

Explanation of Color Lines Based on a Simple Color Image Model

Megan M. Fuller, Jae S. Lim; Massachusetts Institute of Technology, Department of Electrical Engineering and Computer Science, Research Laboratory of Electronics; Cambridge, MA

Abstract

In this paper, we show that a color image model we recently proposed explains the existence of color lines and predicts that they will have a slope of one. We present experimental results verifying this on several image datasets by showing that images segmented into blocks are often well-described by lines of slope one, that pixels with similar local averages fall on a line of slope one, and that, when all the pixels are normalized to have a local average of zero, they all fall on a line of slope one. We also discuss the image formation models that lead to this prediction and address some of the difficulties previously encountered in using image formation models to explain color lines.

Introduction

It has been observed [5] that when the values of the pixels of an image are plotted in RGB space, they tend to form discrete groups that are approximately linear, as shown in Fig. 1. This observation has been used in applications such as segmentation [5] and denoising [8]. However, to our knowledge, no satisfactory physical explanation has been proposed for this phenomenon.

In this paper, we will show that a model we recently derived from a physical image formation model [3, 4] predicts the existence of color lines. First, we will review the basic formulation of that model. We will then use that model to make certain predictions about the existence and properties of color lines, and we will present experimental results verifying these predictions. Finally, we will discuss the physical basis for our model in more detail and address some of the specific difficulties previously encountered in explaining color lines.

Our Model

In [3], we derived the following color image model from the Lambertian image formation model:

$$R(x) = R_{avg}(x) + q(x) \quad (1)$$

$$G(x) = G_{avg}(x) + q(x) \quad (2)$$

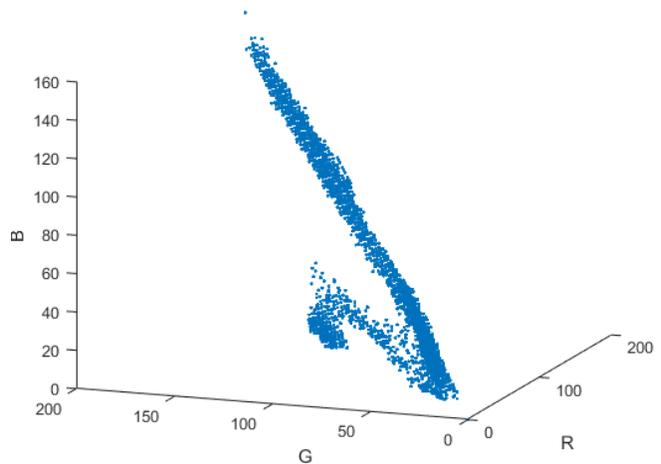
$$B(x) = B_{avg}(x) + q(x) \quad (3)$$

where $R(x)$, $G(x)$, and $B(x)$ are the values of the red, green, and blue channels at location x , $R_{avg}(x)$, $G_{avg}(x)$, and $B_{avg}(x)$ are the values of the local averages of the three color channels, and $q(x)$ is a color-independent residual.

What this model says is that the value of a pixel in a particular color channel can be written as the sum of the color-dependent local average and a color-independent residual. Here, the local average can be computed over any region that contains a single material. In practice, we start with a square centered at the pixel under consideration and use some basic edge-detection to remove



(a) Fragment of image 3 from the Kodak dataset.



(b) Pixel values plotted in RGB space.

Figure 1. Example of color lines.

pixels that are likely from other materials. The method used here is the same as that described in [3].

In [3] and [4], we presented experimental results showing that this model accurately describes a database of images collected with a Sigma DP1 camera, and a separate database collected with a Pentax K-3 II, as well as the standard Kodak dataset [1].

Predictions

If we consider only pixels within a single material, then the local averages of (1)-(3) are no longer functions of position. In this case, our model predicts that the values of the pixels will fall on a line in RGB space. Furthermore, this line will have a slope of one and will intercept the point $(R_{avg}, G_{avg}, B_{avg})$.

Unfortunately, this prediction is difficult to verify because doing so requires segmenting the images into different materials, which is itself a hard problem. Therefore, we will look at several related predictions that are easier to verify.

First, if we segment the images into square blocks of reasonable size, we expect some of these blocks to contain only one material. The pixels within these blocks should fall on a line, and the slope of that line should be one. Second, if we select only pixels that have similar local averages, regardless of their spatial locations, these pixels should also fall on a line of slope one. Finally, if we subtract the local averages from the images, effectively normalizing all pixels to have a local average of zero, then all the pixels of all the images in the dataset should fall on a line of slope one.

