Linear, Color Separable, Human Visual System Models for Vector Error Diffusion Halftoning

Vishal Monga, Wilson S. Geisler, III, and Brian L. Evans, *Senior Member, IEEE*

Abstract— **Image halftoning converts a high-resolution image to a low-resolution image, e.g. a 24-bit color image to a three-bit color image, for printing and display. Vector error diffusion captures correlation among color planes by using an error filter with matrix-valued coefficients. In optimizing vector error filters, Damera-Venkata and Evans transform the error image into an opponent color space where Euclidean distance has perceptual meaning. This paper evaluates color spaces for vector error filter optimization. In order of increasing quality, the color spaces are YIQ, YUV, opponent (by Poirson and Wandell), and linearized CIELab (by Flohr, Kolpatzik, Balasubramanian, Carrara, Bouman, and Allebach).**

Keywords— **Color image display, color quantization, image quality, multifilters**

I. Introduction

Image halftoning converts a high-resolution image to a lower resolution image, e.g. for printing and display. Common examples are converting an eight-bit per pixel grayscale image to a binary image, and a 24-bit color image (with eight bits per pixel per color) to a three-bit color image. Applying grayscale halftoning methods separably to color images does not take the correlation among color planes into account, and can lead to artifacts such as spurious color impulses in the resulting halftones.

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]. Although an 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. The feedback arrangement sharpens the original image by passing low frequencies and amplifying high frequencies. Traditional grayscale error diffused halftones appear sharper than the original and contain highpass noise [2].

Vector error diffusion, first proposed in [3], represents each pixel in a color image as a vector of values. The thresholding step would threshold each vector component separately. The vector-valued quantization error (image) would be fed back, filtered, and added to the neighboring (unhalftoned) color pixels. A matrix-valued error filter could take the correlation among color planes into account. For an RGB image, each error filter coefficient would be a 3×3 matrix. For RGB vector error diffusion, matrix-valued error filter coefficients are adapted in [4] to reduce the mean squared error between the halftone and original. However, mean squared error does not have perceptual meaning in RGB space.

Authors are with the Center for Perceptual Systems, The University of Texas at Austin, Austin, TX 78712 USA. {vishal,bevans}@ece.utexas.edu, geisler@mail.utexas.edu.

A linear color human visual system (HVS) framework [5] designs matrix-valued error filters by transforming the error image into luminance and chrominance components and minimizing a linear model of the human visual response to the quantization noise. In this framework, mean squared error has perceptual meaning. Based on this framework, this paper evaluates four color spaces in which to optimize matrix-valued error filters. We find that the objective and subjective rankings of color spaces agree. The color spaces in order of increasing quality are (1) YIQ space, (2) YUV space, (3) opponent color space [6], [7], and (4) linearized CIELab color space [8]. The first two color spaces are common in video, where Y refers to the luminance component and the other two letters refer to the two chrominance components. The opponent color space is a part of the pattern-color separable vision model [6], [7] that forms the basis for the industry standard Spatial-CIELab (S-CIELab) color space [7]. We did not consider the S-CIELab color space because we are restricting our attention to color spaces based on a linear transformation from RGB. However, the HVS model with the linearized CIELab color space under consideration [8] closely approximates the S-CIELab framework.

Section II linearizes vector error diffusion and defines the vector error filter optimization problem. Section III generalizes the linear color HVS model [5] to compute the vector error filter coefficients. Section IV describes color space transformations and the associated spatial filters. Sections V and VI report objective and subjective quality measures, respectively. Section VII concludes the letter.

II. Vector Color Error Diffusion

We use $\mathbf{m}$ to denote a 2-D spatial index (m_1, m_2) and $\mathbf{z}$ to denote the z-domain variables (z_1, z_2).

