PERCEPTUAL-BASED CORRECTION OF PHOTO RED-EYE

Robert Ulichney
HP Labs, Cambridge, MA, USA; u@hp.com
Matthew Gaubatz
Cornell University, Ithaca, NY, USA; mdg27@cornell.edu

ABSTRACT

Many recent works in digital photograph image enhancement have included methods for detecting camera-flash-induced red-eye artifacts. Less effort, however, has been spent on the correction procedure itself. A variety of corrective procedures have been presented that reduce the chrominance and luminance of detected pixels in different ways. Given that the original pupil color of a subject is often unrecoverable, a simple chrominance de-saturation effectively removes the red hue from the artifact pixels. This work presents the results of a perceptual study designed to find the most visually pleasing target luminance for corrected images. The results of this experiment are analyzed and correlated with the test data to yield simple equations for target luminance. These are in turn combined with a fully-automated procedure designed to minimize intrusive effects associated with pixel re-coloration. Data and equations are presented for hard-copy (printed) and soft-copy (on-screen) images.

KEY WORDS

Digital photography, image enhancement, red-eye

1. Introduction

The appearance of red-eye in flash photography is one of the most prevalent and disturbing artifacts that affects image capture. This problem has long bothered camera users, and is especially common with compact cameras, due to the inherently small angle between the lens and the flash. Figure 1 illustrates the red-eye beacon that shines from a subject's eye in the presence of a flash. This beacon is a cone of radius α shining back at the flash. Its apparent red color is caused by the reflection of the flash off the blood vessels of the subject's retina. The camera will record this red hue if the angle between the flash and camera, β, is not greater than α.

All algorithmic solutions for fixing red-eye can be segmented into two parts: detection and correction. By far the more difficult analytical and computational problem is the detection phase, that is, the process of correctly identifying the precise pixels in the input image that contain red-eye artifacts. Detection solutions are measured by percent correctly identified and false positive rate. Once the red-eye regions are identified, the second phase is correction where the offending pixels are to be changed.

A variety of solutions have been proposed for detection. Several methods [1,10] use a series of classifier-based modules that perform initial candidate selection and verification. Another [2] uses face detection as a starting point for looking for red-eyes, using a fast, robust face detector [3]. A wavelet-based face-detector approach [9] has also been presented. Instead of staring explicitly with faces, skin detection can be employed to narrow the search for red-eye artefacts [5,8]. FotoNation recently reported [4] a scheme with a focus on computational efficiency and low memory usage. These approaches have trade-offs in success rates, and detection solutions continue to evolve.

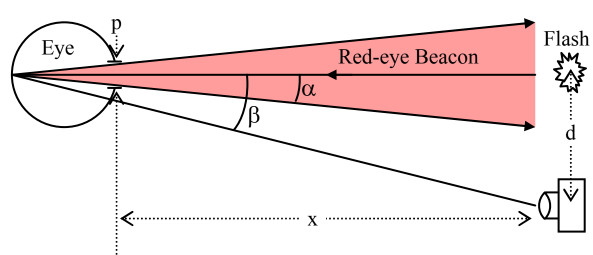

Figure 1. The geometry of photo-flash red-eye.

In almost all previous works, the problem of correction is treated as a separate problem from detection. An interesting solution to the re-coloration problem is given in [6], which leverages color information from a non-flash image. For many current cameras, however, this information is unavailable. The importance of glint preservation is noted in [5]. Several works offer detailed solutions for pixel re-coloration [2,7,8,10] but not from the perspective of perceptual tests.

This paper discusses methods to digitally remedy this problem, and focuses on a means to correct red-eye in a perceptually pleasing way. A perceptual experiment is presented, designed to predict target luminance values for corrected pixels. The results of this experiment are incorporated into an algorithm that operates on pixels containing the detected artifact. This algorithm, which assumes the locations of red-eye pixels are known, can be incorporated into any automated enhancement algorithm that performs detection and correction in two stages.

This paper is organized as follows. Section 2 outlines the perceptual experiment to determine target luminance in corrected pixels. The results are explained in Section 3, the correction procedure is outlined in Section 4. Section 5 concludes the paper.

2. Perceptual Testing

Since the chrominance of the red-eye region is completely incorrect, a natural solution is to de-saturate the effected pixels. This procedure leaves the region grey, and usually much lighter than what would appear natural. Thus, after de-saturating the color of the artifact, the average luminance should be lowered. But the question remains: to what target luminance should the average luminance be lowered? In order to investigate this question, controlled subjects were chosen to evaluate a series of sample images with red-eye correction. The variable in this experiment was the average luminance of the de-saturated pixels.

