Model based color image quantization

T. J. Flohr, B. W. Kolpatzik, R. Balasubramanian,
D. A. Carrara, C. A. Bouman, J. P. Allebach

School of Electrical Engineering
Purdue University
West Lafayette, IN 47907-1285

ABSTRACT

In this paper, we propose a new technique for halftoning color images. Our technique parallels recent work in model-based halftoning for both monochrome and color images; we incorporate a human visual model that accounts for the difference in the responses of the human viewer to luminance and chrominance information. Thus, the *RGB* color space must be transformed to a luminance/chrominance based color space. The color transformation we use is a linearization of the uniform color space $L^*a^*b^*$ which also decouples changes between the luminance and chrominance components. After deriving a tractable expression for total-squared perceived error, we then apply the method of Iterated Conditional Modes (ICM) to iteratively toggle halftone values and exploit several degrees of freedom in reducing the perceived error as predicted by the model.

1 INTRODUCTION

Much recent activity in the area of monochrome digital halftoning has centered upon the use of models for the human visual system. The techniques described in [1]-[7] all produce textures in which the spectral energy is distributed by taking into account the contrast sensitivity of the human viewer as a function of spatial frequency. Even the original error diffusion algorithm [8] does this, although in this case the filter coefficients were designed by trial and error, not by reference to a specific model for the contrast sensitivity of the human viewer.

Although each of the red, green, and blue components of an original color image may be individually halftoned using any one of the above methods, it is known that human viewers are less sensitive to lower frequency chromatic errors than to luminance errors of the same frequency [9]. This fact has been exploited with ordered dither in [10] and [11] and with error diffusion in [12]. The techniques proposed in these papers apply separate luminance and chrominance human visual models to properly transformed color image data. The result is to move the easily perceived luminance error to the chrominance band where these errors are less easily seen.

In this paper, we describe a color halftoning technique that employs a combination of the luminance responses from [1] and [13] and the chrominance response measured in [14]. Our color transformation is a linearized version of the $L^*a^*b^*$ uniform color space; and our error metric is the common squared error criterion which can be minimized using the method described in [2] or by using the method of Iterated Conditional Modes (ICM) described in Sec. 3.3.

2 PERCEPTUAL MODEL

This section describes the model for calculating the perceived halftone image. First we define the opponent color space in which minimization of total-squared perceived error will be performed. This color space is a linearization of the uniform color space $L^*a^*b^*$ which also decouples changes between the luminance and chrominance components. The spatial frequency response of the human visual model is then described. This model is a simple channel-independent model that combines elements of several models reported in the literature.

2.1 Linearized uniform color space

A number of uniform color spaces have been proposed in the literature. For this work, we chose the CIE $L^*a^*b^*$ color space which is more readily linearized than the CIE $L^*u^*v^*$ uniform space. The $L^*a^*b^*$ color space may be

0-8194-1146-9/93/$6.00

expressed in terms of the CIE XYZ color space as [16]

$$\begin{aligned} L^* &= 116\tau\left(\frac{Y}{Y_n}\right) - 16, \\ a^* &= 500\left[\tau\left(\frac{X}{X_n}\right) - \tau\left(\frac{Y}{Y_n}\right)\right], \\ b^* &= 200\left[\tau\left(\frac{Y}{Y_n}\right) - \tau\left(\frac{Z}{Z_n}\right)\right], \end{aligned}$$

where (X_n, Y_n, Z_n) is the D_{65} white point for the XYZ color space and

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

A drawback of the $L^*a^*b^*$ color space is the nonlinearity in the transformation to and from RGB coordinates. Nonlinear color transformations distort the spatially averaged tone of the images, thus leading to halftones that have incorrect average value [17]. This is undesirable since these averages must be equal in both the original and halftone images for faithful tone reproduction. To overcome this problem, we chose to linearize the $L^*a^*b^*$ color space about the D_{65} white point. We denote the linearized color space as $Y_yC_xC_z$, and define it as

$$\begin{aligned} Y_y &= 116\frac{Y}{Y_n} - 16, \\ C_x &= 500\left[\frac{X}{X_n} - \frac{Y}{Y_n}\right], \\ C_z &= 200\left[\frac{Y}{Y_n} - \frac{Z}{Z_n}\right], \end{aligned} \tag{1}$$

where again (X_n, Y_n, Z_n) is the D_{65} white point for the XYZ color space.

An added benefit of this linearization is that it decouples the effect of incremental changes in (Y_y, C_x, C_z) at the white point on the (L^*, a^*, b^*) values. This may be seen by evaluating the gradient of the $L^*a^*b^*$ coordinates with respect to the $Y_yC_xC_z$ coordinates at the white point:

$$\nabla_{(Y_y, C_x, C_z)}(L^*, a^*, b^*)\,|_{D_{65}} = \frac{1}{3}\,\mathbf{I},$$

where $\mathbf{I}$ is the identity matrix. Note that the Y_y component in our color space is proportional to 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 [14] are based. This supports the application of a chrominance spatial frequency response based on her data to the C_x and C_z color components.