A. Matrix Gain Model

Fig. 1 shows a model of vector color error diffusion halftoning after linearizing the system by replacing the quantizer with matrix gain $\tilde{\mathbf{K}}$ and an additive white noise image $\mathbf{n}(\mathbf{m})$ [5]. The matrix gain is related to the amount of sharpening, and the noise image models the quantization error. $\tilde{\mathbf{K}}$ is chosen to minimize the error in approximating the quantizer with a linear transformation, in the linear minimum mean squared error sense

$$\tilde{\mathbf{K}} = \arg\min_{\tilde{\mathbf{A}}} E[\| \mathbf{b}(\mathbf{m}) - \tilde{\mathbf{A}}\ \mathbf{u}(\mathbf{m}) \|^2] \qquad (1)$$

where $\mathbf{b}(\mathbf{m})$ is the quantizer output process (halftone), and $\mathbf{u}(\mathbf{m})$ is the quantizer input process. When $\mathbf{b}(\mathbf{m})$ and $\mathbf{u}(\mathbf{m})$

are wide sense stationary [9], the solution for (1) is

$$\tilde{\mathbf{K}} = \tilde{\mathbf{C}}_{\mathbf{bu}} \tilde{\mathbf{C}}_{\mathbf{uu}}^{-1} \tag{2}$$

where $\tilde{\mathbf{C}}_{\mathbf{bu}}$ and $\tilde{\mathbf{C}}_{\mathbf{uu}}$ are covariance matrices. The linearized vector error diffusion system, shown in Fig. 1 has two inputs (original signal and quantization noise) and one output (the halftone). Using (2), the signal and noise transfer functions are [5]

$$\mathbf{B_s}(\mathbf{z}) = \tilde{\mathbf{K}} \left[\tilde{\mathbf{I}} + \tilde{\mathbf{H}}(\mathbf{z})(\tilde{\mathbf{K}} - \tilde{\mathbf{I}}) \right]^{-1} \mathbf{X}(\mathbf{z}) \tag{3}$$

$$\mathbf{B_n}(\mathbf{z}) = \left[\tilde{\mathbf{I}} - \tilde{\mathbf{H}}(\mathbf{z}) \right] \mathbf{N}(\mathbf{z}) \tag{4}$$

The overall system response is given by

$$\mathbf{B}(\mathbf{z}) = \mathbf{B_s}(\mathbf{z}) + \mathbf{B_n}(\mathbf{z}) \tag{5}$$

B. Error Measure

We form an objective function J that measures the average visually weighted noise energy in the halftone. The output noise is computed by inverse transforming (4):

$$\mathbf{b_n}(\mathbf{m}) = \left[\tilde{\mathbf{I}} - \tilde{\mathbf{h}}(\mathbf{m}) \right] \star \mathbf{n}(\mathbf{m}) \tag{6}$$

We weight the noise energy by a linear spatially invariant matrix-valued HVS model, $\tilde{\mathbf{v}}(\mathbf{m})$, and form

$$J = E \left[\left\| \mathbf{v}(\tilde{\mathbf{m}}) \star \left[\tilde{\mathbf{I}} - \tilde{\mathbf{h}}(\mathbf{m}) \right] \star \tilde{\mathbf{n}}(\mathbf{m}) \right\|^2 \right] \tag{7}$$

Given a linear spatially invariant HVS model $\tilde{\mathbf{v}}(\mathbf{m})$, the problem is to design an optimal matrix-valued error filter

$$\tilde{\mathbf{h}}_{opt}(\mathbf{m}) = \arg \min_{\tilde{\mathbf{h}}(\mathbf{m}) \in \mathcal{C}} J \tag{8}$$

where the constraint $\mathcal{C}$ enforces the criterion that the error filter diffuses all quantization error [10]

$$\mathcal{C} = \left\{ \tilde{\mathbf{h}}(\mathbf{i}),\ \mathbf{i} \in \mathcal{S} \mid \sum_{\mathbf{i}} \tilde{\mathbf{h}}(\mathbf{i}) \mathbf{1} = \mathbf{1} \right\} \tag{9}$$

$\mathcal{S}$ is the set of coordinates for the error filter support, i.e. $\mathcal{S} = \{(1,0), (1,1), (0,1), (-1,1)\}$ for Floyd-Steinberg.

