Adaptive Lens-Dependent Hue Shift Correction

Bruce H. Pillman, Russell J. Palum; Eastman Kodak Company; Rochester, New York/USA

Abstract

As cameras become more compact, lens ray angles become larger, and imagers become more sensitive to ray angle. One problem this creates is lens-dependent spatial hue shift. This presentation discusses the causes of lens-dependent hue shift and describes an effective approach to correct it.

Background: Forms of Hue Shift

This paper focuses on lens-dependent hue shifts, which occur with changes in lens parameters. These show a very low frequency variation in hue as a function of location in an image. The case considered here occurs with an area sensor having a mosaic color filter array (CFA), as do most digital cameras. The most obvious hue shifts occur in scenes that have a neutral background. The neutral background gradually changes from the center of the image toward the edges of the image. With the specific cameras examined in this paper, the left and right sides of a landscape orientation image become bluish.

There are many causes of hue shift, depending on the optical system, sensor, electronics, and their interactions. Primary causes are variations in laydown of colorant in the color filter array, and optical and electrical crosstalk between pixels. Mechanisms for electrical crosstalk between pixels include charge transfer inefficiency (in a CCD sensor), electron diffusion, and temporal effects in the video signal and signal coupling effects between multiple signal channels. The primary cause for lens-dependent hue shift appears to be optical crosstalk, described later in this paper.

Calibration for Hue Shift Correction

In order to focus on lens-dependent hue shift, we first calibrate the imager for a specific set of conditions and see how response varies with lens changes. Calibration accounts for other sources of hue shift, allowing the analysis to focus on lens-dependent effects. The general calibration process is to image a uniform light neutral field and calculate gains to give a correct image. These gains are applicable for effects that are consistent during operation, provided usage conditions are similar to calibration conditions.

One of the best sources for calibration is an integrating sphere. For sphere-based calibration, a few design parameters must be chosen to produce illumination similar to a typical lens for the camera under consideration. The size of the aperture and its distance from the imager are chosen to duplicate the exit pupil size and location for a typical lens. Provided the fixture design keeps the aperture and distance constant, this calibration is extremely consistent. Because there is no lens involved, there is no focus variation, and appropriate selection of the exit port on the sphere prevents vignetting. The illuminant used for this calibration should simulate a real world neutral, such as D55.

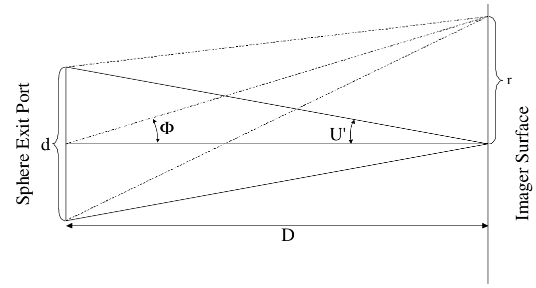

***Figure 1**. Illumination from Lambertian disk*

Given the illumination geometry shown in Figure 1, the illumination on the imager is described by equations (1), (2), and (3) from Reference [1].

$$\Phi = \tan^{-1}\left(\frac{r}{D}\right) \tag{1}$$

$$U' = \tan^{-1}\left(\frac{d}{2D}\right) \tag{2}$$

$$E(L,\Phi,U') = \frac{\pi L}{2}\left[1 - \frac{1+\tan(\Phi)^2 - \tan(U')^2}{\sqrt{\tan(\Phi)^4 + 2\tan(\Phi)^2\left(1-\tan(U')^2\right) + \sec(U')^4}}\right] \tag{3}$$

In these equations, Φ is the angle of the ray from the center of the exit port of the sphere to a point on the imager, r is the radial distance on the imager away from the optical axis, D is the distance from the Lambertian disk to the imager, U' is the half angle the Lambertian disk subtends on the imager at the optical axis, d is the diameter of the Lambertian disk, L is the luminance of the Lambertian disk, and E is the illuminance falling on a specific location on the imager. The controlled source and these equations provide an accurate prediction of the illumination on each pixel of an imager, allowing precise calibration for imager nonuniformity.