Experimental Results

In this section we will present results verifying the predictions listed in the previous section on the images of the Sigma, Pentax, and Kodak datasets.

First, we segmented the images into 16x16 pixel blocks. The intent was that these blocks be large enough that the data within them was meaningful, but small enough that they would generally contain only a single material. We then computed the line of best fit to the pixel values and discarded those blocks that were not well-described by a single line. We selected “well-described” blocks by computing the error of the data to the best-fit line and discarding those blocks where the error was too large. This was necessary in order to discard those blocks that contained two or more objects. We found that we retained 40% of the blocks in the Sigma dataset, 65% of the blocks in the Pentax dataset, and 57% of the blocks in the Kodak dataset. Since we intentionally selected only blocks that were well-described by lines, these results cannot be used to say that lines will form, although the fact that they did in about half the blocks is encouraging. However, we can look at the slopes of the lines that did form and verify that they are close to one, as expected. In particular, the average slopes are shown in Table 1.

Average slopes, images segmented into blocks

	Sigma	Pentax	Kodak
Average Slope R vs G	.9837	1.0252	.9252
R vs B	.8834	1.1017	.8213
G vs B	.8742	1.0442	.8516

Second, we segmented the images based on pixel values. In particular, we computed the edge-sensitive local average of the pixels in the three color channels and chose those pixels whose local averages were similar. We used the global average of the image as our reference point and took pixels whose local averages were within $\pm 1\%$ of the dynamic range of the image of the global average in each color channel to form the group. We found that most of the images in the datasets had sufficiently many pixels in this range for these groups to be meaningful, though we did have to discard a few images that did not have enough pixels, and one that had many small material edges that went undetected when computing the local average.

Some examples are shown in Fig. 2. These figures take all the pixels of the image whose local averages fall within the orange square and plot the value of one color channel as a function of another. The images used to generate these plots are shown in Fig. 3. Although the length and tightness of the lines varied from image to image within the datasets, these examples are representative.

The average slopes of the lines formed by the images in each dataset are shown in Table 2. These results also support the idea that color lines with a slope of one occur in images.

Average slopes, images segmented by local average

	Sigma	Pentax	Kodak
Average Slope R vs G	.8532	.9141	.9112
R vs B	.7716	.9340	.8459
G vs B	.8756	1.0244	.9329

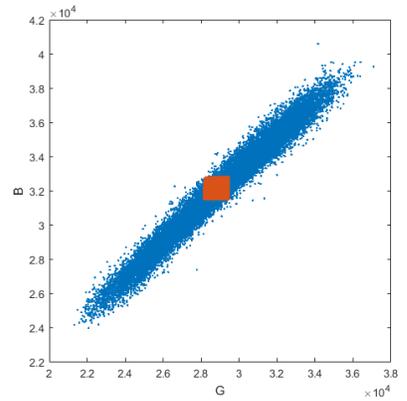
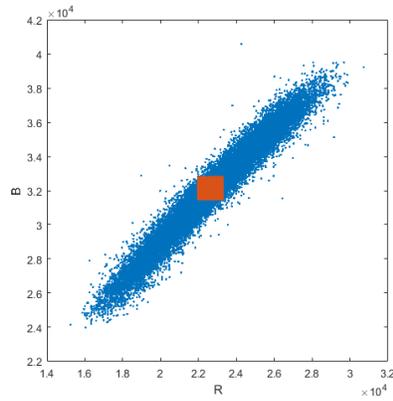
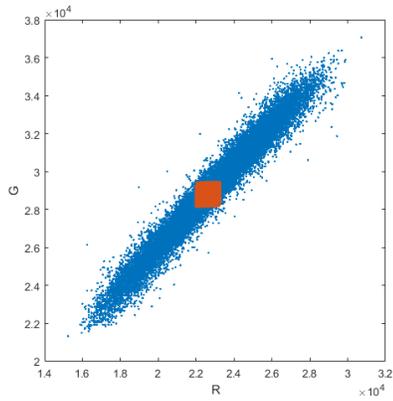
As a final test, we subtracted the local average from the image. This effectively normalizes the local average of all pixels to zero, and so we can look at the values of all the pixels in an image, and, indeed, in a dataset, on a single plot. Fig. 4 shows the 2D histograms of the values of all the normalized pixels in the three datasets. (The narrow horizontal and vertical lines visible in the figures are not meaningful—they are only an effect of how the plots were rendered.) Again, these results show that, once the local averages have been removed, the pixels values tend to fall on a line of slope one in RGB space.