For each of the four HVS models in Section IV, we obtain the optimum filter coefficients by solving a matrix version of the Yule-Walker equations, based on the constraint $\mathcal{C}$ [5]. We use a descent algorithm based on the quasi-Newton Broyden-Fletcher-Goldfarb-Shanno (BFGS) update [11]. This reduces the number of iterations required to converge to the minimum by about half for the same effort per iteration vs. steepest descent [5].

III. Generalized Perceptual Model

The linear color model employed by Damera-Venkata and Evans [5] is based on the pattern color separable model by Wandell *et al.* [6], [7]. They transfer device dependent RGB values into an opponent representation [7], [12]. The three opponent visual pathways are white-black (luminance pathway) and red-green and blue-yellow (chrominance pathways). By x-y, we mean that in value, x is at one extreme and y is at the other.

We generalize this model as a linear transformation $\tilde{\mathbf{T}}$ to a desired color space, which is not necessarily the opponent representation [6] but any one that satisfies pattern color separability, followed by appropriate spatial filtering in each channel. A complete HVS model is uniquely determined by the color space transformation and associated spatial filters. This generalization provides a platform for evaluation of different models in perceptual meaning and error filter quality obtained by minimizing (7).

The linear color model Fig. 2 consists of (1) a linear transformation $\tilde{\mathbf{T}}$, and (2) separable spatial filtering on each channel. Each channel uses a different spatial filter. The filtering in the z-domain is a matrix multiplication by a diagonal matrix $\mathbf{D}(\mathbf{z})$. In the spatial domain, the linear HVS model $\tilde{\mathbf{v}}(\mathbf{m})$ is computed as

$$\tilde{\mathbf{v}}(\mathbf{m}) = \tilde{\mathbf{d}}(\mathbf{m}) \tilde{\mathbf{T}} \tag{10}$$

The following section discusses the transformation $\tilde{\mathbf{T}}$ and the spatial filters $\tilde{\mathbf{d}}(\mathbf{m})$.

IV. Color Transformations and Spatial Filters

Since we are targeting color halftones for display, we first account for the nonlinear response of a CRT to frame buffer values. We pass the RGB values of original image through this nonlinearity to obtain the RGB values of the colors displayed on the monitor before halftoning. This ensures that the colors in the halftone are closest to the colors actually rendered on the monitor. This process is the inverse of gamma correction. We use "gamma uncorrection" as specified by the sRGB standard [13] on the original image prior to halftoning.

We employ transformations to the following color spaces: linearized CIELab [8], opponent [7], YUV, and YIQ. These color spaces are chosen because they all score well in perceptual uniformity [14] and approximately satisfy the requirements for pattern color separability [6]. Since RGB values are device dependent, we perform the color transformations based on an sRGB monitor. The transformation to opponent color space is:

$$\text{sRGB} \longrightarrow \text{CIEXYZ} \longrightarrow \text{opponent representation}$$

The standard transformations from sRGB to CIEXYZ and from CIEXYZ to opponent representation are taken from the S-CIELab [7] code at

http://white.stanford.edu/~brian/scielab/scielab1-1

which is also the source for transformations to the YUV and YIQ representations. The linearized CIELab color space is obtained by linearizing the CIELab space about the D65 white point [8] in the following manner:

$$Yy = 116 \frac{Y}{Y_n} - 16 \tag{11}$$

$$C_x = 500\left[\frac{X}{X_n} - \frac{Y}{Y_n}\right] \tag{12}$$

$$C_z = 200\left[\frac{Y}{Y_n} - \frac{Z}{Z_n}\right] \tag{13}$$

Hence $\tilde{\mathbf{T}}$ is sRGB $\longrightarrow$ CIEXYZ $\longrightarrow$ Linearized CIELab. The $\mathbf{Y}y$ component is similar to the luminance and the $\mathbf{C}_x$ and $\mathbf{C}_z$ components are similar to the R-G and B-Y opponent color components. The original transformation to the CIELab from CIEXYZ is a non-linear one

