Trade-offs Between Color Saturation and Noise Sensitivity in Image Sensors

Poorvi Vora and Cormac Herley
Hewlett-Packard Labs., {poorvi,cormac}@hpl.hp.com

1 Abstract

A color correction matrix is usually used to transform the raw color image obtained from color image sensors to adjust for factors such as variations in illumination and deviations of the actual filter characteristics from the ideal. Color correction matrices that have high condition numbers can greatly amplify noise, while it can be difficult to achieve good color saturation with low condition number matrices. Noise can be a significant problem for digital camera images which have limited bit-depth. We explore the trade-off between noise reduction and color saturation analytically, and using digital camera images. We present an orthonormality criterion to define optimality with respect to noise sensitivity and color saturation in the case of signal-independent, orthogonal sensor measurement noise in each color channel.

2 Introduction

In image acquisition devices color filters are generally placed on top of monochrome sensors to filter the incident light, and give the triplet of pixel values required to make up a color image. Digital cameras, for example, use three or more color filters on a single sensor. The colors that are generated for a particular surface depend on the spectral characteristics of the illuminant, and the transmissivities of the color filters. A color correction stage is usually required to transform the image from the acquired color space; this transformation can be used to produce more or less saturated colors, to adjust for differences that occur because of the illumination, and to allow for deviations of the actual filter characteristics from the ideal. Generally images with saturated colors are more appealing to end users. In this paper we show that, for certain types of color filter choices, color correction matrices with high condition numbers are correlated with the presence of highly saturated colors in the resulting image. A high condition number in the color correction stage results in noise amplification, which is a serious issue for image sensors with limited bit depth - such as those found in consumer digital cameras. We also examine the tradeoff between the color saturation that can be achieved with a color correction matrix, and the noise amplification that it implies.

A number of criteria exist for the evaluation of the color fidelity properties of a particular set of color filters [3, 7]. In addition, there has been work on the color correction and noise sensitivity properties of a color filter set [5, 6, 1, 2], but existing work does not explicitly connect noise sensitivity to filter orthonormality, which is one of the goals of this paper. We present an orthonormality criterion to define optimality with respect to noise sensitivity and color saturation in the case of signal-independent, orthogonal sensor measurement noise in each color channel. This criterion is identical to the criterion of [5], where it was derived for optimality with respect to perturbations in the filter transmissivities. We derive this criterion for a different purpose from that of [5], and in a simpler manner, by maximizing the minimum ratio of expected signal to expected noise. To our knowledge, work on the relationship of this criterion with noise sensitivity and color saturation does not exist.

For color filter sets that are non-optimal with respect to noise sensitivity, a natural solution to the problem of large noise amplification using the linear minimum mean square error (LMMSE), or Wiener, estimate [5] is to increase the estimate of the noise variance in the calculation of the color correction matrix [1]. We demonstrate, using analysis and real digital camera images, that the noise reduction comes at the cost of color saturation. This is why sets of filters that are non-optimal with respect to noise sensitivity, even if span the Human Visual Subspace (HVSS) and have a perfect Vora-Trussell measure [3], present an image capture problem which is difficult to overcome through post-processing.

0-8186-8821-1/98 $10.00 © 1998 IEEE

3 Notation

This paper follows the notation of [3]. Reflective and radiant spectra are represented by finite-dimensional vectors consisting of samples of the spectra as a function of wavelength over the visual range (approximately 400-700 nm.). The analysis presented here does not depend on the number of samples, N, though the results presented were calculated for N = 31 (corresponding to a sample every 10 nm. from 400-700 nm.). The matrix $\mathbf{M}_H$ denotes the 'effective recording system', each column representing the combination of a recording filter response, the recording illuminant and the recording optical path. Similarly the matrix $\mathbf{A}_L$ denotes the combination of the CIE matching functions or the monitor phosphors ($\mathbf{A}$ with a specified viewing illuminant. $\mathbf{t} = \mathbf{A}_L^T\mathbf{f}$ represents the CIE/monitor tristimulus vector of reflective spectrum $\mathbf{f}$ under the viewing illuminant and is the goal of the capture process, because it is a standard representation of the color stimulus of the reflective surface represented by $\mathbf{f}$ in the specified viewing conditions. The vector $\mathbf{n}$ represents measurement noise, and the vector $\mathbf{g}$ represents the measurements, $\mathbf{g} = \mathbf{M}_H^T\mathbf{f} + \mathbf{n}$.