Imagers with a Bayer pattern CFA always have some difference in response between the two green channels (greens on red rows and greens on blue rows) as well as the differences between red, green, and blue channels. The fundamental issue is that greens of one channel have red pixels above and below and blue pixels to the left and right of each green pixel, while the other green channel has the opposite neighborhood. Because crosstalk between pixels is usually different in the horizontal and vertical directions, the change in neighborhood changes the crosstalk.

The split in green channels is driven by crosstalk from other color channels (from red and blue into green). Crosstalk also occurs from green into red and blue, but because each red (or blue) pixel has an identical neighborhood, the crosstalk does not induce particular patterns. Performing a gain calibration to match the two green channels is possible if the green channel differences are constant over a range of illuminants, brightness, and scene colors, and if calibration conditions simulate actual use conditions. Calibration can make things worse if use conditions are very different from calibration conditions or if green channel split varies significantly over use conditions.

Lens-Dependent Hue Shift

After calibration for a nominal condition, hue shift that changes with lens variations is the lens-dependent component of hue shift. The first step toward a robust solution is to understand why hue shift depends on lens parameters. The primary cause of lens-dependent optical crosstalk is the distance from the CFA to the photo sites on the imager. The distance can be comparable to the pixel pitch. As shown in Figure 2, as the distance from the optical axis to a pixel increases, the cone of light rays striking each pixel begins to strike more of the neighboring pixel. Signal sensed in the neighboring pixel is interpreted as having a different color, so this causes shifts in the apparent channel response. Figure 2 is not to scale in order to make the geometric effects clearly visible.

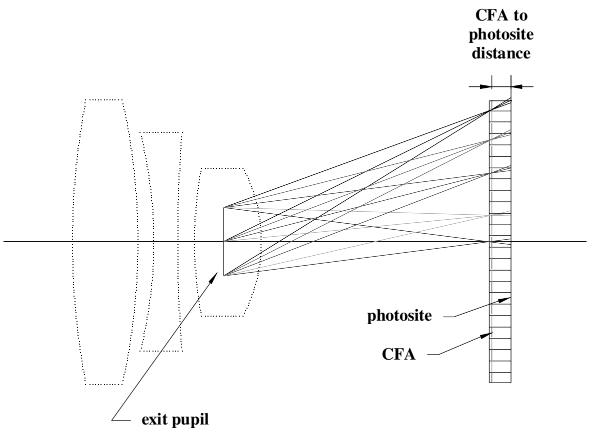

***Figure 2**. Illumination of imager with lens*

The exit pupil diameter and distance from the imager control the cone of light rays striking each pixel on the imager. With most lenses, the distance from the exit pupil to the imager changes with focus changes. Most zoom lenses also change distance with changes in focal length. The f/number of the lens affects the diameter of the exit pupil. Generally, the distance from an exit pupil to an imager has a larger impact on uniformity than does the diameter of the exit pupil, however, both can be significant.

The illuminant and colors in the scene also change crosstalk effects. They change the relative signal level in the different color channels, and hence the signal that leaks into adjacent pixels. The pixel design of the sensor is also significant. Most pixels have asymmetry in their layout, meaning that hue shift will not be a simple function of distance from the optical axis, but will be different horizontally and vertically, and possibly vary between left and right or top and bottom. In general, the problems caused by large ray angles get worse with smaller pixels and smaller fill factors. Lens-dependent hue shift produces response changes of up to 30%.

Correction of Lens-Dependent Hue Shift

Optical crosstalk involves signals leaking between pixels, so convolution-based correction is logical. However, crosstalk is typically less than 10%, so the center coefficient in the kernel is far more significant that any of the others. Correcting the center pixel with a gain correction provides most of the necessary correction. Because a convolution approach requires more detailed modeling and roughly ten times as much processing, we chose to use a gain-based correction.

Our adaptive hue correction, shown in Figure 3, is based on the observation that gain correction surfaces for many lenses have a similar shape, although they are different in magnitude. Lens-dependent hue shift is a very low frequency phenomenon, allowing use of very low-resolution images for storing and estimating the correction. The first step in correction is paxelization (block averaging and subsampling) to create a small 4-channel image from the full resolution CFA image. In this paper, analysis and hue correction images are based on paxelized images of 48 paxels × 72 paxels. We store nominal gain correction surfaces for red and blue, which are scaled to control the correction for red and blue. In addition, the GoB channel is adjusted to match the GoR channel.