2.1 Establishing Equi-spaced Samples

Subjects were presented with five versions of each sample photo. In each version, the corrected red-eye pixels were set to one of five different luminance values. To gather the most useful information from this test, these five luminance values perceptually equi-spaced. The number that represents the pixel brightness value is referred to as “digital count”, Y. It is assumed to be an 8-bit quantity, and is therefore associated with values in the range [0,255]. Equally spaced target lightness values are expressed in CIE L* space, a reasonable representation of perceptually uniform luminance.

These target values must be translated to digital count in order to generate the test images. For on-screen or video display, luminance can be expressed in terms of digital count with the following set of equations:

$$I = Y_0 + (Y_{max} - Y_0)(Y/255)^{\gamma},$$

where Y_{max} = 1.0 and Y_0, a measure of the inability to achieve perfect black to flare, is set to a nominal value of 0.02. The gamma, γ, of the test display is set equal to 1.8. An explanation for the role of γ is given in [11]. The following standard relationship is employed to convert I to L*:

$$L^* = 116\, I^{1/3} - 16$$

Combining these two expressions results in a mapping from digital count to L*, shown in Figure 2. This relationship is for hard copy presentation, where ambient illumination replaces flare as the mitigator at the black end. This type mapping can determine five equally spaced target digital count values, which are in turn used to create test images.

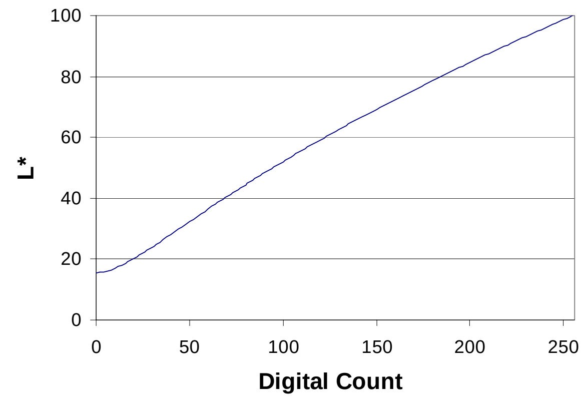

Figure 2. Non-linear mapping between L* and digital count.

Level	L*	Digital Count
0	15.487	0
1	22.732	28
2	29.976	44
3	37.220	62
4	44.465	80

Figure 3. The five test Levels, and associated values of L* and Digital Count.

The luminance values are equi-spaced between black at Y=0 and relatively light grey value of Y=80. Corresponding intervals in L* as shown in Figure 3. Each target luminance was assigned a corresponding value of “Level”, a new unit used in the perceptual test.

2.2 Rating Process

Twenty-five images of single frontal faces with photo flash red-eye were carefully collected to span a representative spectrum of face types and illumination conditions. Each image was scaled and cropped so that the twenty-five faces appeared the same size. Two sets of experiments were performed, one for on-screen viewing and one for hard copy viewing. The images were presented under set lighting conditions, with neutral gray surround.

For both viewing conditions, all 5 versions of the test photo were displayed at the same time. An example soft-copy display is shown in Figure 10. Starting at the upper left image, labelled “B”is Level = 0, followed by “C” for Level = 1, “D” for Level = 2, “E” for Level = 3, and “F” for Level = 4. Image “D” is presented both on the end of the first row and the beginning of the second row to preserve consistent side-by-side comparison. Hard copy images were each printed as 5x7 prints on grey mattes and positioned so that all 5 levels were viewed side-by-side. All 19 subject judges assessed 5 versions of 25 images for both on-screen and hard copy. Each subject was asked to select the version that appeared to have the most natural re-coloration. Subjects were allowed to specify that the most natural corrected luminance lay between two of the presented choices.

3. Interpreted Results

The results of the perceptual tests are summarized in Figure 4 for hard copy and Figure 5 for video screen. The experimental "Level" units can be converted back to digital count, Y, as follows:

$$Y = ((((7.244\ \text{Level} + 31.487)/116)^3 - 0.02)\ 255/0.98)^{1/1.8}$$

For both viewing conditions, the average desired Level roughly corresponds to a digital count of 30. This result can be used to select a pleasing target luminance, denoted f(Y), in detected red-eye pixels via:

$$f(Y) = 30$$

Based on the variety of the dynamic range in the test images, an adaptive approach is considered as well. Intuitively, the ideal target luminance for an extremely bright (almost washed-out) image would be perhaps higher than the target luminance for in a poorly lit scene.

