A "Retained Image Appearance" Metric For Full Tonal Scale, Colorimetric Evaluation Of Photographic Image Stability

Mark McCormick-Goodhart, Henry Wilhelm, and Dmitriy Shklyarov
Wilhelm Imaging Research, Inc.
Grinnell, Iowa USA

Abstract

This paper describes a computational model based on CIELAB colorimetry and full tonal scale evaluation to define objective boundary conditions for "Retained Image Appearance" in a photograph where 100% retention equates to no visual change and 0% retention means total functional loss of image information content. The "Retained Image Appearance" model, I*, examines the color information (hue and chroma) and the black-and-white spatial information (lightness and contrast) in order to characterize the complete life cycle of a photograph as it ages with respect to color, contrast, and lightness. The I* model can also be used to compare tone reproduction accuracy between different renditions of a photographic image and may therefore have applicability in initial print or proof print quality studies as well.

Introduction

Image permanence test methods for photographs routinely extrapolate results from accelerated testing environments to real-world conditions. Endpoint criteria have traditionally been based on densitometric changes which are intended to correlate with noticeable and perhaps objectionable changes in perceived image appearance.[1,2] Unfortunately, densitometry does not adequately characterize modern digital printing systems.[3] Moreover, the consumer generally believes the reported test results are the "end of life" of the photograph. Yet the chosen endpoint denotes only one point in the aging cycle of the photographic record when, in reality, the photo may continue to have informational value and therefore some "remaining life" well beyond this endpoint. Thus, image permanence predictions are subject to criticism not only for the uncertainties caused by the extrapolation of the accelerated test conditions, but also for the chosen endpoint criteria which may be too harsh or too forgiving, depending upon the end-user's requirements.

The I* model can be used to evaluate the complete aging cycle of a photograph as it changes over time, and each calculated value is a measure of image reproduction accuracy at a specific point in time. In other words, how accurate is the image reproduction in its latest state compared to its original state? In order to develop the I* metric we had to understand how colorimetry could be used to quantify tone reproduction accuracy. We considered widely used color difference models such as ΔE and more recent variants such as ΔE2000. Color difference models essentially compare the relationship between two colors in a side-by-side reference situation and weight lightness, chroma, and hue differences in a combined-term metric for overall color difference. Color scientists widely acknowledge that current color difference models are valid in a psychophysical sense for small observed changes in appearance only.[4,5] ΔE values no greater than 8–10 are often considered to be the useful limit. Although larger ΔE values can be reported, they do not continue to correlate in a meaningful visual way to the perceived magnitude of change. Because the loss of photographic tone reproduction accuracy necessarily involves large displacements in hue, chroma, and lightness, we rejected this approach and went back to basic photographic fundamentals by asking the question: what is tone reproduction accuracy? A comprehensive assessment of retained image appearance seemed possible only if full tonal scale evaluation was undertaken and if the attribute of image contrast was added to the colorimetric attributes of lightness, hue, and chroma. Also, the contribution of specific color and contrast errors to the total overall observed change depends on the area occupied by those values in a particular photograph. For example, catalytic fading between magenta and yellow dyes will have greater impact on red hues and skin tones than on other colors not comprised of magenta and yellow dyes. When a system with this failure mode is used to reproduce a scene containing large areas of skin tones (e.g., a close-up portrait) then the retained image appearance over time is likely to be poorer in comparison to a landscape scene containing mostly blues and greens. We wanted the I* metric to be able to calculate image specific results when required rather than being tied solely to an image test target with a predefined array of color patches. In order to accomplish this objective, the extraction of contrast data from the image requires a "nearest neighbor" picture element sampling method.

The I* Metric

Black-and-white photography beautifully conveys spatial information in the form of scene brightness and contrast reproduction while devoid of any chroma and hue informa-

tion (other than subtle hues which may be a coveted feature of the process itself, e.g., sepia toned prints). Color photography adds the hue and chroma attributes to the black-and-white information. It is logical and indeed important to devise a retained image appearance model that analyzes hue and chroma information separately from lightness and contrast information.[6] The color information (hue and chroma) is supplemental to the black-and-white information (lightness and contrast) in the sense that it can only be conveyed if the lightness and contrast information is reasonably intact. Nevertheless, for certain applications the color information accuracy may be essential whereas in other cases it may be of minor importance. A total I* figure of merit needs the flexibility of a weighting factor as shown in the following equation:

$$I^* = \frac{I^*_{color} + \left(I^*_{B\&W} \times \omega\right)}{1+\omega} \qquad (1)$$

The term, ω, is the weighting factor. $I^*_{b\&w}$ is the retained image appearance function for the lightness and contrast information, and I^*_{color} is the retained image appearance function for the color information.

The I^*_{color} Component

$$I^*_{color} = \frac{\sum_{n=1}^{x}\left(I^*_{color}\right)_n}{x}, \text{ where} \qquad (2)$$

$$\left(I^*_{color}\right)_n = 1 - \frac{\left(\Delta a^* b^*\right)_n - \varepsilon}{\left(C^*_i\right)_n}, \text{ for } \left(C^*_i\right)_n > 9.5$$

or

$$\left(I^*_{color}\right)_n = 1 - \frac{\left(\Delta a^* b^*\right)_n - \varepsilon}{9.5}, \text{ for } \left(C^*_i\right)_n \le 9.5$$

or

$$\left(I^*_{color}\right)_n = 1, \text{ for } \left(\Delta a^* b^*\right)_n < \varepsilon$$

$$\Delta a^* b^* = \sqrt{(\Delta a)^2 + (\Delta b)^2}\,,\ C^*_i = \sqrt{a_i^2 + b_i^2}\,,$$

and ε is a positive offset correction factor.