$$L^* = 116f\left(\frac{Y}{Y_n}\right) - 16 \tag{14}$$

$$a^* = 500\left[f\left(\frac{X}{X_n}\right) - f\left(\frac{Y}{Y_n}\right)\right] \tag{15}$$

$$b^* = 200\left[f\left(\frac{Y}{Y_n}\right) - f\left(\frac{Z}{Z_n}\right)\right] \tag{16}$$

where

$$f(x) = \begin{cases} 7.787x + \frac{16}{116} & \text{if } 0 \le x \le 0.008856 \\ x^{1/3} & \text{if } 0.008856 < x \le 1 \end{cases}$$

The values for X_n, Y_n and Z_n are as per the D65 white point [13].

The nonlinearity in the transformation from CIELab distorts the spatially averaged tone of the images, which yields halftones that have incorrect average values [8]. The linearized color space overcomes this, and has the added benefit that it decouples the effect of incremental changes in (Yy, Cx, Cz) at the white point on (L, a, b) values:

$$\nabla_{(Yy,Cx,Cz)}(L^*, a^*, b^*)|_{D_{65}} = \frac{1}{3}\mathbf{I} \tag{17}$$

In the opponent color representation, data in each plane is filtered [7] by $2-D$ separable spatial kernels

$$f = k\sum_i \omega_i E_i \tag{18}$$

where $E_i = k_i \exp(-\frac{(x^2+y^2)}{\sigma_i^2})$. The parameters ω_i and σ_i are based on psychophysical testing and are available in [7]. The spatial filters for Linearized CIELab and the YUV and YIQ color spaces are based on analytical models of the eye's luminance and chrominance frequency response.

Nasanen and Sullivan [15] chose an exponential function to model the luminance frequency response

$$W_{(Y_y)}(\tilde{\rho}) = K(L)e^{-\alpha(L)\tilde{\rho}} \tag{19}$$

where L is the average luminance of display, $\tilde{\rho}$ is the radial spatial frequency, $K(L) = a\,L^b$ and

$$\alpha(L) = \frac{1}{c\ln(L) + d} \tag{20}$$

The frequency variable $\tilde{\rho}$ is defined [8] as a weighted magnitude of the frequency vector $\mathbf{u} = (u, v)^T$, where the weighting depends on the angular spatial frequency ϕ [15]. Thus,

$$\tilde{\rho} = \frac{\rho}{s\,(\phi)} \tag{21}$$

Image	*Linear CIELab*	*Opponent*	*YUV*	*YIQ*
lena	3.3511	3.2912	3.2731	3.2778
pasta	1.4623	1.4566	1.4472	1.4383
fruits	1.3832	1.3730	1.3682	1.3644
peppers	1.9562	1.8766	1.8633	1.8545

TABLE I
NOISE GAIN (IN DB) OF THE FLOYD-STEINBERG ERROR FILTER OVER THE OPTIMUM ERROR FILTER FOR FOUR HVS MODELS.

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

$$s\,(\phi) = \frac{1-\omega}{2}\cos(4\phi) + \frac{1+\omega}{2} \tag{22}$$

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 [16]. The chrominance model we use reflects this [17]:

$$W_{(C_x,C_z)}(\rho) = Ae^{-\alpha\rho} \tag{23}$$

Here, α is determined to be 0.419 and $A = 100$ [17]. Both the luminance and chrominance response are lowpass in nature but only the luminance response is reduced at odd multiples of 45^o (Fig. 3). 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.