2.2 Spatial frequency response

In choosing our model for the spatial dependence of the perception of luminance and chrominance information, we restricted ourselves to a linear, channel-independent model based upon contrast sensitivity measurements for the human viewer. Much experimental data is available for the contrast sensitivity of the human eye in response to spatial variations in luminance. For our luminance model, we used a combination of the models reported by Näsänen and Sullivan *et al* [1] [13]. Näsänen chose an exponential function to model the luminance frequency response:

$$W_{(Y_y)}(\tilde{\rho}) = K(L)\exp\left[-\alpha(L)\tilde{\rho}\right],$$

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

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

$a = 131.6, b = 0.3188, c = 0.525, d = 3.91$, and the subscript Y_y designates the luminance frequency response. We deviated from Näsänen by defining the frequency variable $\tilde{\rho}$ as a weighted magnitude of the frequency vector $\mathbf{u} = (u, v)^T$, where the weighting depends on angular spatial frequency ϕ [1]. Thus,

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

where $\rho = (u^2 + v^2)^{1/2}$ and

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

The symmetry parameter w is 0.7; and the angle ϕ is defined as

$$\phi = \arctan\left(\frac{v}{u}\right).$$

The weighting function $s(\phi)$ effectively reduces the contrast sensitivity to spatial frequency components at odd multiples of 45°.

It is well known that 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. Our chrominance model reflects this. We used an approximation to the experimental results obtained by Mullen [14] to model the response to chrominance. This approximation to the data is due to Kolpatzik and Bouman [18], and is given by

$$W_{(C_x, C_z)}(\rho) = A\exp\left[-\alpha\rho\right],$$

where they determined $\alpha = 0.419$ and $A = 100$. The radial spatial frequency ρ is again defined as $\rho = (u^2 + v^2)^{1/2}$. Note that the subscript (C_x, C_z) denotes chrominance frequency response.

Figure 1 shows plots of the magnitude-squared of the luminance and chrominance frequency responses. Both are lowpass in nature; but only the luminance response is reduced at odd multiples of 45°. This will place more luminance error along diagonals in the frequency domain. The chrominance response has a narrower bandwidth than the luminance response. Using this chrominance response as opposed to identical responses for both luminance and chrominance will allow more lower frequency chromatic error, which will not be perceived by the human viewer.

The linearized color space $Y_y C_x C_z$ preserves the relative weighting of the color components in $L^*a^*b^*$. However, the $L^*a^*b^*$ color space is based on color matching experiments conducted with relatively large color patches. Thus, it is only appropriate for low spatial frequencies. To allow adjustment of the relative weight between the luminance and chrominance responses, we multiplied the luminance frequency response by a weighting factor k. As k increases, more error will be forced into the chrominance components. We investigate the effect of this parameter in Sec. 4.

3 ERROR MINIMIZATION

This section first introduces some notation and briefly discusses the output device model. We then describe the error metric that is minimized in obtaining the digital color halftone image. This metric is the widely used total-squared error criterion. However, like much recent work in halftoning research, the error is the perceived image error based on the perceptual model described in Sec. 2.2. Thus, the actual error image must be implicitly convolved with the point spread function for the human visual system. After describing the error metric, the methodology for minimizing the perceived error between the original image and the halftone image is briefly outlined.

3.1 Notation and device model

Throughout the remainder of this paper, boldface variables will represent vector or matrix valued entities and non-boldface variables will depict scalar entities. The vectors $\mathbf{x} = (x, y)^T \in \Re^2$ and $\mathbf{m} = (m, n)^T \in Z^2$ will specify continuous and discrete spatial coordinates respectively, where $\Re$ and Z are the sets of real and integer numbers. We assume that the output display device is characterized by a lattice of addressable points $\mathbf{x_m} = \mathbf{Xm}$, $\mathbf{m} \in Z^2$, where