The term, ε, ensures that I* equals 1.00 (i.e., 100% color appearance retention) when no change has occurred by compensating for the small error caused by the squaring of residual Δa and Δb measurements that are otherwise within instrumental error limits. A default value for ε would typically be approximately 0.5. $(C^*_i)_n$ is the initial chroma of the n^{th} sampled picture element of x total elements which is calculated by the formula above using only the initial a* and b* values, denoted a_i and b_i, respectively. Δa is the difference between the initial (i.e., reference print) and final (i.e., compared print) a* values, and Δb is the difference between the initial and final b* values.

A logical and unambiguous boundary condition where I^*_{color} equals zero was derived by considering some essential aspects of the color information in a photograph:

First, as chroma goes to zero all color information is lost. When the color in a photograph fully desaturates, the photo becomes a black-and-white photo, and the $(I^*_{color})_n$ value of each picture element must mathematically approach zero as this boundary condition is approached. Thus, the $(I^*_{color})_n = 0$ boundary must expand or contract as a function of the initial chroma of each picture element (hence, the inclusion of the $(C^*_i)_n$ term in the denominator of the equation).

Second, some color scientists classify only four unique or primary hues; blue, green, red, and yellow, preferring to think of cyan as a mixture of blue and green and magenta as a mixture of blue-red. However, color photography is so fundamentally built on the principles of additive and subtractive color, that we choose to elevate cyan and magenta to equal status as unique hues. With six fundamental hues, the LCH color space of CIELAB can be divided mathematically into 60 degree sectors. As a* and b* collapse to zero, chroma decreases, but another consequence is that the determination of hue angle becomes increasingly less accurate. Hue angle is determined by calculating the arctan(a*/b*) value. As perfect gray is reached where chroma equals zero, the hue angle becomes undefined (i.e., division by zero). A practical solution to make the determination of hue well-resolved and without discontinuity is to add a seventh sector of color for gray and near grays. Thus, we define seven sectors of color as shown in Figure 1. These sectors categorize hue into seven zones; cyan, magenta, yellow, blue, green, red, and "gray". The gray sector denotes picture elements with initial chroma less than or equal to 9.5. This choice was determined by considering the CIELAB concept of just noticeable differences for the standard human observer, then setting chroma large enough so that a standard observer can accurately identify the hue of a given color with enough precision for colors differing by approximately 3 degrees in hue angle. In other words, when initial chroma exceeds 9.5, a picture element has lost all essence of grayness and acquires the color of one of the six other hue sectors to the standard observer. Retained "grayness" appearance goes to zero.

Third, when a picture element in a photograph retains chroma level but fails to reproduce hue correctly, then color information is not merely lost. It may become a completely false color which is an even more extreme loss of color accuracy than full loss of chroma. Negative values are possible, and their mathematical significance is that they represent falsely encoded color data. Consider the consequence of a hue shift for a picture element from location 1 to location 1 prime, and finally to 1 double prime as shown in Figure 1. After

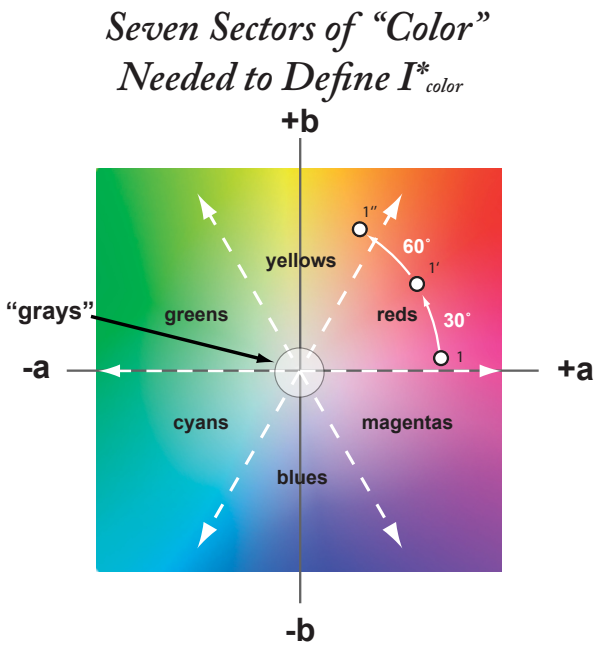

Figure 1. Perfectly neutral gray has no chroma and therefore an undefined hue angle in LAB (LCH) colorspace. Defining a special sector for grays allows a smooth transition for the calculation of $(I^*_{color})_n$ of each n^{th} picture element of x total number of elements.

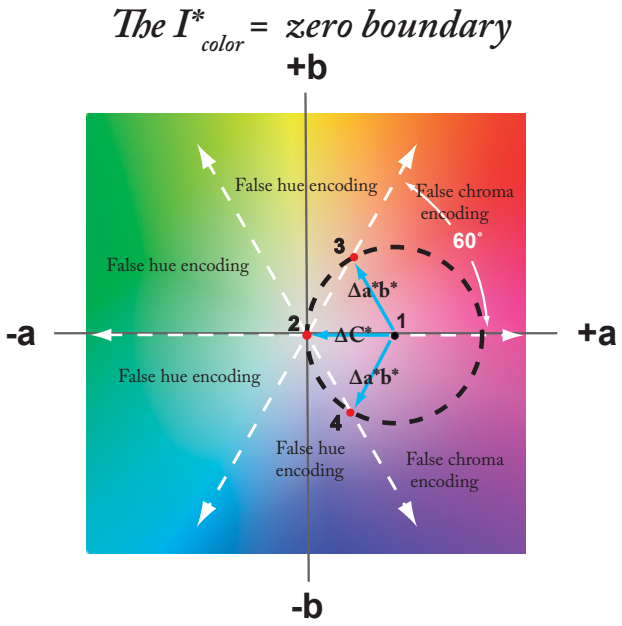

Figure 2. The $(I^*_{color})_n = 0$ boundary condition for the n^{th} picture element located at point 1 is shown by the dashed circle. When the boundary perimeter is exceeded, $(I^*_{color})_n$ becomes a negative value.