V. Simulation Results

We evaluate four models based on linearized CIELab, opponent, YUV, and YIQ color spaces using four images. Table I gives the noise shaping gain NG of the optimal filter over the Floyd-Steinberg filter in decibels [5]:

$$NG = 10\log_{10}\left(\frac{J_{fs}}{J_{opt}}\right) \tag{24}$$

Here, J refers to the value of the objective function given by (7). The linearized CIELab model outperforms the other three models. The other three models have similar performance. Matrix-valued error filters for the four HVS models are available at *http://www.ece.utexas.edu/~vishal/halftoning.html*.

VI. Subjective Testing

We develop a subjective assessment procedure for comparing the models based on a pair-comparison task. On each trial, the observer saw an original image above a pair of halftone images positioned to the left and right. The images were displayed on a sRGB monitor. The observer

was forced to pick the halftone that looked most similar to the original. For each original image, observers made six preference judgements, one for each possible pairing of the four halftone images generated using the four different HVS models.

The relevant data from the assessment procedure is the proportion of trials where one given halftoning method is preferred to the other. If one method is not preferred over the other, then the measured proportion should (ideally) be 0.5. Further, to conclude that one method is significantly better, we require it to be picked over the other 60% or more of the time. Using the binomial probability distribution, the total number of trials in the study was picked so that the binomial proportion could be estimated to within a tolerance of 0.03 with 95% confidence. This requirement resulted in a total of 960 comparisons [18], [19] for each of the six possible pairings of the four halftoning methods. The total comparisons/trials were obtained from sixty subjects who judged 96 comparisons each.

Let A, B, C and D be the 4 halftoning methods (or HVS models) to be compared. Based on the subjective test data, we build

$$\mathbf{P} = \begin{bmatrix} 0.5 & p_{AB} & p_{AC} & p_{AD} \\ p_{BA} & 0.5 & p_{BC} & p_{BD} \\ p_{CA} & p_{CB} & 0.5 & p_{CD} \\ p_{DA} & p_{DB} & p_{DC} & 0.5 \end{bmatrix}$$

where p_{AB} is the proportion that method A was preferred to method B. The entries in $\mathbf{P}$ satisfy $P(k,l) = 1 - P(l,k)$. The data from the subjective assessment procedure gives $\mathbf{P}$ as

$$\mathbf{P} = \begin{bmatrix} 0.5 & 0.625 & 0.833 & 0.854 \\ 0.375 & 0.5 & 0.563 & 0.621 \\ 0.167 & 0.437 & 0.5 & 0.521 \\ 0.146 & 0.379 & 0.479 & 0.5 \end{bmatrix}$$

This results in a unique rank order of the methods i.e. $A \geq B \geq C \geq D$, where A refers to the HVS model based on Linearized CIELab, B refers to the model based on the opponent color space, C to the YUV model and D to YIQ. Note that if a rank order of the form $A \geq B \geq C \geq D$ exists, and we have $p_{AB} \geq 0.6$, then we would automatically have $p_{AC} \geq 0.6$ and so forth, and we can conclude that A is significantly better than or preferred to B, C, and D. Similarly, we can find if B is better than C and so on. From the entries in $\mathbf{P}$, we hence conclude that the HVS model based on the linearized CIELab [8] color space is the clear winner. In descending subjective quality,

$$\text{Linearized CIELab} \gg \text{Opponent} > \text{YUV} \approx \text{YIQ}$$

The objective rankings in Table I and subjective rankings given by $\mathbf{P}$ agree. In Table I, YUV is slightly better than YIQ in three of the four images. Likewise, the "perceptual" distinction in the subjective test between YUV and YIQ is not appreciable. The subjective test is available online at *http://www.ece.utexas.edu/~vishal/cgi-bin/test.html*.