$\mathbf{X}$ is the periodicity matrix, the columns of which comprise a basis for the lattice. At each point on this lattice, we assume that a spot with intensity profile $p(\mathbf{x})$ can be generated for each of the three RGB primaries. For simplicity, we assume that

$$p(\mathbf{x}) = \begin{cases} 1, \mathbf{x} \in \chi, \\ 0, \text{else}, \end{cases}$$

where χ denotes the unit cell associated with the lattice $\mathbf{X}$. It should be pointed out that at sufficiently large viewing distances or high device resolutions, the exact shape of the output spot profile $p(\mathbf{x})$ will be inconsequential, since it will be dominated by the point spread function associated with the model for the human viewer. The vector $\mathbf{f}[\mathbf{n}]$ shall denote the continuous-tone digital color image we wish to display and $\mathbf{f}(\mathbf{x})$ shall represent a continuous-space version of $\mathbf{f}[\mathbf{n}]$. Likewise, the halftone versions will be denoted by the digital image $\mathbf{g}[\mathbf{n}]$ and the continuous-space image $\mathbf{g}(\mathbf{x})$. Note that arguments of functions of continuous-space variables are enclosed in parentheses, whereas arguments of functions of discrete-space variables are enclosed in brackets. Finally, perceived entities will be marked with a tilde and subscripts to the image vectors will denote the color space in which the images are represented. For example, $\tilde{\mathbf{f}}_{(Y_y,C_x,C_z)}(\mathbf{x})$ will specify the perceived image in the color space $Y_yC_xC_z$.

3.2 Error metric

If the output device could render continuous-tone directly, the displayed image would be given by

$$\mathbf{f}_{(R,G,B)}(\mathbf{x}) = \sum_{\mathbf{n}} \mathbf{f}_{(R,G,B)}[\mathbf{n}]p(\mathbf{x} - \mathbf{X}\mathbf{n}).$$

Since the output device has only bi-level capability with respect to each primary, we must display

$$\mathbf{g}_{(R,G,B)}(\mathbf{x}) = \sum_{\mathbf{n}} \mathbf{g}_{(R,G,B)}[\mathbf{n}]p(\mathbf{x} - \mathbf{X}\mathbf{n}),$$

where each element of the vector $\mathbf{g}[\mathbf{n}] = 0$ or 1. The goal of the halftoning process is to minimize the perceived difference between $\mathbf{f}_{(Y_y,C_x,C_z)}(\mathbf{x})$ and $\mathbf{g}_{(Y_y,C_x,C_z)}(\mathbf{x})$. Transforming these images to the $Y_yC_xC_z$ space, we obtain

$$\mathbf{f}_{(Y_y,C_x,C_z)}(\mathbf{x}) = \sum_{\mathbf{n}} \mathbf{f}_{(Y_y,C_x,C_z)}[\mathbf{n}]p(\mathbf{x} - \mathbf{X}\mathbf{n}), \tag{2}$$

and

$$\mathbf{g}_{(Y_y,C_x,C_z)}(\mathbf{x}) = \sum_{\mathbf{n}} \mathbf{g}_{(Y_y,C_x,C_z)}[\mathbf{n}]p(\mathbf{x} - \mathbf{X}\mathbf{n}).$$

We then define the matrix $\mathbf{W}(\mathbf{x})$ as

$$\mathbf{W}(\mathbf{x}) = \begin{bmatrix} w_{(Y_y)}(\mathbf{x}) & 0 & 0 \\ 0 & w_{(C_x)}(\mathbf{x}) & 0 \\ 0 & 0 & w_{(C_y)}(\mathbf{x}) \end{bmatrix},$$

where

$$\begin{aligned} w_{(Y_y)}(\mathbf{x}) &= \mathcal{F}^{-1}\left[W_{(Y_y)}(\mathbf{u})\right], \\ w_{(C_x)}(\mathbf{x}) &= \mathcal{F}^{-1}\left[W_{(C_x)}(\mathbf{u})\right], \\ w_{(C_y)}(\mathbf{x}) &= \mathcal{F}^{-1}\left[W_{(C_y)}(\mathbf{u})\right], \end{aligned}$$

and $\mathcal{F}^{-1}$ denotes the inverse continuous-space Fourier transform defined by $f(\mathbf{x}) = \int F(\mathbf{u})e^{j2\pi\mathbf{u}^T\mathbf{x}}d\mathbf{u}$. The perceived displayed continuous-tone image may then be expressed as

$$\tilde{\mathbf{f}}_{(Y_y,C_x,C_z)}(\mathbf{x}) = \mathbf{W}(\mathbf{x}) * \mathbf{f}_{(Y_y,C_x,C_z)}(\mathbf{x}) = \int_{\mathbf{t}} \mathbf{W}(\mathbf{t})\mathbf{f}_{(Y_y,C_x,C_z)}(\mathbf{x} - \mathbf{t})d\mathbf{t}. \tag{3}$$