The correlation between preferred luminance value for a particular image and the average luminance in a box surrounding each detected eye (where both luminance values were represented in digital count) was computed to analyze this relationship. This computation was repeated for a wide range of box sizes, normalized relative to the radius of the artifact. (The "radius" of the box, which is square, refers to half the width of one of the sides.) Results of this computation, averaged over all the images, are presented in Figure 6.

A surprisingly convenient result is revealed: for both viewing conditions, the correlation peaked at a radius ratio of around one. The resulting implication is that the desired target luminance is most strongly related to the average luminance of the red-eye artifact, instead of the luminance of some other region. This relationship is fortuitous because the output of a detection algorithm already specifies the pixels to be corrected and thus a correction module does not need any more data that that to derive a well-chosen target luminance.

Using this, the rule for finding the target luminance given the average luminance of the red-eye, Y_{av}, is

$$f(Y) = 0.167 \times Y_{av} + 11.523,$$

as plotted in Figure 7.

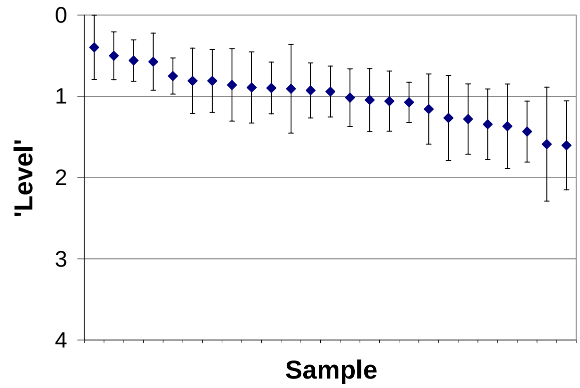

Figure 4. Average user luminance preferences for hard copy images.

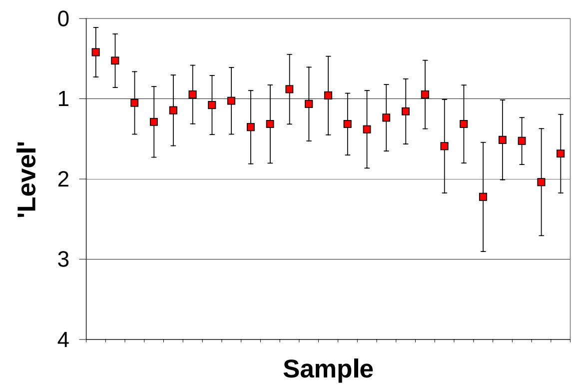

Figure 5. Average user luminance preferences for soft copy images.

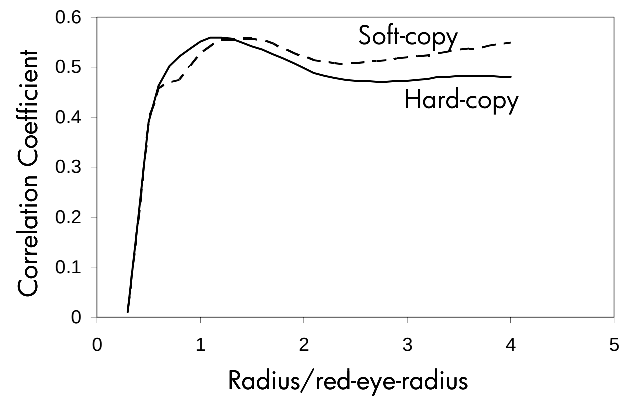

Figure 6. Patch radius vs. correlation coefficient. As the plot extends left to right, the mean luminance is computed over a larger and larger patch of pixels.

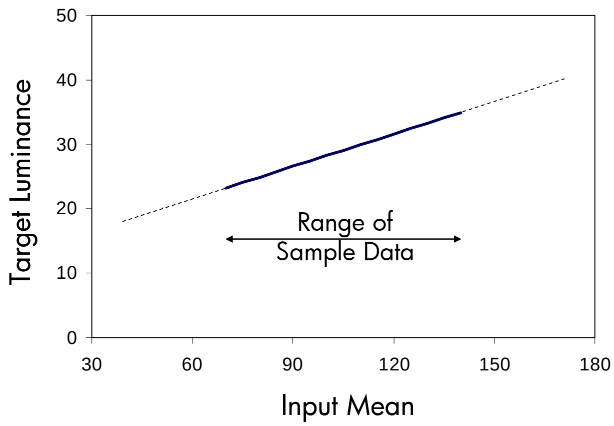

Figure 7. Model mapping input patch mean luminance Y_{av} to target mean luminance f(Y).