These results verify the predictions of the previous section. Taken together, they provide good evidence that the values of the pixels of an image will fall on distinct lines of slope one, as expected.

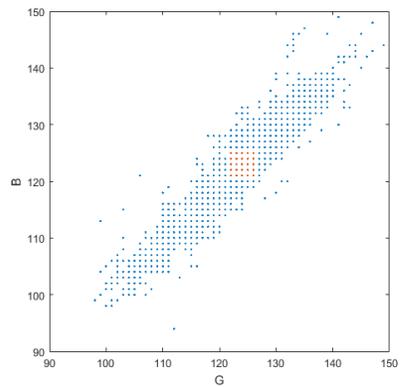
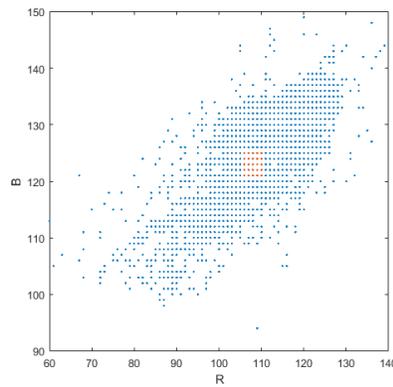
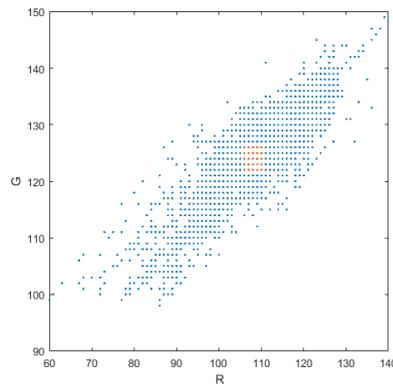
Discussion

We note here that, although we are particularly interested in the Lambertian image formation model, (1)-(3) could be derived from any image formation model that is a separable function of position and wavelength. This includes the simplified Oren-Nayar model [6] and the Torrance-Sparrow model [7], assuming there is only specular reflection and the Fresnel coefficient is approximately constant over angle in the region of interest [4].

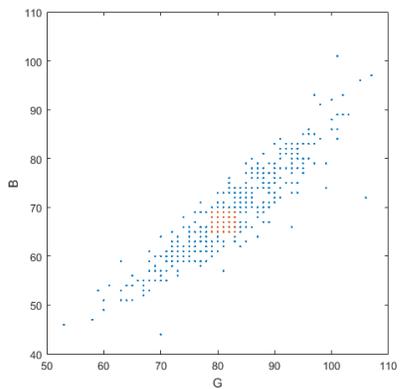
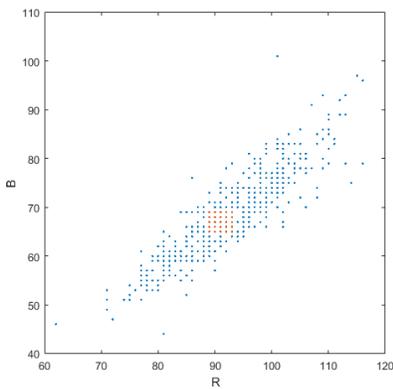
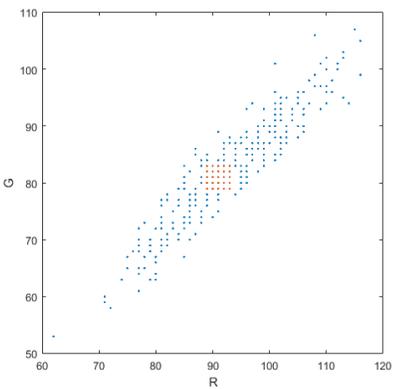
We also note that when Omer and Werman proposed their color-line model, they rejected the idea that it was explainable by the Lambertian model [5]. This rejection was based on two observations: First, the color “lines” were not actually straight lines, and, second, that they did not intercept the origin, as the Lambertian model would predict. Therefore, they required, on average, six points to define the clusters of data they referred to as



(a) Pixel values of image 28 from the Sigma dataset.



(b) Pixel values of image 12 from the Pentax dataset.



(c) Pixel values of image 5 from the Kodak dataset.

Figure 2. Examples showing the values of pixels with similar local averages.



(a) Sigma image 28.

(b) Pentax image 12.

(c) Kodak image 5.

Figure 3. Images used to generate the examples in Fig. 2.

“lines.”