This expression should be interpreted as a matrix-vector product in which each multiplication between elements of W and p is replaced by convolution. Combining (2) and (3) gives

$$\tilde{\mathbf{f}}_{(Y_y,C_x,C_z)}(\mathbf{x}) = \sum_{\mathbf{n}} \tilde{\mathbf{P}}(\mathbf{x} - \mathbf{X}\mathbf{n})\mathbf{f}_{(Y_y,C_x,C_z)}[\mathbf{n}]. \tag{4}$$

where

$$\tilde{\mathbf{P}}(\mathbf{x}) = \mathbf{W}(\mathbf{x}) * p(\mathbf{x}),$$

is a matrix for which each element is the convolution of the corresponding element of $\mathbf{W}(\mathbf{x})$ with the scalar $p(\mathbf{x})$. Similarly, we may express the perceived halftone image as

$$\tilde{\mathbf{g}}_{(Y_y,C_x,C_z)}(\mathbf{x}) = \sum_{\mathbf{n}} \tilde{\mathbf{P}}(\mathbf{x} - \mathbf{X}\mathbf{n})\mathbf{g}_{(Y_y,C_x,C_z)}[\mathbf{n}]. \tag{5}$$

As a measure of the difference between $\tilde{\mathbf{f}}_{(Y_y,C_x,C_z)}$ and $\tilde{\mathbf{g}}_{(Y_y,C_x,C_z)}$, we use the total-squared perceived error

$$\begin{aligned} E &= \|\tilde{\mathbf{e}}_{(Y_y,C_x,C_z)}(\mathbf{x})\|^2, \\ &= \int_{\mathbf{x}} (\tilde{\mathbf{e}}_{(Y_y,C_x,C_z)}(\mathbf{x}))^T (\tilde{\mathbf{e}}_{(Y_y,C_x,C_z)}(\mathbf{x}))d\mathbf{x}, \end{aligned} \tag{6}$$

where

$$\tilde{\mathbf{e}}_{(Y_y,C_x,C_z)}(\mathbf{x}) = \tilde{\mathbf{f}}_{(Y_y,C_x,C_z)}(\mathbf{x}) - \tilde{\mathbf{g}}_{(Y_y,C_x,C_z)}(\mathbf{x}). \tag{7}$$

A block diagram for the calculation of (6) appears in Fig. 2.

3.3 Error minimization technique

The method we used to minimize (6) was Iterated Conditional Modes (ICM). Starting with a random initial image, we iteratively visit each pixel $\mathbf{n}_0$ in the image and toggle $\mathbf{g}[\mathbf{n}_0]$ to each of its eight color possiblities. To determine the effect on the overall error E of changing just $\mathbf{g}[\mathbf{n}_0]$, we substitute (4) and (5) into (7) and separate out the term $\mathbf{n} = \mathbf{n}_0$ under the summation, obtaining

$$\tilde{\mathbf{e}}_{(Y_y,C_x,C_z)}(\mathbf{x}) = \sum_{\mathbf{n}\neq\mathbf{n}_0} \tilde{\mathbf{P}}(\mathbf{x} - \mathbf{X}\mathbf{n})\mathbf{e}_{(Y_y,C_x,C_z)}[\mathbf{n}] + \tilde{\mathbf{P}}(\mathbf{x} - \mathbf{X}\mathbf{n}_0)\mathbf{e}_{(Y_y,C_x,C_z)}[\mathbf{n}_0]. \tag{8}$$

Using this in (6) then leads to

$$\begin{aligned} E &= \int_{\mathbf{x}} \sum_{\mathbf{n}\neq\mathbf{n}_0} \sum_{\mathbf{m}\neq\mathbf{n}_0} \mathbf{e}^T_{(Y_y,C_x,C_z)}[\mathbf{n}]\tilde{\mathbf{P}}^T(\mathbf{x} - \mathbf{X}\mathbf{n})\tilde{\mathbf{P}}(\mathbf{x} - \mathbf{X}\mathbf{m})\mathbf{e}_{(Y_y,C_x,C_z)}[\mathbf{m}]d\mathbf{x} \\ &+ 2\mathbf{e}^T_{(Y_y,C_x,C_z)}[\mathbf{n}_0] \sum_{\mathbf{n}} \mathbf{C}_{\tilde{\mathbf{P}},\tilde{\mathbf{P}}}[\mathbf{n}_0 - \mathbf{n}]\mathbf{e}_{(Y_y,C_x,C_z)}[\mathbf{n}] + \mathbf{e}^T_{(Y_y,C_x,C_z)}[\mathbf{n}_0]\mathbf{C}^T_{\tilde{\mathbf{P}},\tilde{\mathbf{P}}}[\mathbf{0}]\mathbf{e}_{(Y_y,C_x,C_z)}[\mathbf{n}_0], \end{aligned} \tag{9}$$