We use a number of symbols in the rest of this paper which we define here and also where they occur:
$\mathbf{R}$ and $\mathbf{R}_n$ - the correlation matrices of $\mathbf{f}$ and $\mathbf{n}$ respectively
$\mathbf{C}$ - the color correction matrix
λ_i and ζ_i - eigenvalues of $(\mathbf{M}_H^T\mathbf{R}\mathbf{M}_H)\mathbf{R}_n^{-1}$ and $\mathbf{M}_H^T\mathbf{R}\mathbf{M}_H$ respectively
σ^2 - noise variance
κ - condition number of $\mathbf{M}_H^T\mathbf{R}\mathbf{M}_H$

4 A criterion for noise sensitivity

The LMMSE estimate of the tristimulus values from the measurements [5] for zero-mean signal and zero-mean noise is:

$$\mathbf{A}_L^T\mathbf{R}\mathbf{M}_H(\mathbf{M}_H^T\mathbf{R}\mathbf{M}_H + \mathbf{R}_n)^{-1}\mathbf{g} \tag{1}$$

and the color correction matrix, $\mathbf{C}$, is:

$$\mathbf{C} = \mathbf{A}_L^T\mathbf{R}\mathbf{M}_H(\mathbf{M}_H^T\mathbf{R}\mathbf{M}_H + \mathbf{R}_n)^{-1} \tag{2}$$

where $\mathbf{R}$ is the expected value of $\mathbf{f}\mathbf{f}^T$ and $\mathbf{R}_n$ that of $\mathbf{n}\mathbf{n}^T$. The incorporation of non-zero signal and noise means does not change the basic results of our analysis.

The min-max method of [5] may be used to analyze the ratio of expected signal power to expected noise power ($S_{exp}N_{exp}R$) to provide an expression for the noise sensitivity in terms of filter orthogonality as follows. The value of $S_{exp}N_{exp}R$ assuming signal-independent and signal-uncorrelated noise is:

$$S_{exp}N_{exp}R = \frac{E[Trace\ (\mathbf{C}\mathbf{M}_H^T\mathbf{f}\mathbf{f}^T\mathbf{M}_\mathbf{H}\mathbf{C}^T)]}{E[Trace\ (\mathbf{C}\mathbf{n}\mathbf{n}^T\mathbf{C}^T)]}$$

where E represents the expectation operator. The above expression may be equated to:

$$\frac{Trace\ (\mathbf{C}\mathbf{M}_H^T\mathbf{R}\mathbf{M}_\mathbf{H}\mathbf{C}^T)}{Trace\ (\mathbf{C}\mathbf{R}_n\mathbf{C}^T)}$$

From the theory of matrix inequalities,

$$\lambda_{min} \leq S_{exp}N_{exp}R \leq \lambda_{max} \tag{3}$$

where λ_{min} and λ_{max} are minimum and maximum eigenvalues respectively of $(\mathbf{M}_H^T\mathbf{R}\mathbf{M}_H)\mathbf{R}_n^{-1}$.

With no more knowledge about the nature of the individual matrices, an optimal solution is one where $\lambda_{min} = \lambda_{max}$ or, equivalently,

$$(\mathbf{M}_H^T\mathbf{R}\mathbf{M}_H)\mathbf{R}_n^{-1} = c\mathbf{I}$$

for a constant c. This implies that:

$$\mathbf{M}_H^T\mathbf{R}\mathbf{M}_H = c\mathbf{R}_n$$

If $< \mathbf{x}, \mathbf{y} >' = \mathbf{x}^T\mathbf{R}\mathbf{y}$ is defined as a pseudo-inner-product, an optimal solution is one where

$$< \mathbf{M}_{Hi}, \mathbf{M}_{Hj} >' = cE(\mathbf{n}_i\mathbf{n}_j)$$

If the noise variables are orthogonal to one another, the optimal effective recording filters are R-orthogonal (i.e. orthogonal with respect to the inner-product $<,>'$. In addition if the noise is isotropically distributed, or the noise variables are orthonormal, optimal effective recording filters have equal norms as induced by the inner product $<,>'$ and are R-ortho*normal.*

In the rest of this paper, we assume that the noise variables are orthogonal, i.e., $\mathbf{R}_n = \sigma^2\mathbf{I}$. This assumption is reasonable because measurement noise is usually zero-mean and uncorrelated across channels. Our results may be easily extended to the case of correlated, non-zero-mean noise variables.