4. Correction System

The detection system must identify the regions of red-eye to be corrected. The job of the correction system is to (1) de-saturate the region, and (2) set the average luminance to the rule in Figure 7. It would appear unnatural to set the entire red-eye region to one luminance level, i.e. such that the eye region contains absolutely no variations in brightness. Instead, to preserve the subtle luminance structure in the eye, each red-eye luminance value is multiplied by the ratio of target luminance over original mean luminance.

It is also important to taper the correction for both de-saturation and luminance adjustment to avoid inducing an artificial hard edge in the eye. A taper that extends over 10% of the diameter of the eye region, D, was found to achieve good results. An example cross-section of such a tapered correction mask is shown in Figure 8, along with an example of an eye with and without tapering in Figure 9. The values of the tapered mask fall in the range [0,1].

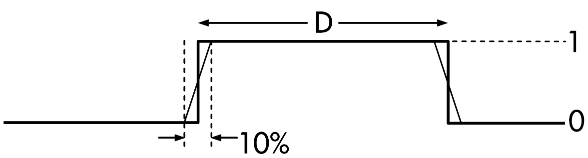

Figure 8. Example cross-section of a tapered correction mask.

Let p represent the value of the tapered mask a single location. This value represents the percentage by which the pixel chrominance or luminance will be reduced, according to the given scheme. Let Y, Cb and Cr be the original pixel values. Equations for the modified chrominance values for this pixel are given as the following:

$$Cb' = (1 - p) * Cb$$

$$Cr' = (1 - p) * Cr$$

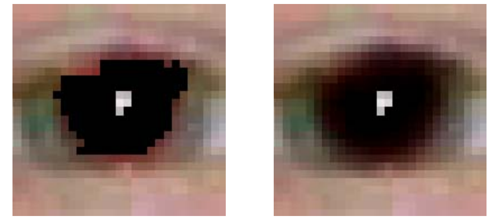

Figure 9. Example red-eye defect with non-tapered correction (left) and tapered correction (right).

Recall Y_{av} represents the mean pixel luminance (in digital count) for all pixels in the immediate vicinity of the detected artifact. The adjusted pixel luminance is then given by

$$Y' = (1 - p) *Y + p*f(Y)/Y_{av},$$

where $f(Y) = 0.167 \times Y_{av} + 11.523$. This equation can therefore be rewritten

$$Y' = (1 - p) *Y + p*(0.167 + 11.523/Y_{av}).$$

A schematic diagram of the correction procedure is given in Figure 11. In the proposed implementation, the results of the perceptual experiment were used to determine the desired luminance map f(Y).

5. Conclusion

Automatic correction of photo red-eye will continue to grow in importance as flash-to-lens distances will remain small for a large number of capture devices. It is important to recognize that the problem of detection and correction can be separated, and that research to improve the performance of both parts is an on-going process.

This paper presents a scheme for achieving natural looking corrections based on perceptual testing. The algorithm de-saturates the color of affected pixels, and reduces the average luminance to a target level, determined by the testing. Only the pixels immediately surrounding the detected artifact are required to determine the target luminance. The border of the region of modified pixels is softened with a simple taper.

6. Acknowledgements

The authors would like to thank Andrew Juenger for giving advice regarding mapping digital count to actual display luminance and Kristyn Falkenstern for executing the subject evaluations.

References:

[1] H. Luo, J. Yen and D. Tretter, An efficient redeye detection and correction algorithm, *Proc. IEEE Conf. on Pattern Recognition,* Cambridge, UK, 2004.

[2] M. Gaubatz and R. Ulichney, Automatic red-eye detection and correction, *Proc. IEEE Conf. on Image Processing*, Rochester, NY, 2002, 804-807,

[3] P. Viola and M. Jones, Robust real-time object detection, Tech. Rep. 2001/1, Compaq Cambridge Research Laboratory, Feb. 2001.

[4] P. Corcoran, et al., Automated in-camera detection of flash eye-efects, *IEEE Digest of Technical Papers, Int. Conf. on Consumer Electronics*, 2005, 367-368.

[5] J. Schildkraut and R. Gray, A fully automatic redeye detection and correction algorithm, *Proc. IEEE Conf. on Image Processing*, Rochester, NY, 2002, 801-803.

[6] X. Miao and T. Sim, Automatic red-eye detection and removal, *Proc. IEEE Conf. on Multimedia and Expo*, Taipei, Taiwan, 2004, 1195-1198.

[7] A. Patti, K. Konstantinides, D. Tretter and Q. Lin, Automatic digital red-eye reduction, *Proc. IEEE Conf. on Image Processing*, Chicago, IL , 1998.