where $\mathbf{C}_{\tilde{\mathbf{P}},\tilde{\mathbf{P}}}(\mathbf{y}) = \int \tilde{\mathbf{P}}^T(\mathbf{x})\tilde{\mathbf{P}}(\mathbf{x} + \mathbf{y})d\mathbf{x}$ is the autocorrelation function of $\tilde{\mathbf{P}}(\mathbf{x})$ and $\mathbf{C}_{\tilde{\mathbf{P}},\tilde{\mathbf{P}}}[\mathbf{m}] = \mathbf{C}_{\tilde{\mathbf{P}},\tilde{\mathbf{P}}}(\mathbf{X}\mathbf{m})$. Equation (9) is a quadratic function of the error $\mathbf{e}_{(Y_y,C_x,C_z)}[\mathbf{n}_0]$, and thus may easily be evaluated for the eight choices for $\mathbf{g}[\mathbf{n}_0]$. We pick that value which minimizes E, move on to the next pixel, and repeat the process. The algorithm converges when no changes are made to $\mathbf{g}[\mathbf{n}_0]$ during a complete pass through the image. Note that this method is similar to DBS. The only differences are that DBS employs a slightly different search heuristic, and uses an efficient lookup-table based implementation to evaluate the effect of changes to $\mathbf{g}[\mathbf{n}]$.

The key to the ICM algorithm is the autocorrelation of the visual point spread function. This autocorrelation is obtained by taking the inverse 2 D Discrete Fourier transform of the sampled magnitude-squared frequency response described in Sec. 2.2. In order for the computation in (9) to be tractable, the spatial extent of the autocorrelation must be limited. Simply truncating it leads to large sidelobes in the frequency responses which limit the ability of ICM to move the spectral error energy into the stopbands. This impedes the convergence of ICM. To rectify this dilemma, we smoothly tapered the autocorrelation with a 2-D Hamming window, thus reducing sidelobes in the luminance and chrominance stop bands. Section 4 demonstrates the importance of the windowing step.

4 RESULTS

The continuous-tone color original image we chose to halftone was the "Balloon" girl image [1]. To show the texture of the color halftones produced by our ICM procedure, we present monochrome images containing the Y_y and C_x components of the final halftone images. 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 then transformed back to the RGB space for display. To show only the C_x component of the final halftone images, both the Y_y and C_z components are set to zero. The resulting vector is then transformed back to the RGB space. This resulting RGB vector will not be a monochrome image due to the chromaticity of the C_x component. Therefore, the RGB components are normalized to the maximum chromaticity $C_x = [0, 1, 0]$ to be printed in this black and white medium. Note that the C_y component could also be shown, but similar results would be obtained.

When ICM is used to produce color halftones, several degrees of freedom play a role in the minimization of the perceived error. One degree comes from the relative weightings in the $Y_yC_xC_z$ color transformation in (1). This transformation weights the C_x component more than the C_z component, thus this relative weighting will move error from the C_x component to the C_z component. The frequency responses for both the luminance and chrominance components give another degree of freedom for reducing perceived error. The luminance frequency response described in Sec. 2.2 has the reduced diagonal response that is inherent in many luminance models. Therefore, some luminance error in the halftones produced by ICM will be placed along the diagonals. The luminance response is also bandlimited, so error will appear in the high frequencies of luminance. The chrominance response is also lowpass; thus chrominance error will be moved to high frequencies, but not as high as the luminance errors. The last degree of freedom in our model is the weighting factor k between the luminance and chrominance frequency responses. As k increases, more error will be moved from luminance to chrominance. Our experimental results are intended to show how ICM exploits these degrees of freedom in minimizing the error.

Figures 3-6 compare color images halftoned with error diffusion and ICM. Figures 3 and 4 show the Y_y components. In comparing these two figures, we see that the worm-like textures common to error diffused images is not present with ICM. Surprisingly, the texture in the ICM image is more visible than that in the error diffusion image in the regions of the dress and the balloon to the left of the girl. At this time, we do not have an explanation for this phenomenon, although it is interesting to note that these are both high luminance regions of the image. The dress is a bright white and the balloon is a bright yellow.

Figures 5 and 6 show the C_x components for the same two images. When displayed as luminance variation, the texture is much more visible than when shown in proper chromatic form. Comparing Figs. 3 and 5 for error diffusion, we see that both the Y_y and C_y components exhibit similar texture. In contrast, we see in Figs. 4 and 6 that with ICM the C_x component has much lower frequency texture than the Y_y component. When viewed as true color images, the effect is that the overall texture is less visible in the ICM image than in the error diffusion image, except on the dress and the yellow balloon as noted previously.

Figures 7 and 8 show a similar but more dramatic effect as the weighting k of luminance relative to chrominance is varied. In Fig. 7 which shows the Y_y component, we see that the texture steadily decreases with increasing values of k. At the same time, the texture in the C_x component steadily increases. When viewed in true color, the effect on the balloons is particularly striking. Here the texture shifts from being relatively contrasty and monochromatic to being softer and multihued. Overall, we have found that $k = 4$ appears to be a good value to use.