The value of $S_{exp}N_{exp}R$ before color correction is

$$\frac{Trace\ \mathbf{M}_H^T\mathbf{R}\mathbf{M}_H}{Trace\ \mathbf{R}_n} = \frac{Trace\ \mathbf{M}_H^T\mathbf{R}\mathbf{M}_H}{3 \times \sigma^2}$$

$$= \frac{1}{3} Trace\ (\mathbf{M}_H^T \mathbf{R} \mathbf{M}_H) \mathbf{R}_n^{-1} = \frac{1}{3} \sum \lambda_i$$

This implies that the ratio of the value of $S_{exp}N_{exp}R$ before and after color correction is:

$$\frac{S_{exp}N_{exp}R}{\sum \frac{\lambda_i}{3}} \tag{4}$$

Equations (3) and (4) can be used to obtain bounds on the factor by which the $S_{exp}N_{exp}R$ increases after color correction.

$$\begin{aligned} \frac{\lambda_{min}}{\lambda_{max}} &\leq \frac{3 \times \lambda_{min}}{\sum \lambda_i} \\ &\leq \frac{3 \times S_{exp}N_{exp}R}{\sum \lambda_i} \\ &\leq \frac{3 \times \lambda_{max}}{\sum \lambda_i} \leq \frac{\lambda_{max}}{\lambda_{min}} \end{aligned}$$

Since the noise variables are orthogonal,

$$\frac{\lambda_{max}}{\lambda_{min}} = \frac{\frac{\zeta_{max}}{\sigma^2}}{\frac{\zeta_{min}}{\sigma^2}} = \frac{\zeta_{max}}{\zeta_{min}}$$

where ζ_i is an eigenvalue of $\mathbf{M}_H^T \mathbf{R} \mathbf{M}_H$. The condition number (ratio of maximum to minimum eigenvalue) of $\mathbf{M}_H^T \mathbf{R} \mathbf{M}_H$ and its inverse provide upper and lower bounds on the ratio of the values of $S_{exp}N_{exp}R$ before and after color correction. We refer to the condition number of $\mathbf{M}_H^T \mathbf{R} \mathbf{M}_H$ as κ in the rest of this paper, and propose its use as a measure of filter orthogonality and noise sensitivity. A perfect value of κ is unity, and indicates an orthogonal set of color filters. Larger values of κ indicate 'less orthogonal' filters and larger noise amplification in general.

The condition number of the color correction matrix, equation(2), is distinct from the value κ. The condition number of the color correction matrix is a measure of the amount of noise amplification, for a specific viewing illuminant $\mathbf{L}$ and a specific estimated value of σ. An increase in the value of σ decreases the condition number of the color correction matrix, and hence also the noise amplification, but at the cost of color saturation - we discuss this in detail in the next section. Hence the condition number of the color correction matrix is not an accurate predictor of the image quality from a specific filter set. Calculating the condition number of the color correction matrix for $\sigma = 0$ may be a means of estimating the noise amplification and color saturation trade-off, but it may only be done when the filters are to be evaluated for a fixed viewing illuminant, which is not always the case (digital camera and scanner output images may need to be rendered for many different viewing conditions).

5 Color saturation and noise

A saturated color is defined as one whose tristimulus values are similar to those of a radiant spectrum with power concentrated in one wavelength. Thus, clearly, the tristimulus values of a saturated color match the values of the color matching functions at a particular wavelength, and are hence such that a single tristimulus value is dominant except at either end of the visual wavelength range. Colors of low saturation are represented by tristimulus values that are closer to being equal to another.

The combination of the color correction matrix and the sensors is the effective sensor set. The closer the spectral transmissivities of this combination are to the color matching functions, the more saturated the output colors. While the LMMSE estimate of the *tristimulus values* provides the expression for the color correction matrix, equations (1, 2), $\mathbf{R} = \mathbf{I}$ and $\mathbf{R}_n = \mathbf{0}$ provide the *best estimate of* $\mathbf{A}_L$ *as a linear combination of* $\mathbf{M}_H$. Thus $\mathbf{R} = \mathbf{I}$ and $\mathbf{R}_n = 0$ provide the most saturated colors, though not necessarily the best tristimulus estimates. Increasing the value of σ better conditions the matrix $(\mathbf{M}_H^T \mathbf{R} \mathbf{M}_H + \sigma^2 \mathbf{I})^{-1}$ and hence the color correction matrix of equation (2), but decreases color saturation. The effect of increasing the value of σ in the calculation of the color correction matrix is more pronounced when $\mathbf{M}_H^T \mathbf{R} \mathbf{M}_H$ is ill-conditioned.