a 60 degree shift the picture element now occupies the same relative position in a new hue sector. For example, what was originally observed to be a red color element in the print now appears to be a yellow or magenta color element, depending on the direction of the shift. Thus, color accuracy is totally lost as the shift reaches 60 degrees. $(I^*_{color})_n$ must again approach zero as hue shifts approach 60 degrees, and hue shifts greater than 60 degrees should cause $(I^*_{color})_n$ to become negative, thus indicating false color encoding. To summarize, when $(I^*_{color})_n < 0$ for the n^{th} picture element of x in a photograph, the color accuracy in that location of the image is not merely lost. It is now falsely colored.

Finally, the preceding facts can be combined as illustrated in Figure 2 to derive an I^*_{color} equation that accommodates chroma changes, hue changes, and Δa*b* changes in gray sector picture elements. Equation 2 achieves a seamless mathematical transition in the way a picture element is treated right at the boundary of the gray sector and as the color moves into one of the other six hues. There is no discontinuity in the math calculation at this transition point, and this feature is a notable result of having also defined a 60 degree hue angle shift as one pathway to reach an $(I^*_{color})_n = 0$ condition.

The $I^*_{B\&W}$ Component

$$I^*_{b\&W} = \frac{\sum_{n=1}^{x} \gamma_n}{x}, \text{ where} \qquad (3)$$

$$\gamma_n = \frac{(\Delta L_i)_n}{(\Delta L_f)_n}, \text{ for } \frac{(\Delta L_f)_n}{(\Delta L_i)_n} > 1$$

or

$$\gamma_n = \frac{(\Delta L_f)_n}{(\Delta L_i)_n}, \text{ for } \frac{(\Delta L_f)_n}{(\Delta L_i)_n} \le 1$$

γ_n is the contrast retention factor of the group neighboring the n^{th} of x picture elements. In the simplest calculation, a pair of picture elements is used, and $(\Delta L_i)_n$ is the initial lightness difference (i.e., the reference measurement at time t = 0) between the n^{th} pair and $(\Delta L_f)_n$ is the final lightness difference (i.e., the comparison measurement made at time t > 0). Contrast sampling at this location in the image can be improved using a nearest neighbor sampling technique that compares a picture element to its surrounding neighbors. Figure 3 illustrates the nearest neighbor sampling approach. Equation 4 shows that γ_n at the n^{th} picture element location is the arithmetic mean of the ΔL initial differences compared to the final differences. The sampling order between the m^{th} and n^{th} picture elements must be consistent because ΔL must be assigned positive or negative status. Keeping track of positive and negative ΔL values allows positive or negative contrast relationships

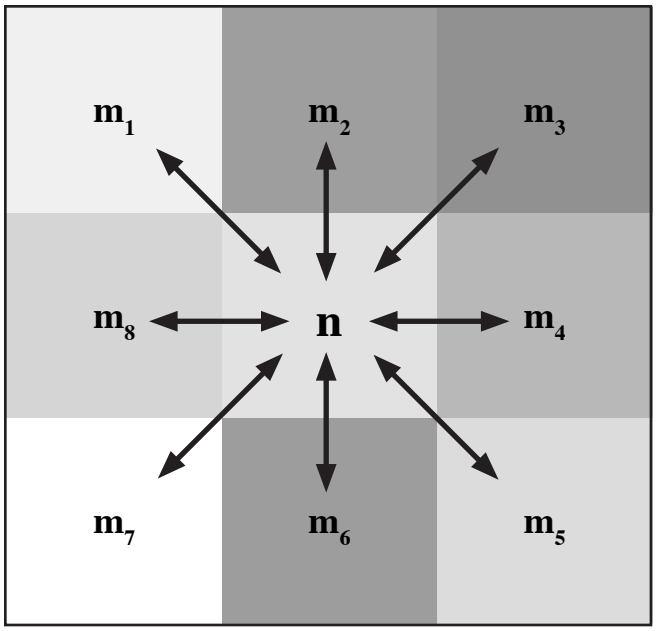

Figure 3. Nearest neighbor sampling method. ΔL initial and ΔL final values between picture elements, n and m_1, n and m_2, etc., are divided and averaged to calculate γ_n according to Equation 4.

$$\gamma_n = \frac{\sum_{n=1}^{8} \gamma_{mn}}{8}, \text{ where} \qquad (4)$$

$$\gamma_{mn} = \frac{(\Delta L_i)_{mn}}{(\Delta L_f)_{mn}}, \text{ for } \frac{(\Delta L_f)_{mn}}{(\Delta L_i)_{mn}} > 1$$

or

$$\gamma_{mn} = \frac{(\Delta L_f)_{mn}}{(\Delta L_i)_{mn}}, \text{ for } \frac{(\Delta L_f)_{mn}}{(\Delta L_i)_{mn}} \leq 1$$

to be determined. Sign reversal of contrast signifies a tonal inversion in an image which is unlikely but possible. Similar to falsely encoded color information, inverted contrast means falsely encoded black-and-white spatial information.