The final figure illustrates the importance of the Hamming window. It shows the Y_y and C_x components of a color halftone image produced when the autocorrelation function of the visual point spread function is simply truncated rather than being smoothly tapered. Although the algorithm has converged reasonably well over most of the image after 200 iterations, there is a conspicuous artifact in the lower left corner which would only disappear after many more iterations. In contrast, when the Hamming window is applied, satisfactory convergence is obtained with 20 to 30 iterations.

5 CONCLUSIONS

The color halftoning technique described herein effectively reduces the total-squared perceived error in both the luminance and chrominance components of the resulting halftone images by implicitly convolving properly color

[1] Image provided through courtesy of Eastman Kodak Company.

transformed image data with a visual luminance/chrominance point spread function. The color transformation and separate luminance/chrominance frequency responses offer several degrees of freedom for allocating the error between the continuous-tone original and halftone images to different spatial frequencies in luminance and chrominance. Our experimental results demonstrate that the method of Iterated Conditional Modes (ICM) can effectively exploit these degrees of freedom in minimizing the error. Although this implementation is computationally intensive, the same ideas could be applied with Direct Binary Search (DBS) at a greater cost in terms of complexity in code, but a much lower computational cost. The unexplained tendency of ICM to produce textures at high luminance levels that are more visible than those of error diffusion suggests that further investigation of the visual model is needed. Another possibility is that at high luminance levels the exact display spot profile, which was ignored in this work, becomes more important. This would be consistent with results observed earlier when using DBS to halftone monochromatic images [2].

6 REFERENCES

[1] J. R. Sullivan, L. A. Ray, and R. Miller, "Design of Minimum Visual Modulation Halftone Patterns," *IEEE Trans. Syst. Man. Cyb.*, Vol. 21, No. 1, January 1991, pp. 33-38.

[2] M. Analoui and J. P. Allebach, "Model-based Halftoning by Direct Binary Search," *Proceedings of the 1992 SPIE/IS&T Symposium on Electronic Imaging Science and Technology*, San Jose, CA, February 9-14, 1992, pp. 96-108.

[3] J. B. Mulligan and A. J. Ahumada, Jr., "Principled Halftoning Based on Models of Human Vision," *Proceedings of the 1992 SPIE/IS&T Symposium on Electronic Imaging Science and Technology*, San Jose, CA, February 9-14, 1992, pp. 109-121.

[4] C. H. Chu and W. Watunyuta, "Designing Dither Patterns Using A Human Visual Model," *Proceedings of the 1992 SPIE/IS&T Symposium on Electronic Imaging Science and Technology*, San Jose, CA, February 9-14, 1992, pp. 134-143.

[5] B. Kolpatzik and C. A. Bouman, "Optimized Error Diffusion Based on a Human Visual Model," *Proceedings of the 1992 SPIE/IS&T Symposium on Electronic Imaging Science and Technology*, San Jose, CA, February 9-14, 1992, pp. 152-164.

[6] T. N. Pappas, D. L. Neuhoff, and N. Seshadri, "Least-Squares Model-Based Halftoning," *Proceedings of the 1992 SPIE/IS&T Symposium on Electronic Imaging Science and Technology*, San Jose, CA, February 9-14, 1992, pp. 165-176.

[7] D. A. Carrara, M. Analoui, and J. P. Allebach, "Recent Progress in Digital Halftoning," *Proceedings of IS&T's 8th Congress on Advances in Non-Impact Printing Technologies*, Williamsburg, VA, October 25-30, 1992, pp. 265-269.

[8] R. W. Floyd and L. Steinberg, "An Adaptive Algorithm for Spatial Greyscale," *Proc. SID*, Vol. 17, No. 2, 1976, pp. 75-77.

[9] D. H. Kelly, "Spatiotemporal Variation of Chromatic and Achromatic Contrast Thresholds," *J. Opt. Soc. Am.*, Vol. 73, 1983, pp. 742-750.

[10] J. B. Mulligan, "Digital Halftoning Methods for Selectively Partitioning Error into Achromatic and Chromatic Channels," *Proceedings of the 1990 SPIE/IS&T Symposium on Electronic Imaging Science and Technology*, Santa Clara, CA, February 12-14, 1990, pp. 261-270.

[11] J. B. Mulligan and A. J. Ahumada, Jr. "Principled Methods for Color Dithering Based on Models of the Human Visual System," *SID International Symposium, Digest of Technical Papers*, Boston, MA, May 17-22, 1992, pp. 194-197.

[12] R. Miller and J. R. Sullivan, "Color Halftoning Using Error Diffusion and a Human Visual System Model," *Proceedings of the SPSE 43rd Annual Meeting*, Rochester, NY, 1990, pp. 149-152.

[13] R. Näsänen, "Visibility of Halftone Dot Textures," *IEEE Trans. Syst. Man. Cyb.*, Vol. 14, No. 6, 1984, pp. 920-924.