VII. Conclusion

For grayscale images, contrast sensitivity functions are linear models of the HVS response vs. spatial frequency. These linear HVS models have been explicitly used in grayscale halftoning for more than a decade to improve halftone quality [20]. For color images, similar linear HVS models exist for luminance channels and certain chrominance channels, which can be applied after a color space transformation. In this letter, we evaluate four color HVS models for color halftoning by vector error diffusion using objective measures and subjective testing. The color HVS model based on transformation to the linearized CIELab [8] color space and spatial filters for the luminance frequency response in [15] and the chrominance frequency response in [17] yields the best error filter for vector error diffusion.

These results may be improved by the choice of a more "perceptually accurate" transformation $\tilde{\mathbf{T}}$ (if available). A similar analysis could be carried out for a printing application. For example, we could first convert the CMY image to a CIEXYZ image [12], then CIEXYZ to the desired color space, and finally apply the spatial filters. The reported procedure for subjective testing of color halftones is general and could be used to compare halftones generated by different methods.

References

[1] R. W. Floyd and L. Steinberg, "An adaptive algorithm for spatial graysacle," *Proc. SID*, vol. 17, no. 2, pp. 75–77, 1976.

[2] R. Ulichney, *Digital Halftoning*. MIT Press, 1987.

[3] H. Haneishi, H. Yaguchi, and Y. Miyake, "A new method of color reproduction in digital halftone image," in *Proc. IS&T Conf.*, (Cambridge, MA), May 1993.

[4] L. Akarun, Y. Yardimci, and A. E. Cetin, "Adaptive methods for dithering color images," *IEEE Trans. on Image Processing*, vol. 6, pp. 950–955, July 1997.

[5] N. Damera-Venkata and B. L. Evans, "Design and analysis of vector color error diffusion halftoning systems," *IEEE Trans. on Image Processing*, vol. 10, pp. 1552–1565, Oct. 2001.

[6] A. B. Poirson and B. A. Wandell, "Appearance of colored patterns: Pattern-color separability," *J. Opt. Soc. Amer. A*, vol. 10, pp. 2458–2470, Dec. 1993.

[7] X. Zhang and B. A. Wandell, "A spatial extension of CIELab for digital color image reproduction," *SID Tech. Dig.*, pp. 731–734, 1996.

[8] 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*, vol. 1913, pp. 270–281, 1993.

[9] H. Stark and J. W. Woods, *Probability, Random Processes and Estimation Theory for Engineers*. Prentice Hall, 1986.

[10] N. Damera-Venkata, *Analysis and Design of Vector Error Diffusion Systems for Image Halftoning*. PhD thesis, Dept. of ECE, Univ. of Texas, Austin, TX 78712 USA, Dec. 2000, www.ece.utexas.edu/~bevans/students/phd/niranjan/.

[11] D. P. Bertsekas, *Nonlinear Programming*. Athena Scientific, 1995.

[12] M. D. Fairchild, *Color Appearance Models*. Addison-Wesley, 1998.

[13] "What is sRGB? - Introduction to the standard default RGB color space developed by Hewlett-Packard and Microsoft." http://www.srgb.com/aboutsrgb.html, 1999.

[14] C. A. Poynton, "Frequently asked questions about colour." http://www.inforamp.net/~poynton/ColorFAQ.html, 1999.

[15] J. R. Sullivan, L. A. Ray, and R. Miller, "Design of minimum visual modulation halftone patterns," *IEEE Trans. Sys. Man. Cyb.*, vol. 21, pp. 33–38, Jan. 1991.

[16] D. H. Kelly, "Spatiotemporal variation of chromatic and achromatic contrast thresholds," *Journal Opt. Soc. Amer. A*, vol. 73, pp. 742–750, 1983.

[17] B. Kolpatzik and C. Bouman, "Optimized error diffusion for high quality image display," *Journal of Electronic Imaging*, vol. 1, pp. 277–292, Jan. 1992.
[18] G. E. P. Box, J. S. Hunter, and W. G. Hunter, *Statistics for Experimenters*. Wiley, 1978.
[19] "STATLETS: The statistical analysis software." *http://www.statlets.com/usermanual/sect5_4_3.htm*.
[20] S. H. Kim and J. P. Allebach, "Impact of human visual system models on model based halftoning," *IEEE Trans. on Image Processing*, vol. 11, pp. 258–269, Mar. 2002.