Figures 4a–4j are examples of images which have been digitally modified to illustrate $\mathbf{I}^*_{B\&W}$ as it progresses from 100% image retention to approximately 0%. $\mathbf{I}^*_{B\&W}$ calculates with symmetrical parity between gamma increases or decreases in the tone curve (i.e., the slope of output lightness plotted versus input lightness). When gamma goes to zero all tonal information is lost, and this situation constitutes an $\mathbf{I}^*_{B\&W} = 0$ boundary condition. No image information is discernible. Figure 4j has a faint discernible image because $\mathbf{I}^*_{B\&W}$ is approximately three percent. As gamma approaches infinity the image becomes a silhouette and possesses no continuous tone. The condition is analogous to a bit-mapped image with two bit depth of information encoding. In figure 4e, this residual information content is still enough to identify some aspects of the original scene. The tonal breakpoint in figure 4e occurred at a fortuitous value. However, with the high probability that the tonal break will not occur at an optimum value and with the full loss of continuous tone details in the image, gamma approaching infinity also defines for all practical purposes another $\mathbf{I}^*_{B\&W} = 0$ boundary condition. Thus, the basic linear symmetry to the gamma function is the essential feature of Equation 3. Consider a totally inverted tonal scale with gamma = –1 compared to the original image shown in Figure 4a. The print would look like a photographic film negative. Equation 3 computes a –100% value in this case which signifies the falsely encoded image data. This perfectly inverted condition is, of course, extremely unlikely to occur in an actual photograph, but small "flat spots" and slightly negative slopes within the full tonal scale of a print are possible as colorant mixtures fade or change unevenly in a modern digital print. The metric therefore assigns negative values to image areas of inverted contrast.

Threshold Values for ΔL

Although equations 3 and 4 are conceptually correct, a practical implementation of $\mathbf{I}^*_{B\&W}$ requires a threshold treatment for nearest neighbor picture elements that have very small initial or final ΔL values. The plain solid background of a passport photo, for example, would present a large image area where lightness is uniform and little or no visual contrast is observed. Without a threshold evaluation method, the measurement of very small ΔL changes in these areas would give rise to large deviations when dividing ΔL_i and ΔL_f terms, and consequently, an $\mathbf{I}^*_{B\&W}$ valuation which is too low. Hence, the $\mathbf{I}^*_{B\&W}$ component requires a threshold value for ΔL to compare with actual ΔL measurements. A set of conditional instructions for the γ_{mn} calculations is also needed in order to properly evaluate areas of uniform lightness in an image. The instruction set determines the correct formula to use for the γ_{mn} calculation based on the initial and final ΔL values compared to the chosen threshold value which we denote as δL. Because ΔL can be positive or negative, the threshold value, δL, can also be considered to have positive or negative direction, so the comparison is made using absolute values. Figure 5 is a schematic diagram that maps the full instruction set. There are four main paths, A, B, C, and D. Path A has two branches, A1 and A2. The arrows in the diagram trace initial and final states for ΔL in relation to $+\delta L$ and $-\delta L$. These states determine the path and possible branch to use when selecting the appropriate γ_{mn} calculation. For example, Path A is the nominal case where good contrast exists and the measured initial and

a. $I^*_{B\&W} = 100\%$ *b.* $I^*_{B\&W} \sim 75\%$ *c.* $I^*_{B\&W} \sim 50\%$ *d.* $I^*_{B\&W} \sim 25\%$ *e.* $I^*_{B\&W} = 0\%$

f. $I^*_{B\&W} = 100\%$ *g.* $I^*_{B\&W} \sim 75\%$ *h.* $I^*_{B\&W} \sim 50\%$ *i.* $I^*_{B\&W} \sim 25\%$ *j.* $I^*_{B\&W} \sim 3\%$

Figure 4. Series of black-and-white images digitally adjusted in L channel to create progressive loss in retained image appearance, $I^*_{B\&W}$. *If these had been color photos,* I^*_{color} *would also have diminished to zero because hue and chroma would not have been sustainable with such large changes in contrast.*

final states for $|\Delta L|$ are equal to or greater than the threshold $|\delta L|$. We have programmed a spreadsheet program so that $|\delta L|$ can be chosen by the user. A value between 0.5 and 2.0 appears reasonable for $|\delta L|$. More testing is needed to refine the optimum value. The results shown in figures 8–11 used $|\delta L| = 1.35$. Too small a value contributes to an underestimation of retained image appearance (i.e., an assessment that is overly critical of the final image condition). Too large a value desensitizes the model to image appearance problems that would be manifested as noticeable image mottle and graininess in the aged print.

A Lightness Error Factor

Equation 3 evaluates contrast without an additional compensation factor for changes in overall lightness level. For this reason, Equation 3 is more accurately described as a "retained information content" metric than a "retained image appearance" metric. Figure 6 illustrates this point. Photos b, c, and d in Figure 6 are tone reproductions of the same scene at the same contrast retention but at different overall lightness levels. Their tone curves were linearly scaled to lower the contrast from the photo shown in Figure 6a to those in 6b, c, and d. The reduced contrast compressed the dynamic range of the scene which in turn allowed variations in overall lightness level to be introduced without clipping any highlight or shadow detail. Had highlight or shadow detail been sacrificed to varying amounts, γ_n values of the affected picture elements would differ from those elements still on the linear tonal curve, and the $I^*_{B\&W}$ calculation for 6b, c, and d would not have stayed at 50%. Photos 6b, c, and d all retain the same information content which can be extracted, amplified, and reproduced again with greater contrast and adjusted lightness level in a new reproduction. Information content is identical though overall lightness appearance of the three photos is different. Equation 5 includes a lightness error factor, β_n, in order to account for changes in lightness of the picture elements in the photograph.

