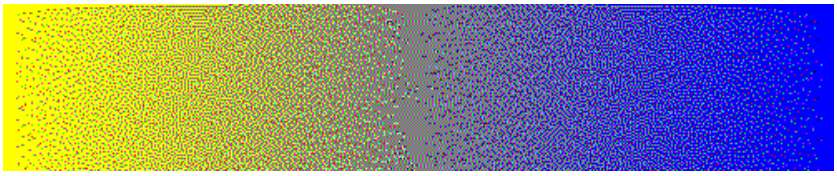

(f) Color TDED with 2-row serpentine scan

Figure 4. A color ramp and its halftone images. The halftone in (c) is courtesy of Prof. Jan P. Allebach and Mr. Ti-chiun Chang at Purdue University.

For color TDED, the conventional raster scan (Fig. 4(e)) still shows the tendency for dots/holes to line up in horizontal or diagonal worms, particularly at extreme levels. Although the serpentine scan (Fig. 4(d)) almost completely removes directional artifacts, it can only be executed as a serial process. A 2-row serpentine scan[8] employed in Fig. 4(f) generates comparable results to Fig. 4(d) but has more parallelism.

The results in Figs. 5 and 6 explain the role of color HVS model. Note that the green diagonal worms in the shadow region (extreme levels) under the roof that appear in the Floyd-Steinberg halftone in Fig. 5(a) are removed for the most part in Fig. 5(b). Monochrome images corresponding to the Y_y and C_x components of

(a) Floyd-Steinberg Halftone (b) TDED Halftone

Figure 5. (a) Floyd-Steinberg and (b) TDED halftones of the house image

the Floyd-Steinberg and TDED halftones are presented in Fig. 6(c) through (f). The Y_y component is obtained by converting the final halftone to the $Y_yC_xC_z$ color space and setting both the C_x and C_z components to zero, resulting in $(Y_y, 0, 0)$. The resulting vector is transformed back to the RGB space for display. To show only the C_x component of the halftone images both the Y_y and C_x components are set to zero. The resulting vector is then transformed back to the RGB space. This vector will not in general be a monochrome image due to the chromaticity of the C_x component. Therefore, the RGB components are normalized to maximum chromaticity $C_x = [0, 1, 0]$ to be printed in a black and white medium. Comparing Fig. 6(c) and (d) for Floyd-Steinberg error diffusion, we see that both the Y_y and C_x components exhibit similar texture. In contrast, we see in Fig. 6(e) and (f) that the C_x component has much lower frequency texture than the Y_y component, which is from the more aggressive HVS filtering of the chrominance planes. When viewed as true color images, the overall halftone texture is much less visible in the TDED halftone.

7. ENHANCEMENTS

In addition to using error filter weights that depend on the input level, error filter shape and size may also be varied for diffusion of errors. Li and Allebach[8] use wider matrices to improve rendering of shadow and highlight regions. With a wider matrix, the current pixel to be binarized will be affected by a larger area of the halftone region. Fig. 7 shows color error filter shapes (for each color plane) that are dependent on the input level $a \in$ [0,255]. To keep the computational complexity low, we use an error-weighting matrix that still has four non-zero terms, as does the Floyd-Steinberg filter, but with variable locations. The locations of the weights were determined empirically. The results presented in Section 6 use the error filter shapes in Fig. 7.

Many other approaches reduce directional artifacts in error diffusion. Fan proposed an error weighting matrix that allows the quantizer error to propagate further back on the next line to reduce worm artifacts.[21] Other approaches used larger error-weighting matrices/filters for highlight and shadow regions.[9,22] These ideas may all be combined into designing error filter weights and shape for each input level.

8. CONCLUSION

An input level (tone) dependent color halftoning algorithm was proposed. A linear channel-separable color HVS model is used to design visually optimum error filters for each color plane. Several error diffusion artifacts were reduced. The accuracy of color rendition is also greater. To further improve homogeneity of halftone textures, color DBS[23] may be incorporated into the design of input-level dependent color error filters. Future work could investigate the design of optimum matrix valued filters[24] for tone dependent vector error diffusion.

(a) Y_y component of the Floyd-Steinberg halftone

(b) C_x component of the Floyd-Steinberg halftone

(c) Y_y component of the TDED halftone

(d) C_x component of the TDED halftone

Figure 6. Luminance and chrominance components of the halftones in Fig. 5.