Another way of seeing how increasing σ in the color correction matrix calculation decreases color saturation, especially for highly non-orthogonal filters, is as follows. A set of non-orthogonal filters with non-negative transmissivities is a set of overlapping filters, and hence a set of filters where the support of $\mathbf{M}_{Hi}$ is large. Because the support of $\mathbf{M}_{Hi}$ is large, and the transmissivities non-negative, the values of $\mathbf{M}_H^T \mathbf{f}$ are more nearly equal to one another than are the tristimulus values of $\mathbf{f}$. Increasing the value for σ in the matrix calculation tends to take the matrix $(\mathbf{M}_H^T \mathbf{R} \mathbf{M}_H + \sigma^2 \mathbf{I})$ closer to a multiple of the identity and hence also the matrix $(\mathbf{M}_H^T \mathbf{R} \mathbf{M}_H + \sigma^2 \mathbf{I})^{-1}$. Hence, the color-corrected measurements using a larger value of σ are $\mathbf{A}_L^T \mathbf{R} \mathbf{M}_H \mathbf{d}$ (equation 1) where the individual values of $\mathbf{d}$ are closer to one another than the tristimulus values of $\mathbf{f}$. Again, because the support of $\mathbf{M}_{Hi}$ is large, $\mathbf{M}_H \mathbf{d}$, which represents a non-zero linear combination of the effective scanning filters, is not a saturated color. Thus, an increase in the estimate of σ results in a decrease in color saturation, though it may reduce noise amplification.

6 Results and Discussion

We have studied many filter sets to understand the use of κ to characterize the trade-off between noise amplification and color saturation. We present here representative results with two specific filter sets used in digital cameras - a cyan, magenta, yellow, green (CMYG) filter set and a red, green, blue (RGB) filter set. The value of κ for the CMYG filters with respect to the Vrhel-Gershon-Iwan data set [4] is 34092, and with $\mathbf{R} = \mathbf{I}$ the value of κ is 138. The high values of κ for the CMYG filters predict:

1. Large noise amplification when $\sigma = 0$, represented by a large condition number for the color correction matrix for a specific viewing illuminant.
2. A large trade-off between noise amplification and color saturation as σ is increased, represented by a large decrease in condition number of the color correction matrix for a specific viewing illuminant.

For comparison, the corresponding values of κ for the RGB filters are 216 and 3.58. The low values of κ for the RGB filters predict that the noise amplification for $\sigma = 0$ will be low and that increasing σ will not significantly impact noise amplification or color saturation.

The following chromaticity plots (Figures 1-4) show the trade-off between color saturation and noise amplification for the CMYG camera. The raw data for the plots was taken with a CMYG mosaic arranged in a simple repeating pattern of period 2 pixels in each direction. The raw data was demosaiced with nearest neighbour interpolation in each color plane. The data was then color corrected for the CIE D65 viewing illuminant using the LMMSE estimate for $\sigma = 0$ (Fig. 1) and $\sigma = 10$ (Fig. 2), the latter corresponding to a signal to noise ratio of 25.5. The actual signal to noise ratio for the raw data is about 128, but a smaller value was used to illustrate the effect of increasing σ. The condition numbers of the *color correction matrices* (note - these are not values of κ, but indicate the extent of noise amplification for the particular viewing illuminant and the particular choice of σ) are 5.67 and 2.67 respectively. The high condition numbers and the decrease in condition number is consistent with the prediction. The condition numbers for the corresponding matrices for the RGB camera are 1.81 and 1.50, also consistent with predictions.

The horseshoe-shaped boundary of the plots of Figures 1-4 represents the entire chromaticity space possible with physically realizable spectra. Points on its boundary represent the most saturated possible colors. The vertices of the triangle defining the boundary of the points plotted in Figure 1 are the chromaticity coordinates of the phosphors for which the image was rendered ($\mathbf{A}$). Hence, the entire triangle defines the space of chromaticities possible with the specified monitor. The triangle is pretty well-covered by the points in this image. Hence, the image represented by the points in Figure 1 is highly saturated and contains a large proportion of the colors within the gamut of the monitor.

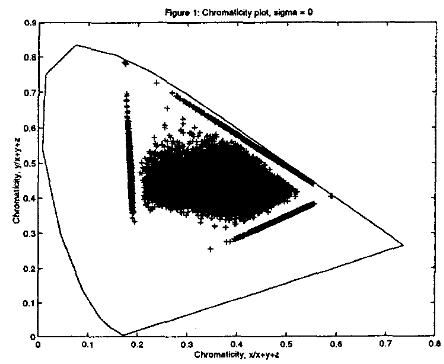

Figure 1: Chromaticity plot, $\sigma = 0$

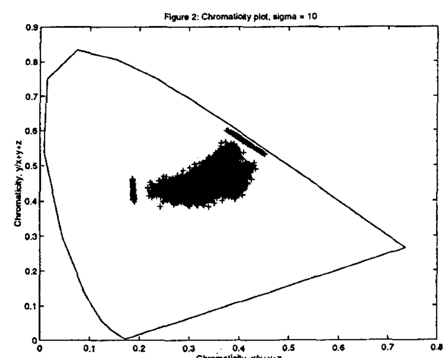

Figure 2: Chromaticity plot, $\sigma = 10$

The points of Figure 2 are concentrated towards the centre of the horseshoe, which represents white and grey colors and more desaturated colors (colors with tristimulus values that are close to equal). Hence, Figure 2 represents a highly desaturated image. Further, the triangle representing the monitor chromaticity space is not covered by the points in this image.