$$I^*_{b\&W} = \frac{\sum_{n=1}^{x}\left[\gamma_n - (\gamma_n \times \beta_n)\right]}{x}, \text{ where} \qquad (5)$$

for nearest neighbor sampling,

$$\gamma_n = \frac{\sum_{n=1}^{8}\gamma_{mn}}{8}, \text{ (defined in equation 4)}$$

and

$$\beta_n = \frac{\left|(L_i)_n - (L_f)_n + \sum_{m=1}^{8}\left((L_i)_m - (L_f)_m\right)\right|}{(8+1)\times 100}$$

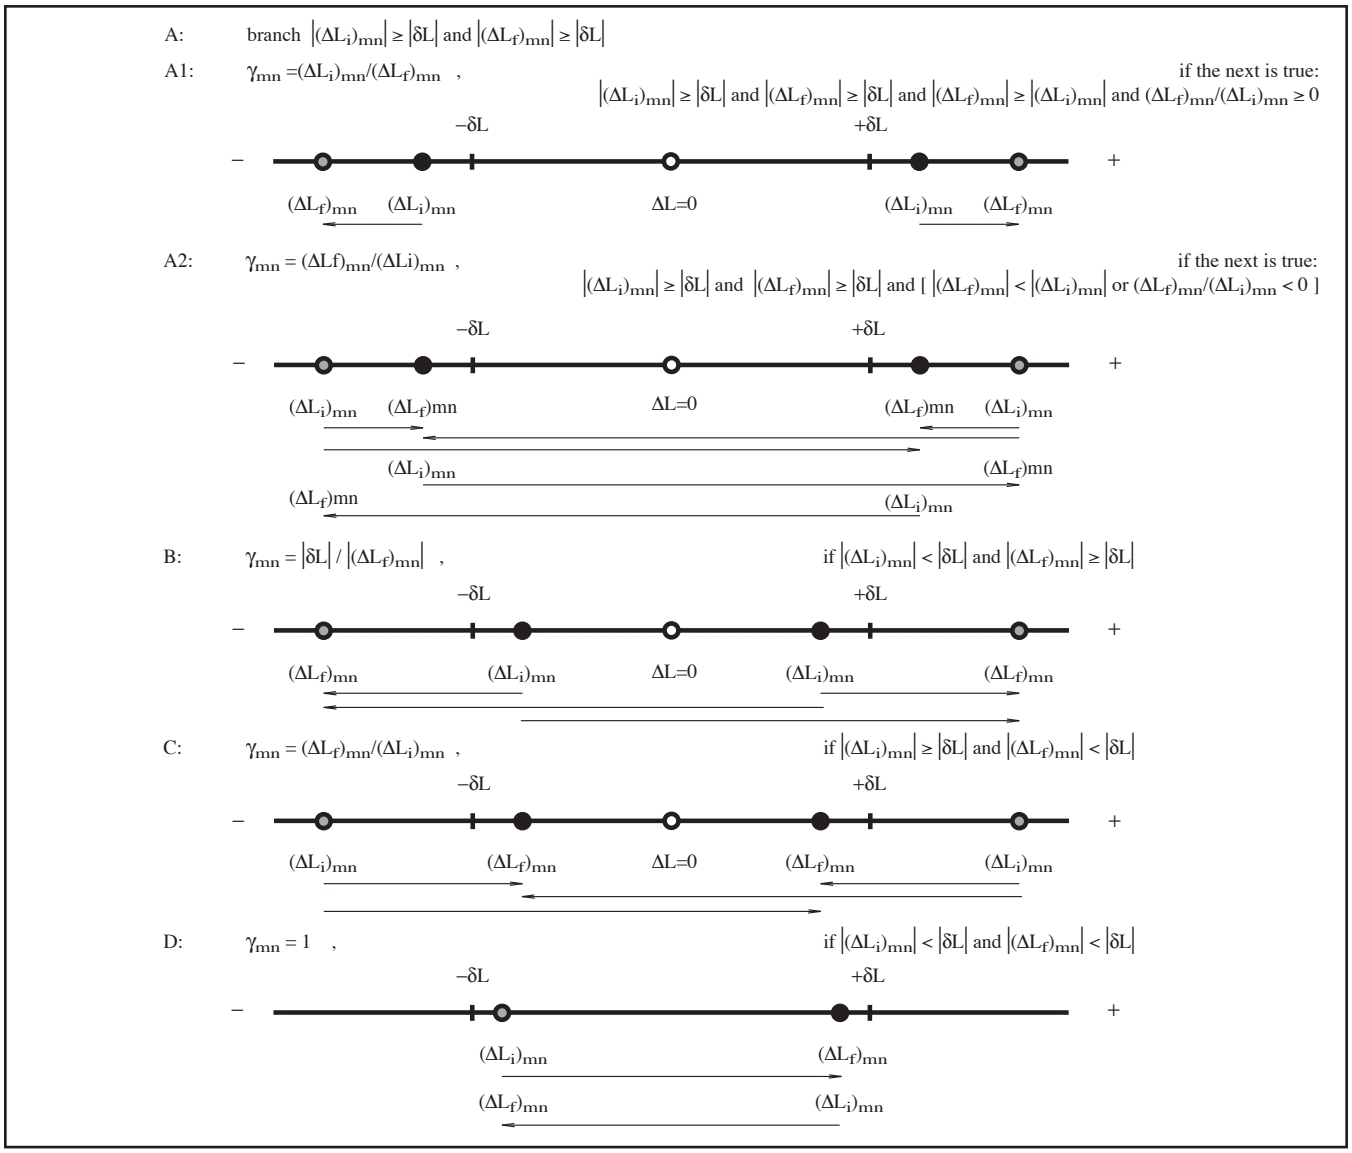

Figure 5. Schematic diagram of branched instruction set used to select the appropriate γ_{mn} calculation.