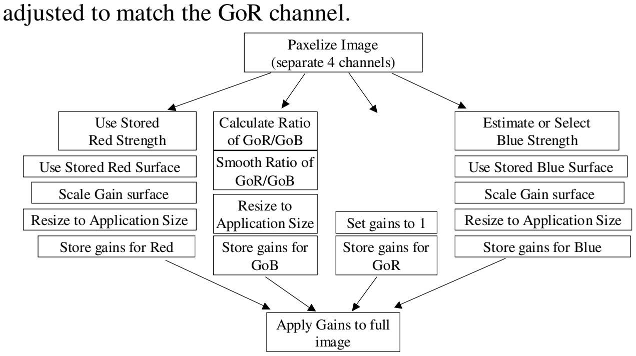

***Figure 3**. Adaptive hue correction flow*

This approach has the advantage of allowing the overall shape of red, green, and blue corrections to be independent. It also allows good precision for correction, as the gain surfaces can be scaled essentially continuously. It also simplifies adjustment of the gain surface to one degree of freedom for red and one degree of freedom for blue. Finally, it allows very good correction of the GoB channel with a very simple objective function, minimizing GNU.

Scaling of the red and blue gain surfaces is shown in Equation (4), where S_1 is strength, I is 4096 (1.0), G_i is the initial surface, G_s is the scaled surface, and x and y are indices for row and column.

$$G_s(x, y) = [G_i(x, y) - I] \times S_1 + I \tag{4}$$

Example initial red and blue gain surfaces are shown in Figures 5 and 6. The initial red gains range from 1.0 to 1.03, a fairly modest correction. The blue gains range from 1.0 to 1.1. The gains are stored as integers, with 4096 representing a gain of

1.0. Because the red gain correction for this camera was small, we decided to leave it fixed, simplifying the user interface. The range of the blue correction is large enough to justify adjustment for different conditions.

The gain surface correction described here corrects hue shift to within several percent. Examples of hue shift profiles are shown in Figures 6 and 7. The profiles show the GoR/blue ratio for the central rows of an image of a daylight neutral flat field. The UnCorr profiles show a substantial drop in GoR/blue toward the edges. The Corr profiles show the response after correction.

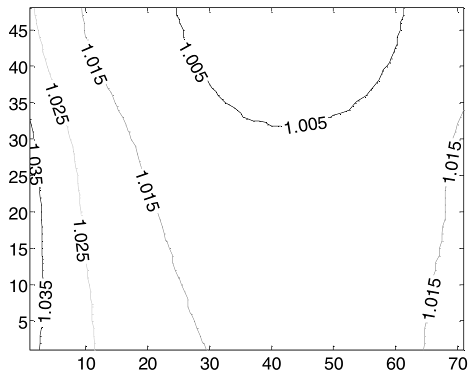

***Figure 4**. Initial red gain surface*

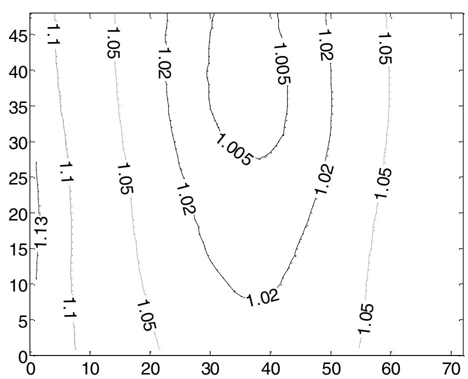

***Figure 5**. Initial blue gain surface*

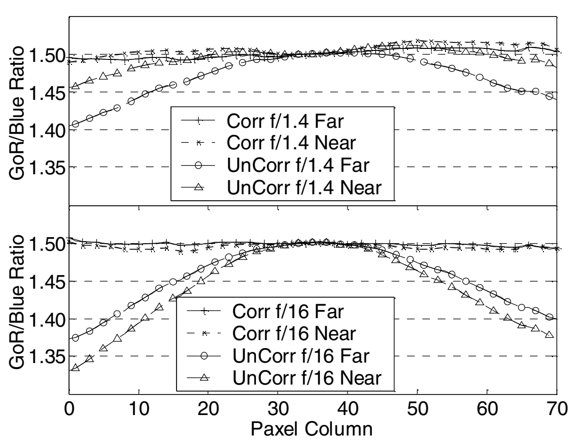

***Figure 6**. GoR/blue profiles, 85 mm Macro lens*

Figure 6 shows profiles for an 85 mm Macro lens at different f/numbers and focus distances. The magnitude of correction changes with f/number and with focus distance. Figure 7 shows the profiles for a 28 mm lens, for different focus distances and f/numbers. Focus and f/number changes have a smaller impact with this lens.