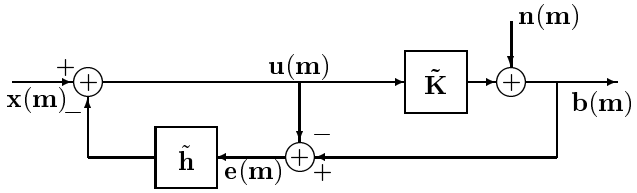

Fig. 1. Linearized model of vector color error diffusion. The quantizer has been replaced by a linear transformation by $\tilde{\mathbf{K}}$ plus additive noise, $\mathbf{n}(\mathbf{m})$, that is uncorrelated with $\mathbf{u}(\mathbf{m})$. The original image is $\mathbf{x}(\mathbf{m})$, and the halftone is $\mathbf{b}(\mathbf{m})$.

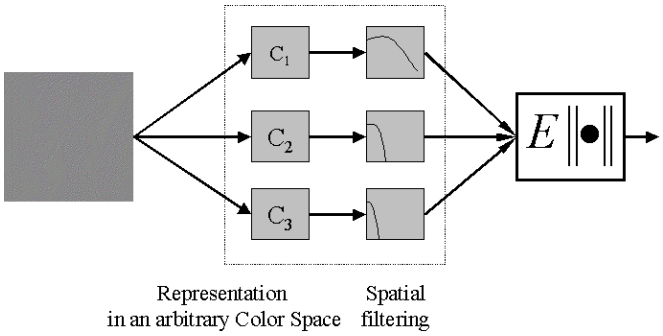

Fig. 2. Generalized Linear Color Model for the Human Visual System

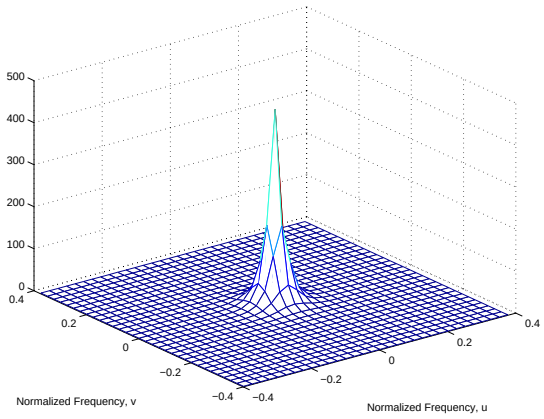

(a)

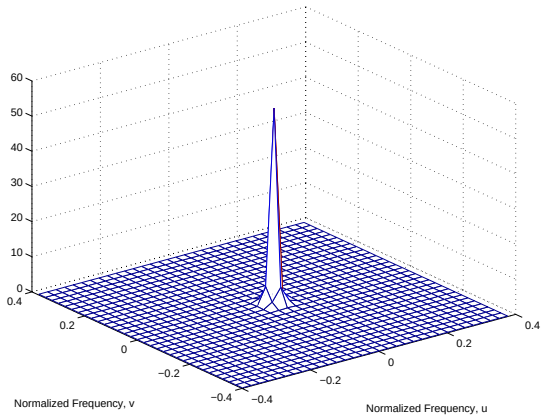

(b)

Fig. 3. Spatial frequency responses (a) luminance $W_{(Y_y)}(\mathbf{u})$ and (b) chrominance $W_{(C_x,C_z)}(\mathbf{u})$

(a) Original peppers image

(b) using linearized CIELab

(c) using opponent color space

(d) using YUV

(e) using YIQ

Fig. 4. Results of halftones of peppers generated by optimizing the vector error filter in the four color spaces. Tiff versions of these images are available at *http://www.ece.utexas.edu/~vishal/halftoning.html*.