The lightness error factor could be characterized as a "retained lightness" factor if the term were [1-β]. Equation 5 produces less change in $\mathbf{I}^*_{B\&W}$ when picture elements retain their original lightness level with greater accuracy, and changes to both contrast and lightness produce lower $\mathbf{I}^*_{B\&W}$ values than comparable amounts of contrast or lightness invoked separately. Some viewers may prefer the appearance of photo 6d over 6b. However, the metric gives a higher rating to 6b because all picture elements taken collectively have less lightness error when compared to the original photo in 6a. The discrepancy may in part be due to the specific image which has important highlights in the face and the shirt that we expect to be of high lightness levels. The metric would require scene identification intelligence to determine that those highlights have special significance greater than the shadow values. A viewer preference of 6d over 6b may also be an indication that people generally accept muddy shadow detail better than muddy highlights. If the latter hypothesis is true, then Equation 5 could be modified further to weight highlight lightness accuracy higher than shadow lightness accuracy. Psychophysical testing of this hypothesis may lead to an improved metric in the future.

Limit Values for False Encoding

False encoding values significantly less than –100% for both $(I^*_{color})_n$ and γ_n picture elements are possible. Negative picture elements cancel positive elements in a mathematical sense and possibly in terms of human visual assessment of the print. However, the extent to which false encoding affects the viewer's judgement of print quality does not nec-

essarily correlate with the assessment of overall accuracy. For example, a water damaged corner in a swellable polymer inkjet print may leave 10% of the print surface area with high falsely encoded picture elements while the remaining 90% of the print is in perfect condition. The print might be deemed unacceptable for one end-user yet still hold excellent information content for another person. If falsely encoded data were truncated by substituting a limit value (e.g., 0, –1, etc.) into the summation and averaging of the *I** result, then the final result may better represent retained image accuracy. Yet allowing a greater limit below zero may correlate better with viewer ranking of prints as they become seriously degraded. Psychophysical studies are needed to study this aspect of the I* model. A false encoding limit value is simple to program. The value gets substituted at the picture element level during calculations.

Experimental Results

Figures 8–11 show *I** data for two inkjet print systems. A portrait of a mother and child was printed, and a generic test target having L channel ramps of red, green, blue, cyan, magenta, and yellow, plus grays and skin tones was also printed (see Figure 7). Figures 8 and 9 are light-fastness results for System A, and Figures 10 and 11 are the System B results. System A exhibited a severe loss of yellow dye. The color fidelity was rapidly lost and the prints appeared to turn excessively blue. However, yellow dye does not contribute significantly to the overall contrast of an image, so the prints retain a large amount of original contrast and lightness levels. System B is a more fade resistant system and is losing color information more proportionally to lightness and contrast. The System B prints still look in reasonably good condition at the 10 year extrapolated fading time, and this visual assessment is consistent with the I* metric result so far. The tests are continuing.

The I* values have been calculated with and without a limit value for false encoding. Compare graph lines b and c ($\mathbf{I}^*_{B\&W}$) and lines d and e (I^*_{color}). The chosen limit value was 0% for both I^*_{color} and the γ_n calculations. Few if any picture elements have reached the γ_n false encoding limit in either print system, so the $\mathbf{I}^*_{B\&W}$ plot was essentially unchanged by the inclusion of the limit value, and the lines overlap. The use of an encoding limit shows the greatest effect in the I^*_{color} plots for System A (Figures 8 and 9).

Each graph in Figures 8–11 also has a second y-axis which plots the percentage of picture elements that have reached or exceeded a minimum quality value. The minimum quality value can be but doesn't have to be the same as the false encoding limit value. Figures 8–11 use $(I^*_{color})_n = 0$ and $\gamma_n = 0.2$ (20%) as the respective minimum quality values. Plotting the percentage of elements that are lower than a minimum quality level is one way to evaluate the distribution of picture element quality.

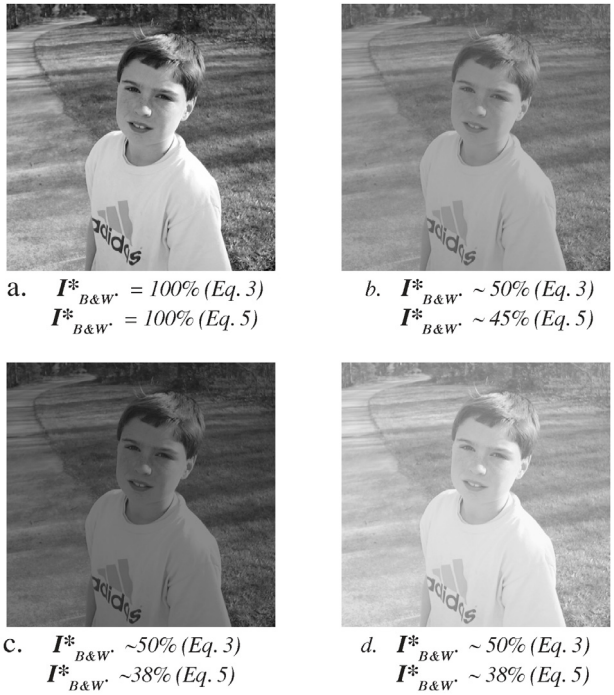

Figure 6. $I^*_{B\&W}$, calculations using Equation 3 which evaluates contrast only and Equation 5 which evaluates contrast plus lightness.