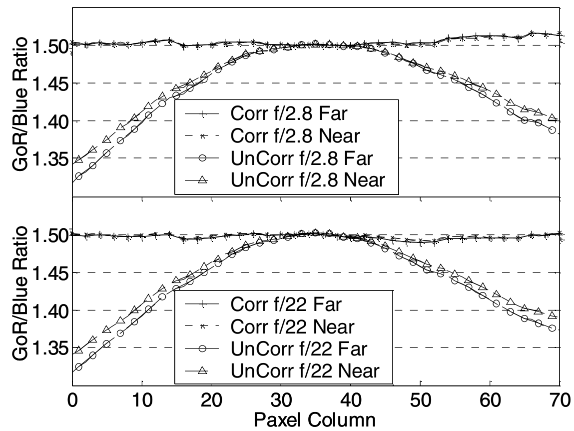

***Figure 7**. GoR/blue profiles, 28 mm lens*

Because the blue correction strength is a single parameter, it is fairly amenable to a simple slider-based user adjustment. This approach works well for users who have time and skill, but is labor-intensive. It is possible to use data reported from the lens (such as aperture, focal length, and focus distance) to estimate correction strength. Unfortunately, this breaks down with interchangeable lenses. Lenses of different designs have different exit pupil locations, even if they have the same focal length. This means adjustment based only on lens data is not completely reliable. For best reliability, estimation of the blue strength correction should be based on the image data itself.

Automatic strength estimation must estimate blue correction strength to reasonable tolerances. It is not necessary to be perfectly accurate, because manual control is also available. The enabling insight for correction strength estimation is that lens-dependent hue shift is primarily a horizontal gradient in the blue/green ratio, as shown in Figure 8.

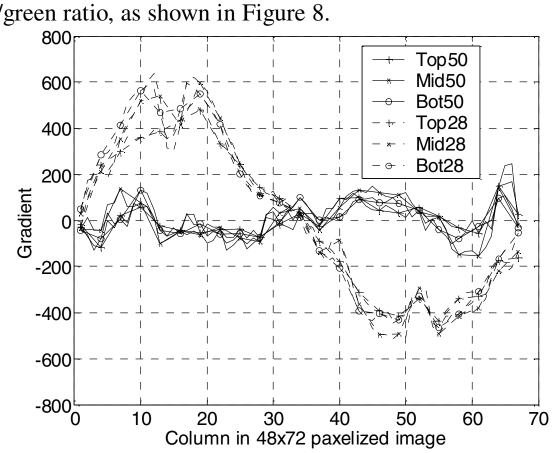

***Figure 8**. Blue/green gradient, nominal*

The ratio image is scaled by 4096 and the gradient values are scaled by 64, to make sure numerical precision does not affect the gradient estimates. The gradient is somewhat noisy but nearly flat for the case requiring little or no hue shift correction (solid lines). The gradient is not constant across the image even when much correction is required (dashed lines). The gradient is mostly left to right, not top to bottom. The change in hue is most visible when color is less saturated and in uniform areas that have little edge content. Thus, the estimation algorithm seeks a global estimate for this gradient, while it rejects busy and colored areas. The flow of analysis is shown in Figure 9.