Figures 3 and 4 show the chromaticity distribution of a small patch of one color in the images represented by Figures 1 and 2 respectively.

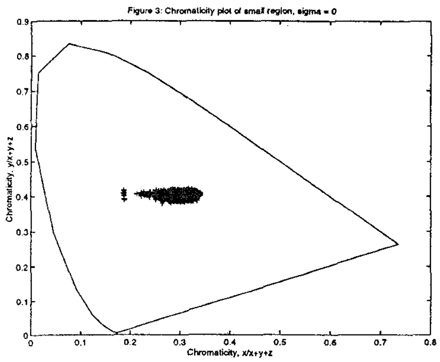

Figure 3: Chromaticity plot, $\sigma = 0$, small patch

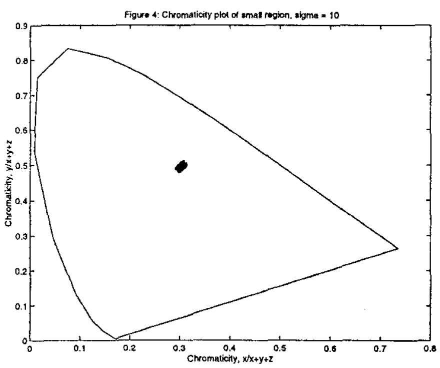

Figure 4: Chromaticity plot, $\sigma = 10$, small patch

The average values of the distribution of the NTSC RGB values are (123, 117, 184) and (51, 72, 49) respectively. The estimated standard deviations of these values are (20 6.5 6.1) and (2.3 2.9 1.5) respectively. The average values of the chromaticity distribution are (0.3 0.4 0.29) and (0.3 0.5 0.2) and the estimated standard deviations are (0.0148 0.0054 0.0157) and (0.0019 0.0026 0.0041) respectively. Both the numerical values and the figures show the existence of larger chromatic noise when σ is low.

Visual examination of color prints also shows that saturation in the upper picture is far superior to that in the lower picture, and that the noise in the lower picture is less visible - as predicted.

7 Conclusions

We have shown that the noise amplification in a captured image is related to the degree of orthogonality of the filters when measurement noise variables are orthogonal to one another. We have suggested the use of the condition number, κ, of the inner product matrix of the effective recording filters with themselves as a criterion for filter orthogonality. In cases where the set of color filters in an image acquisition device is suboptimal, we have shown that noise suppression by increasing the noise estimate in the calculation of the color correction matrix results in loss of color saturation, and that the trade-off between noise amplification and color saturation can be measured by κ. We have illustrated the trade-off as well as the use of κ with digital camera images.

References

[1] Sharma, G. and Trussell, H. J., *Figures of merit for color scanners*, IEEE Trans. Image Proc., vol. 6, no. 7, Jul. 1997, pp. 990-1001.

[2] Sharma, G., Trussell, H. and Vrhel, M. *Optimal non-negative color scanning filters*, in review, IEEE Trans. Image Proc.

[3] Vora, P. L., and Trussell, H. J., *Measure of Goodness of a Set of Color Scanning Filters*, Journal of the Optical Society of America - A, Vol. 10, No. 7, July 1993.

[4] M. J. Vrhel, R. Gershon, and L. Iwan, *The Measurement and Analysis of Object Reflectance Spectra*, Color Research and Application, vol. 19, no. 1, pp. 4-9, Feb 1994

[5] Vrhel, M. J., and Trussell, H. J., *Filter Considerations in Color Correction*, IEEE Transactions on Image Processing, Vol. 3, No. 2, pp. 147-161, Mar. 1994.

[6] Vrhel, M. J., and Trussell, H. J., *Optimal Color Filters in the Presence of Noise*, IEEE Transactions on Image Processing, Vol. 4, No. 6, pp. 814-823, June 1995.

[7] Wolski, M., Allebach, J. P., Bouman, C. A. and Walowit, E., *Optimization of sensor response functions for colorimetry of reflective and emissive objects*, IEEE Trans. Image Proc., vol5, no. 3, pp. 507-517, Mar. 1996.