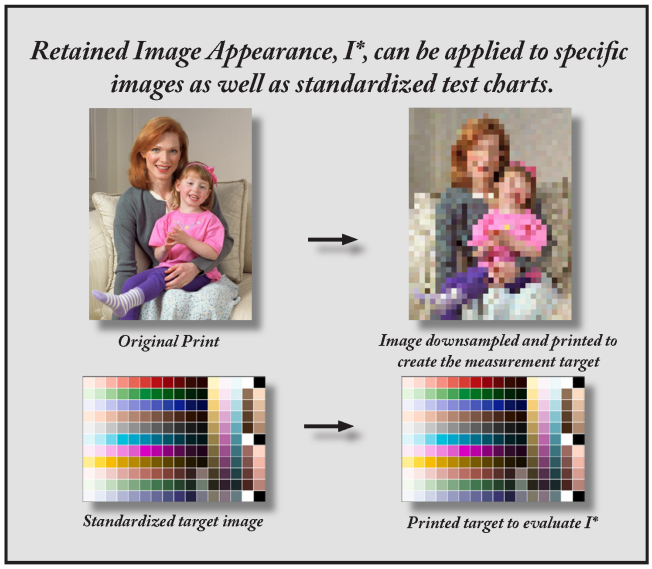

Figure 7. The pictorial image and the generic test target image used to print test samples on Systems A and B.

Percent minimum quality used in conjunction with the I* calculations may prove useful when defining image quality limits that viewers may tolerate in specific applications.

Finally, a comparison of the pictorial image results to the generic test target results reveals generally similar

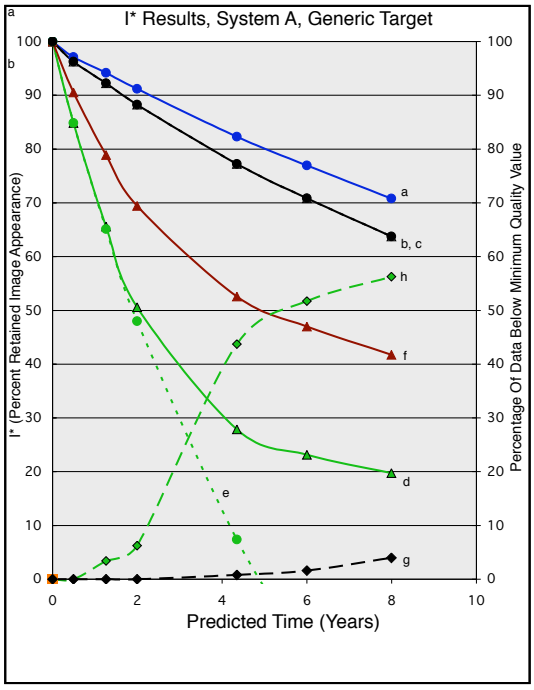

Figure 8. System A, Generic Target. Line a: $I^*_{B\&W}$ Equation 3 (contrast only), Line b: $I^*_{B\&W}$ Equation 5 (lightness and contrast) with false encoding limit = 0%, Line c: $I^*_{B\&W}$ Equation 5, false encoding not limited, Line d: I^*_{color} with false encoding limit = 0%, Line e: I^*_{color} with no false encoding limit, Line f: I^* as per Equation 1 with ω =1, Line g: γ_n percentage at minimum quality value (20%), Line h: $(I^*_{color})_n$ percentage at minimum quality value (0%).

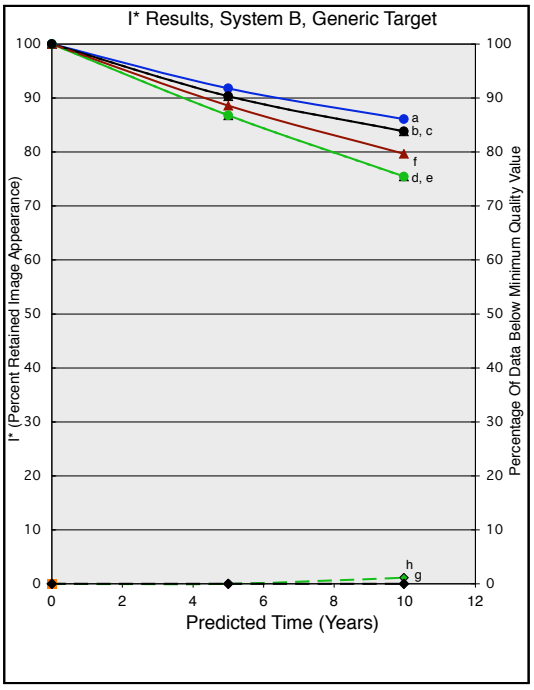

Figure 10. System B, Generic Target: Lines a–g as per Figure 8. Lines b and c as well as d and e are essentially identical because no picture elements have reached the false encoding limits.