The results presented here indicate that color lines are in fact straight lines in many instances. We believe the difference is due to the nature of the datasets—Omer and Werman present results for a dataset which they indicate had been compressed aggressively, while our datasets were uncompressed. In the case of the Sigma and Pentax datasets, we used the most raw images available from the camera. Therefore, it is reasonable to conclude that the non-linear nature of the groups observed by Omer and Werman is due primarily to camera post-processing and is not inherent in the original images.

The notion that the Lambertian model requires the color lines to pass through the origin is only correct if the camera is measuring values that are linearly related to the scene irradiances. However, using the camera calibration function of Debevec and Malik [2], we found that both the Sigma and Pentax cameras were reporting values related almost linearly to the logarithm of the scene irradiances, and we assume the camera used to collect the Kodak dataset was doing something similar. Once this transformation is accounted for, there is no longer any reason to expect the color lines to intercept the origin; rather, as we have shown, we should expect a slope of one.

Conclusion

In this paper, we have shown that the color image model we presented previously leads to color lines with a slope of one. We have presented experimental results from three different cameras verifying this. These results show that the color lines model of Omer and Werman is physically justified and that the complicated nature of the lines they found is due primarily to post-processing done in the camera, meaning a simpler model could potentially be used when raw data is available.

References

- [1] Kodak lossless true color image suite. <http://r0k.us/graphics/kodak/>. Accessed: 11-29-16.
- [2] P. E. Debevec and J. Malik. Recovering high dynamic range radiance maps from photographs. In ACM SIGGRAPH 2008 Classes, SIGGRAPH '08, pages 31:1–31:10, 2008.
- [3] M. Fuller and J. S. Lim. A color image model with applications to denoising. In Proc. IS&T Electronic Imaging Conf., 2017.
- [4] M. Fuller and J. S. Lim. Derivation and experimental verification of a simple model for color images. JIST, March/April 2018. accepted.
- [5] I. Omer and M. Werman. Color lines: image specific color representation. In CVPR, volume 2, pages 946–953, 2004.
- [6] M. Oren and S. K. Nayar. Generalization of lambert’s reflectance model. In Proc. of the 21st annual conference on computer graphics and interactive techniques, pages 239–246, 1994.
- [7] K. E. Torrance and E. M. Sparrow. Theory for off-specular reflection from roughened surfaces. Journal of the optical society of America, 57(9):1105–1114, 1967.
- [8] C. Tsai, W. Tu, and S. Chien. Efficient natural color image denoising based on guided filter. In IEEE ICIP, pages 43–47, Quebec City, QC, September 2015.

Author Biography

Megan M. Fuller received her B.S. in electrical engineering from Brigham Young University in 2012 and her M.S. in electrical engineering from the Massachusetts Institute of Technology in 2014. She is currently a graduate student at MIT working towards a Ph.D. in electrical engineering. Her work has focused on image modeling and restoration.

Jae S. Lim received the S.B., S.M., E.E., and Sc.D. degrees in EECS from the Massachusetts Institute of Technology in 1974, 1975, 1978, and 1978, respectively. He joined the M.I.T. faculty in 1978 and is currently a Professor in the EECS Department.

His research interests include digital signal processing and its applications to image, video and speech processing. He has contributed a number of articles to journals and conference proceedings. He is a holder of more than 40 patents in the areas of advanced television systems and signal compression. He is the author of a textbook, Two-Dimensional Signal and Image Processing. He is the recipient of many awards including the Senior Award from the IEEE ASSP Society and the Harold E. Edgerton Faculty Achievement Award from MIT. He is a member of the Academy of Digital Television Pioneers and a fellow of the IEEE.