[14] K. T. Mullen, "The Contrast Sensitivity of Human Color Vision to Red-Green and Blue-Yellow Chromatic Gratings," *J. Physiol. 359*, 1985, pp. 381-400.

[15] R. S. Gentile, E. Walowit, and J. P. Allebach, "Quantization and Multilevel Halftoning of Color Images for Near-Original Image Quality," *J. Opt. Soc. Am.*, Vol. 7, No. 6, June 1990, pp. 1019-1026.

[16] G. Wyszecki and W. S. Stiles, *Color Science: Concept and Methods, Quantitative Data and Formulae*, Wiley, New York, 1982.

[17] R. S. Gentile, J. P. Allebach, and E. Walowit, "Quantization of Color Images Based on Uniform Color Spaces," *J. of Imaging Technology*, Vol. 16, No. 1, February 1990, pp. 11-21.

[18] B. W. Kolpatzik and C. A. Bouman, "Optimized Error Diffusion for Image Display," *J. of Electronic Imaging*, Vol. 1, No. 3, July 1992, pp. 277-292.

7 FIGURES

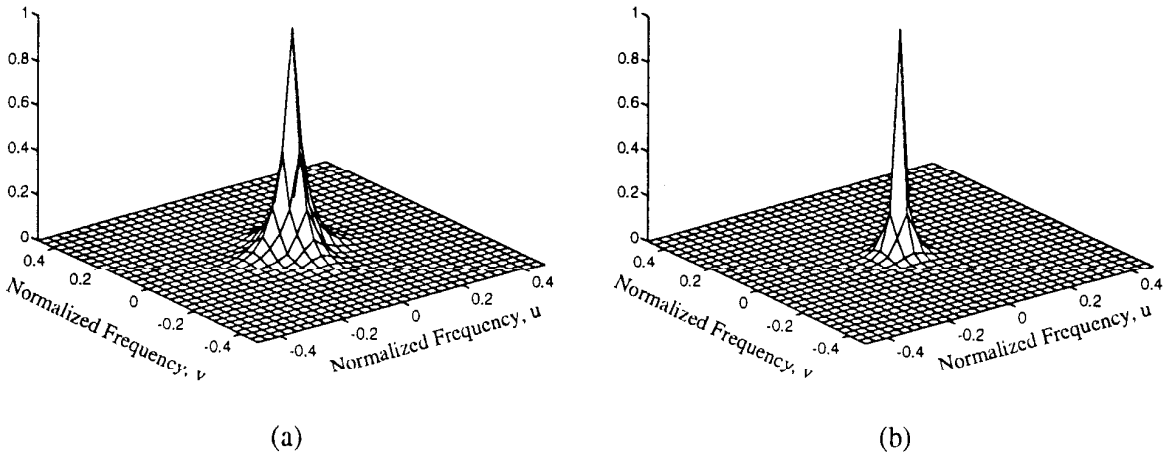

Figure 1: Spatial frequency responses (magnitude-squared): (a) luminance $|W_{Y_y}(\mathbf{u})|^2$, (b) chrominance $|W_{(C_x,C_z)}(\mathbf{u})|^2$.

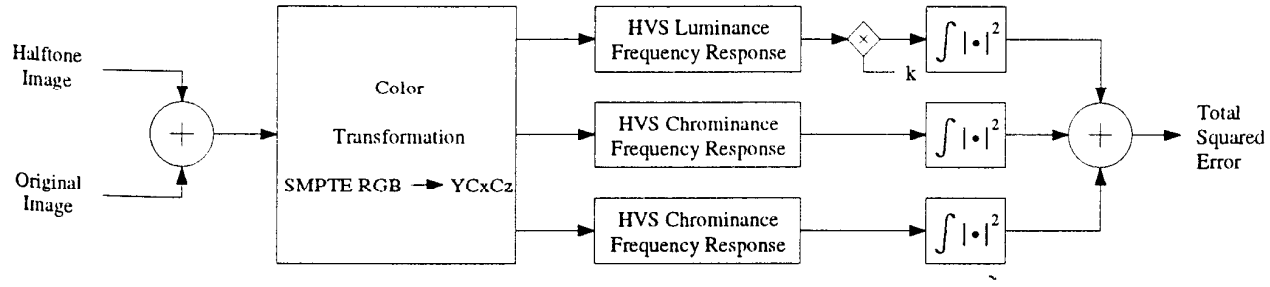

Figure 2: Block diagram for calculating perceptual error metric.

Figure 3: Y_y component of error diffusion color halftone: Balloon halftoned by error diffusion.

Figure 4: Y_y component of ICM color halftone: Balloon halftoned in $Y_y C_1 C_2$ with separate luminance/chrominance frequency responses and $k = 4$.