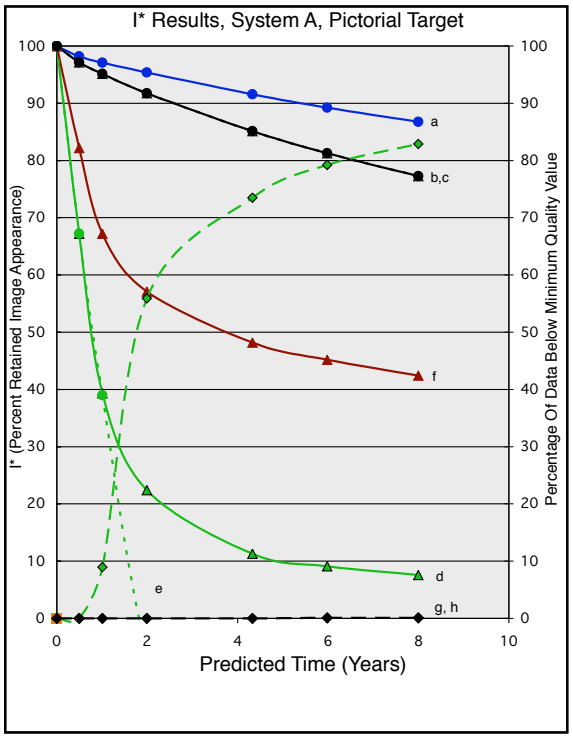

Figure 9. System A, Pictorial Target: Lines a–g as per Figure 8.

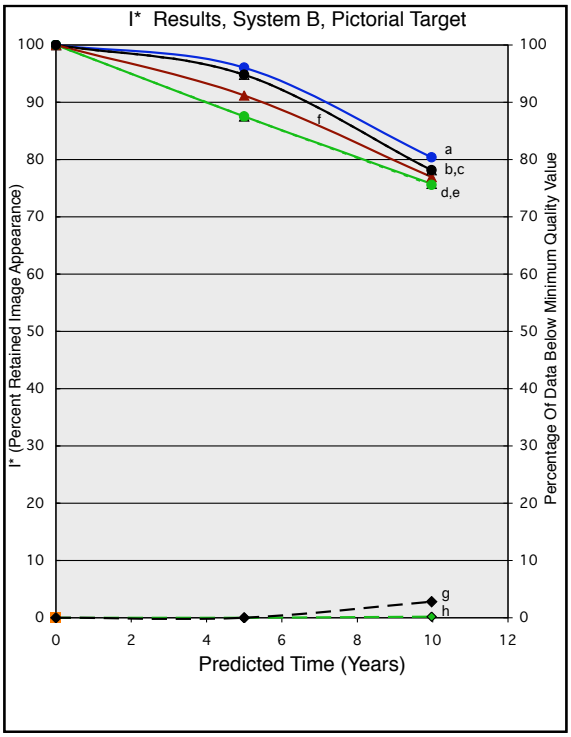

Figure 11. System B, Pictorial Target: Lines a–g as per Figure 8.

behavior but also some significant differences. The differences are especially true for System A because the portrait of mother and child has more areas of reds, flesh tones, near grays, and pale beige background that are adversely affected by the loss of yellow dye.

Conclusion

The mathematics of the I* metric have been presented in this paper. Theory of the boundary conditions, threshold values for contrast, a lightness error factor, false encoding theory and limits, and percent minimum quality values pertaining to color and contrast information, all derived from full tonal scale analysis, have been discussed. The authors programmed the I* metric with these features in a spreadsheet program, and the experimental results to date indicate that the I* metric is generating data consistent with visual observations of print aging in terms of color and contrast losses. Further psychophysical studies are needed in order to better understand the impact of false encoding values on perceived image quality and how typical observers weight color information, especially memory colors such as skin tones, versus black-and-white information when judging overall retention of image appearance.

Wilhelm Imaging Research, Inc. intends to make the I* test target files and analysis software described in this paper available to other researchers.

References

1. American National Standards Institute, Inc., *ANSI IT9.9–1996, American National Standard for Imaging Media – Stability of Color Photographic Images – Methods for Measuring*, American National Standards Institute, New York, New York, 1996. See also: *ISO 10977:1993(E) Photography – Processed colour films and paper prints – Methods for measuring image stability*, 1993.

2. Henry Wilhelm and Carol Brower (contributing author), *The Permanence and Care of Color Photographs: Traditional and Digital Color Prints, Color Negatives, Slides, and Motion Pictures*, Preservation Publishing Company, Grinnell, Iowa, 1993. Although it was recognized that assessment of the full tonal scale was desirable to fully characterize both light stability and dark storage stability behavior of color photographs, the difficulty of preparing properly calibrated full tonal scale test targets with the analog optical enlarger technology available at the time precluded putting this into general practice (see discussion on pages 77–80).

3. Mark McCormick-Goodhart and Henry Wilhelm, *A New Test Method Based on CIELAB Colorimetry for Evaluating the Permanence of Pictorial Images*. Available at <www.wilhelm-research.com> as a PDF file <WIR_CIELAB_TEST_2003_07_25.pdf>, June 16, 2003.

4. Mark D. Fairchild, *Color Appearance Models*, Addison Wesley Longman, Inc., ISBN 0-201-6346-3, MA, 1998.

5. Günther Wyszecki and W.S. Stiles, *Color Science Concepts and Methods, Quantitative Data and Formulae*, Second Edition, John Wiley & Sons, Inc., ISBN 0-471-02106-7, NY, 2000.

6. Mark McCormick-Goodhart and Henry Wilhelm, "Progress Towards a New Test Method Based on CIELAB Colorimetry for Evaluating the Image Stability of Photographs," presentation given at IS&T's 13th International Symposium on Photofinishing Technology, February 8, 2004. Available at <www.wilhelm-research.com> as a PDF file <WIR_ISTpresen2004_02MMG_HWr.pdf>, February, 2004.