[8] R. Schettini, F. Gasparini, F. Chazli, A modular procedure for automatic redeye correction in digital photographs, *Proc. SPIE Conf. On Color Imaging: Processing, Hardcopy, and Applications IX*, San Jose, CA, 2004.

[9] S. Ioffe, Red eye detection with machine learning, *Proc. IEEE Conf. on Image Processing*, Barcelona, Spain, 2003, 871-874.

[10] L. Zhang, Y. Sun, M. Li and H. Zhang, Automated red-eye detection and correction in digital photographs, *Proc. IEEE Conf. on Image Processing*, Singapore, 2004, 2363-2366.

[11] C. Poynton, The rehabilitation of gamma, *Proc. SPIE Human Vision and Electronic Imaging III*, San Jose, CA, 1998, 232-249.

Figure 10. Example sequence for on-screen test.

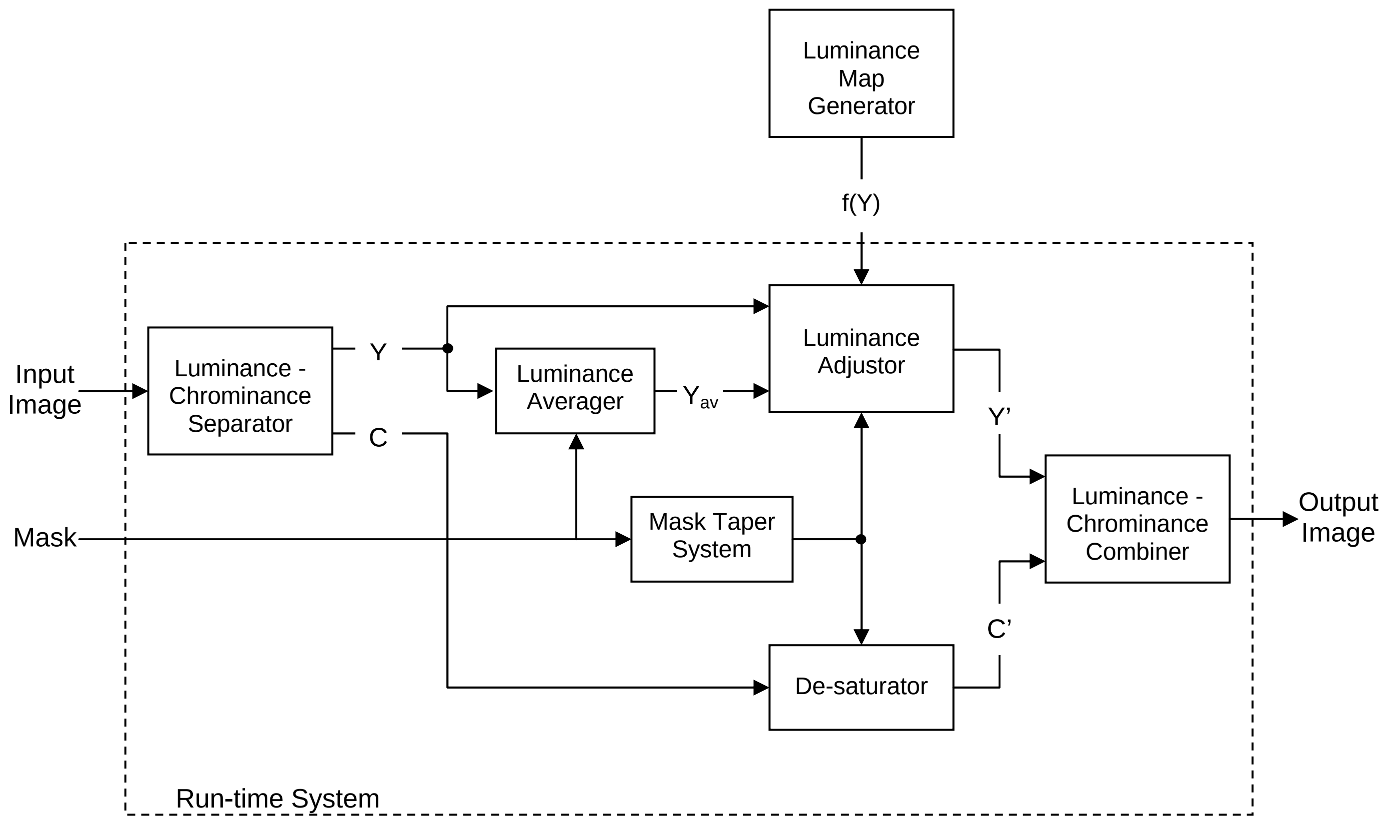

Figure 11. Schematic of correction procedure.