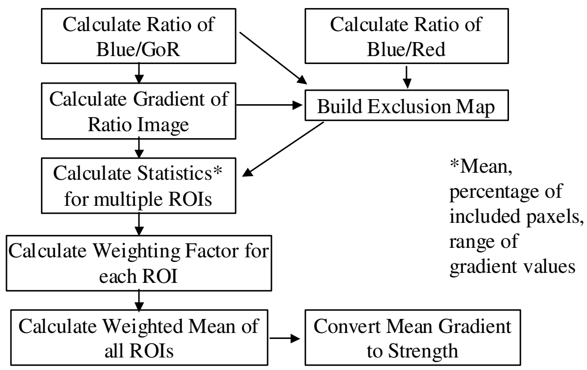

***Figure 9**. Estimation of correction strength*

The first step in this process is to calculate the ratio of the paxelized blue and GoR channels. We use a simple 1D filter to calculate the horizontal gradient in the ratio image. The third step is to build an exclusion map. This map is based on the blue/GoR ratio, gradient, and blue/red ratio. Paxels where either ratio is far from nominal or where the gradient is far from zero are rejected as being too busy or too highly colored. We compute statistics for the valid paxels in the gradient image for multiple regions of interest (ROI). To save processing, we locate the ROIs only where the gradient may be significant. The ROI locations are illustrated in Figure 10, with a plot of gradient normalization factors superimposed over the ROI locations.

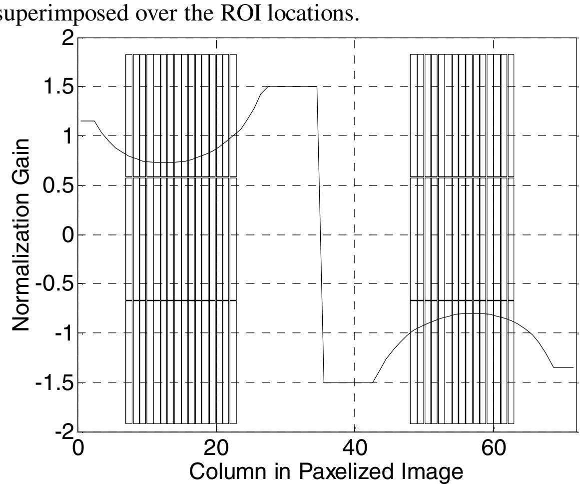

***Figure 10**. ROIs for gradient estimation*

The global gradient estimate is a weighted average of the mean gradient from each ROI. Depending on the scene content (and thus the statistics in each ROI), the global gradient estimate effectively comes from different ROIs. In order to produce a consistent gradient value, even when using different regions of interest, we normalize gradient values in different ROIs to have similar magnitude. This is particularly significant for correcting the change of sign between the left and right halves of the imager. To develop the gradient normalization gains shown in Figure 10, gradients from many images of flat fields under different conditions were averaged and fit with a smooth curve. Gains were calculated to bring the gradient at each ROI position to a nominal level.

The weighting factors for each ROI are calculated as shown in Equations (5)–(7). In these equations, r is the range of valid gradient values in an ROI, W_r is the weighting factor based on the range, p is the percentage of paxels included in an ROI, W_p is the weighting factor based on the range, and W is the final weighting factor for an ROI.

$$W_r = \begin{cases} 100; & r \le 4000 \\ 100\frac{r-4000}{1000-4000}; & 4000 < r < 10000 \\ 0; & r \ge 10000 \end{cases} \tag{5}$$

$$W_p = \begin{cases} 0; & p \le 0 \\ 100\frac{p}{80}; & 0 < p < 80 \\ 100; & p \ge 80 \end{cases} \tag{6}$$

$$W = \min(W_r, W_p) \tag{7}$$

Further Improvements

This algorithm corrects roughly 90% of a significant problem. Further improvements require more flexibility for control of the shape of correction, probably as a function of size of the exit pupil. Adjusting red channel correction strength would improve accuracy. The automatic strength estimation algorithm estimates correction strength within about 10%, which is sufficient for the current workflow. It is deliberately biased slightly toward undercorrection rather than overcorrection, and is designed to degrade gracefully toward no correction when the scene is too busy to make a good estimate.

References

[1] Rudolph Kingslake, Applied Optics and Optical Engineering, Volume II, ISBN 0-12-408602-0.

Author Biography

Bruce Pillman received a B.S. in chemical engineering from Northwestern University and an MS in electrical engineering from the University of Rochester. He has worked in research at Eastman Kodak Company since 1982. Since 1992, he has worked in development, simulation, and testing of consumer and professional-oriented digital imaging systems.