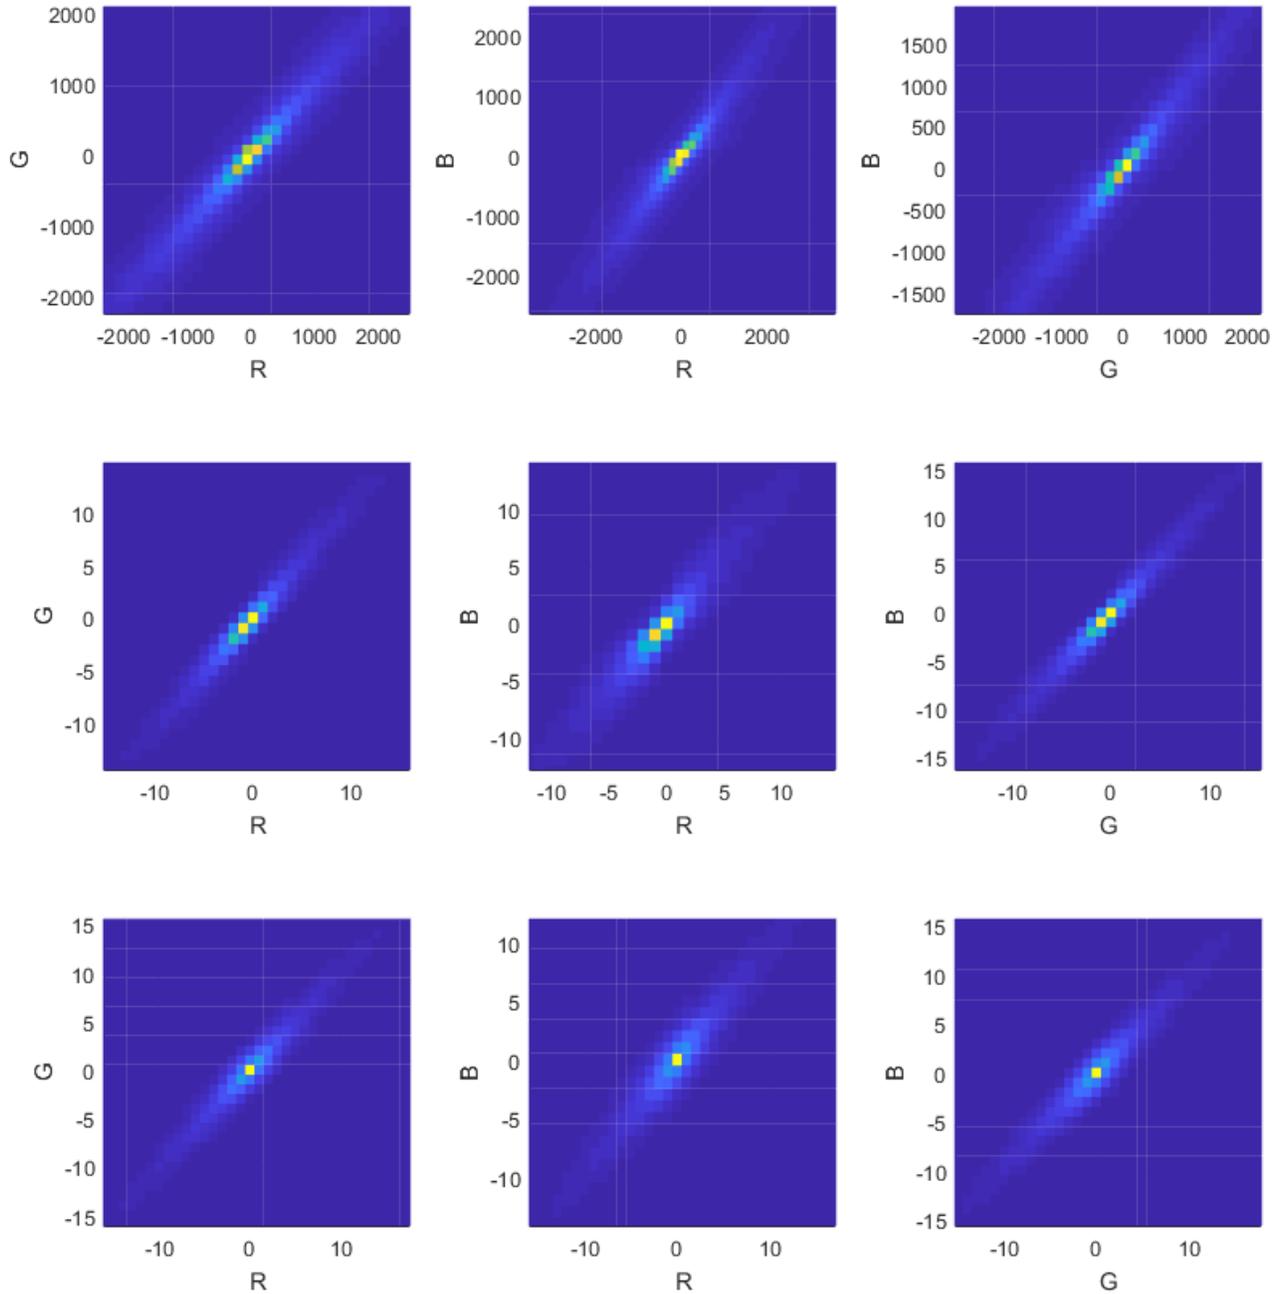


Figure 4. Plots of pixel values from the Sigma (top), Pentax (middle) and Kodak (bottom) datasets