Figure 5: C_z component of error diffusion color halftone: Balloon halftoned by error diffusion.

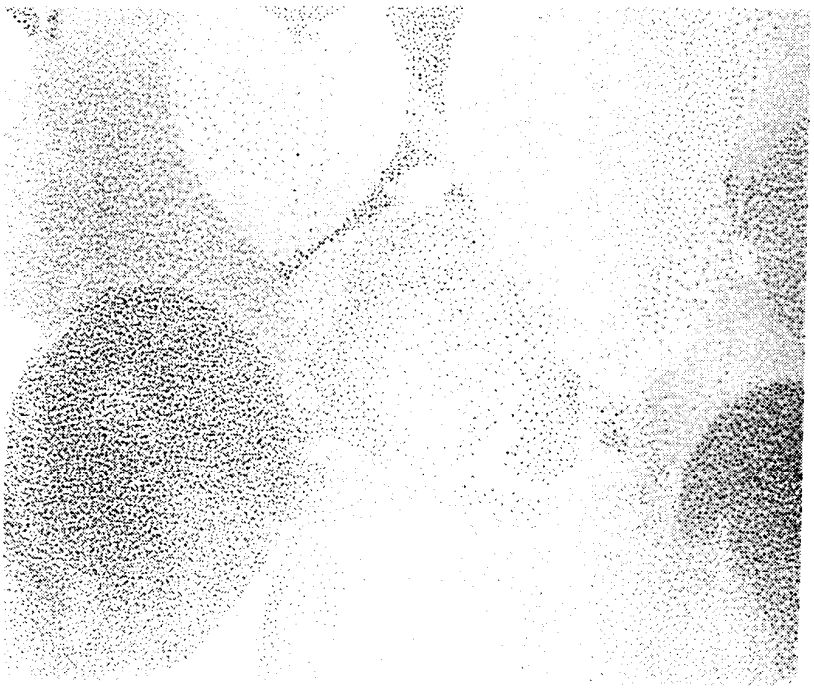

Figure 6: C_x component of ICM color halftone: Balloon halftoned in $Y_y C_x C_z$ with separate luminance/chrominance models and $k = 4$.

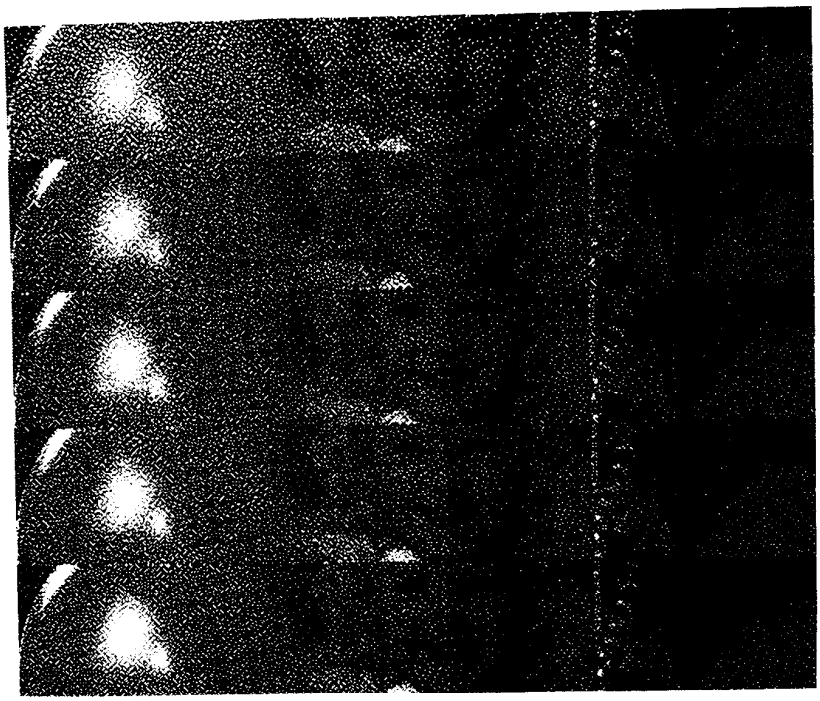

Figure 7: Effect of varying k on Y_y component: Balloon halftoned in $Y_y C_x C_z$ with separate luminance/chrominance frequency responses, and from top to bottom, $k = 1, 2, 4, 8, 30$.

Figure 8: Effect of varying k on C_x component: Balloon halftoned in $Y_y C_x C_z$ with separate luminance/chrominance frequency responses, and from top to bottom, $k = 1, 2, 4, 8, 30$.

Figure 9: Effect of simple truncation of point spread autocorrelation function (Y_y component | C_x component): Balloon halftoned in $Y_y C_x C_z$ with separate luminance/chrominance frequency responses and $k = 4$.