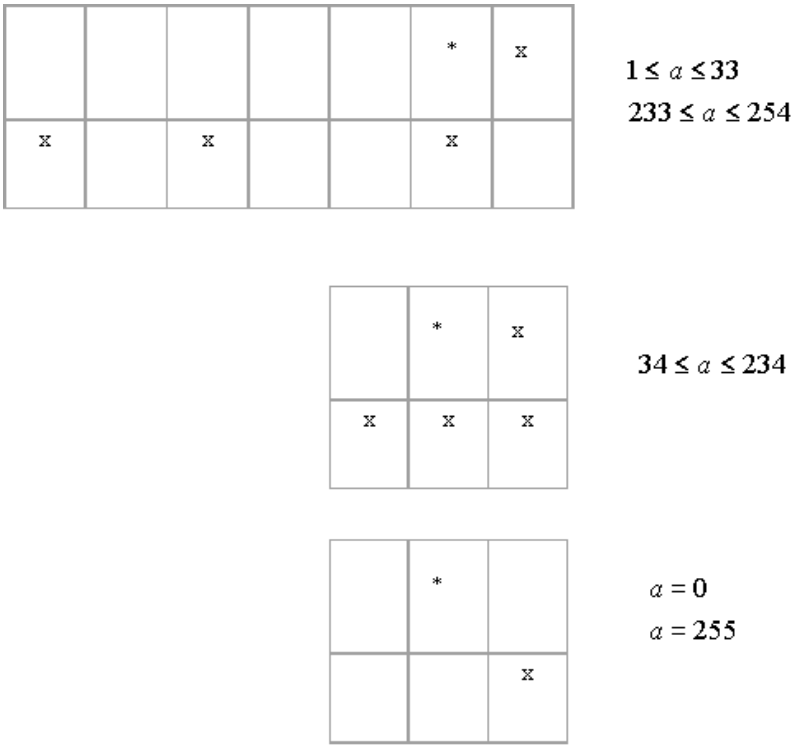

Figure 7. Input level dependent error filters, where * denotes the current pixel.

REFERENCES

1. R. Floyd and L. Steinberg, "An adaptive algorithm for spatial grayscale," *Proc. Soc. Image Display*, vol. 17, 1976.
2. B. L. Evans, V. Monga, and N. Damera-Venkata, "Variations on error diffusion : Retrospectives and future trends," *Proc. SPIE Color Imaging: Processing, Hardcopy and Applications VIII*, vol. 5008, pp. 371–389, Jan. 2003.
3. N. Damera-Venkata and B. L. Evans, "Adaptive threshold modulation for error diffusion halftoning," *IEEE Trans. on Image Processing*, vol. 10, no. 1, pp. 104–116, Jan. 2001.
4. J. Sullivan, R. Miller, and G. Pios, "Image halftoning using a visual model in error diffusion," *J. Opt. Soc. Am. A*, vol. 10, no. 8, pp. 1714–1724, Aug. 1993.
5. R. Eschbach, "Error-diffusion algorithm with homogeneous response in highlight and shadow areas," *J. Electronic Imaging*, vol. 6, pp. 1844–1850, July 1997.
6. P. Wong, "Adaptive error diffusion and its application in multiresolution rendering," *IEEE Trans. on Image Processing*, vol. 5, no. 7, pp. 1184–1196, July 1996.
7. R. Ulichney, *Digital Halftoning*, MIT Press, 1987.
8. P. Li and J. P. Allebach, "Tone dependent error diffusion," *SPIE Color Imaging: Device Independent Color, Color Hardcopy, and Applications VII*, vol. 4663, pp. 310–321, Jan. 2002.
9. R. Eschbach, "Reduction of artifacts in error diffusion by means of input-dependent weights," *J. Electronic Imaging*, vol. 2, no. 4, pp. 352–358, Oct. 1993.
10. T. J. Flohr, B. W. Kolpatzik, R. Balasubramanian, D. A. Carrara, C. A. Bouman, and J. P. Allebach, "Model based color image quantization," *Proc. SPIE Human Vision, Visual Proc. and Digital Display IV*, 1993.
11. V. Monga, W. S. Geisler, and B. L. Evans, "Linear, color separable, human visual system models for vector error diffusion halftoning," *IEEE Signal Processing Letters*, vol. 10, pp. 93–97, Apr. 2003.
12. J. Mannos and D. Sakrison, "The effects of a visual fidelity criterion on the encoding of images," *IEEE Trans. on Info. Theory*, vol. 20, no. 4, pp. 525–536, July 1974.
13. V. Ostromoukhov, "A simple and efficient error-diffusion algorithm," *Proc. Int. Conf. on Comp. Graphics and Interactive Tech.*, Aug. 2001.
14. M. Analoui and J. P. Allebach, "Model based halftoning using direct binary search," *Proc. SPIE Human Vision, Visual Processing, and Digital Display III*, Feb. 1992.

15. K.T. Mullen, "The contrast sensitivity of human color vision to red-green and blue-yellow chromatic gratings," *Journal of Physiology*, 1985.
16. M. D. Fairchild, *Color Appearance Models*, Addison-Wesley, 1998.
17. J. Sullivan, L. Ray, and R. Miller, "Design of minimum visual modulation halftone patterns," *IEEE Trans. on Systems, Man, and Cybernetics*, vol. 21, no. 1, pp. 33–38, Jan. 1991.
18. D. H. Kelly, "Spatiotemporal variation of chromatic and achromatic contrast thresholds," *Journal Opt. Soc. Amer. A*, vol. 73, pp. 742–750, 1983.
19. B. Kolpatzik and C. Bouman, "Optimized error diffusion for high quality image display," *Journal of Electronic Imaging*, vol. 1, pp. 277–292, Jan. 1992.
20. G. Sharma, *Digital Color Imaging Handbook*, CRC Press, 2002.
21. Z. Fan, "A simple modification of error diffusion weights," *Proc. IS&T 46th Annual Conference*, May 1993.
22. P. Stucki, "MECCA—a multiple-error correcting computation algorithm for bilevel hardcopy reproduction," Research Report RZ1060, IBM Research Laboratory, Zurich, Switzerland, 1981.
23. U. A. Agar and J. P. Allebach, "Model based color halftoning using direct binary search," *Proc. SPIE Color Imaging: Processing, Hardcopy and Applications VI*, 2000.
24. N. Damera-Venkata and B. L. Evans, "Design and analysis of vector color error diffusion halftoning systems," *IEEE Trans. on Image Processing*, vol. 10, no. 10, pp. 1552–1565, Oct. 2001.