

A Real-time Error-free Color-correction Facility for Digital Consumers

Rodney Shaw
White Rose Digital, Aptos, CA 95003

ABSTRACT

It has been well known since the earliest days of color photography that color-balance in general, and facial reproduction (flesh tones) in particular, are of dominant interest to the consumer, and significant research resources have been expended in satisfying this need. The general problem is a difficult one, spanning the factors that govern perception and personal preference, the physics and chemistry of color reproduction, as well as wide field of color measurement specification, and analysis. However, with the advent of digital photography and its widespread acceptance in the consumer market, and with the possibility of a much greater degree of individual control over color reproduction, the field is taking on a new consumer-driven impetus, and the provision of user facilities for preferred color choice now constitutes an intense field of research. In addition, due to the conveniences of digital technology, the collection of large data bases and statistics relating to individual color preferences have now become a relatively straightforward operation. Using a consumer preference approach of this type, we have developed a user-friendly facility whereby unskilled consumers may manipulate the color of their personal digital images according to their preferred choice. By virtue of its ease of operation and the real-time nature of the color-correction transforms, this facility can readily be inserted anywhere a consumer interacts with a digital image, from camera, printer, or scanner, to web or photo-kiosk. Here the underlying scientific principles are explored in detail, and these are related to the practical color-preference outcomes. Examples are given of the application to the correction of images with unsatisfactory color balance, and especially to flesh tones and faces, and the nature of the consumer controls and their corresponding image transformations are explored.

1. INTRODUCTION

1. 1 Background

During the past two years at this same conference, the author has described a consumer digital-image enhancement interface based entirely on an intuitive user-friendly *image-choice-only* operation [1, 2]. This ease of operation by an unskilled user has also been demonstrated practically as allowing critical-path navigation through as many as a million-plus possible individual image-quality states within a matter of seconds. It has been observed that the core image-enhancement methodology is restricted to *determinate* pixel mapping, allowing error-free transport through these many image-quality states, and thus, for example, an exact path-reversal back to the original image without accumulated signal-to-noise deterioration other than by inevitable quantization error. This was made possible by the definition of a set of orthogonal image vectors, and describing all excursions in image-quality space in terms of a fixed visual-pixel model. The author also demonstrated the crucial role of linearizing the increments of the image vectors, and showed that each vector could be represented in terms of these identical absolute increments. Further the practical concept of the Just Appreciable Visual Difference (or JAVD) was introduced, relating to a consumer preference scale-increment for the image quality of digital images (not to be confused with the term JND, as used extensively in visual psychology). In the above terms, our previous contributions concentrated heavily on the image quality vectors relating to brightness and contrast. In this present contribution we now concentrate exclusively on the problem posed by specifying color variations in this way, and defining an entirely practical color scale and interval so that color selection can be as easily presented for personal choice to the average consumer, as could the other basic image-quality variables. This posed a quite considerable challenge, since the literature on color reproduction and color preference is a weighty one, and the mere task of specifying a color transformation methodology is far from straightforward. However, our approach was to bypass this large body of knowledge and to attempt to apply exactly the same analytical approach to the individual color (RGB) vectors that had worked so successfully with the brightness and contrast vectors, since the latter had proved especially robust in the absence of formal image brightness or contrast theory, as such

The practical range of the set of orthogonal vectors, as selected to encompass the key physical variables associated with a digital image, was determined in a pragmatic manner based on the reality of the statistical image-base that it was desired to describe. For this several thousand typical consumer digital images were originally selected, covering all realistic image-quality levels, acquisition devices (cameras, scanners, video-phones), scene and lighting types, and their relevant statistics were compiled. This same validation set of images likewise covered the practical gamut of scene (and facial) types, and spanned a wide range of practical image-quality levels and conditions of color balance. Thus the pictorial database was already at hand to test the hypothesis that color-correction could be treated in exactly the same manner as image brightness and contrast.

2. USER CONTROLS FOR COLOR CORRECTION

2. 1 The Color Scale

The schematic shown in Figure 1a illustrates the broad hypothesis behind the present approach. The main assumption is that the defined orthogonal vectors for brightness, contrast and color can all be described in the same space, and although the ranges of these vectors may differ to match the overall extent of the practical data set, all may be defined on the same incremental scale within this image-quality volume. Naturally, in the color case, the scale may be applied equally to each of the color (RGB) components as appropriate. Within this latter assumption, a basic increment of color balance is no different than a basic increment in image brightness. The construction of a set of increments (JAVDs) to encompass the practical ranges of brightness and contrast, were discussed in detail in our most recent contribution at this same conference [2].




Figure 1a. Left: A schematic representation of the orthogonal image-quality vectors relating to brightness, contrast and color: the essential hypothesis, tested widely on the sample set, is that the same increment scale can be applied to each of these dimensions.

Figure 1b. Right: The pixel mapping function used on the individual color components. For generality the scales have been normalized to unity, but typically in practice they would range from 0 to 255.

Figure 1b shows this basic increment of color change, in terms of the corresponding (normalized) pixel mapping function, to be applied appropriately to the R, G or B components. In this same sense the mapping function to the left of the forty-five degree line may be thought of an increase in this component, with that to the right being a decrease. For example, and for descriptive convenience, assuming that the pixel mapping functions of Figure 1b relate to the red component, according to our present incremental vector notation we think of these as PlusR and MinusR, respectively. Color corrections based on the increments of Figure 1b were tested on a large sample sub-set of the overall consumer images described above, where this sub-set consisted of the wide range of color misbalances. The effectiveness of the color corrections and the statistics of the necessary degree of correction were noted, and it was that found successful corrections could be made to the overwhelming majority of consumer images originally having unsatisfactory color balance.

The conclusions from this validation exercise confirmed the orthogonal assumptions of the overall methodology, and can be summarized as follows:

- a) The color correction interval as represented in Figure 1b was confirmed as representing the basic interval of color (JAVD) below which the user became indifferent to the degree of change.
- b) Suitable (RGB) combinations of these unit changes allowed satisfactory practical corrections of color balance according to user preference for digital images having original misbalances that ranged from subtle to severe.
- c) The overall change made by a combinations of these RGB intervals is independent of the order in which they have been applied.
- d) All such changes are substantially error-free, and thus reversible by applying the sequence of changes (in any order) of opposite effect.
- e) Typically the number of combined interval changes for the majority of consumers to reach final preference fell in the range of one to four overall, with a very small minority of images calling for as many as ten such changes.

With hindsight, the fundamental reason for these ‘well-behaved’ characteristics of the practical color corrections arises from the much more restricted range that is necessary for effecting satisfactory color change, than the ranges found for brightness and contrast change. Thus, the orthogonal assumption underlying Figure 1a was most closely observed in practice for changes in color. This is not to say that users attempting to correct facial or overall preferred appearance did not sometimes first make significant changes in brightness and contrast (in fact, by definition, for many image types it was simply not possible to do so), but merely to state that color changes *per se* were found to have all the practical virtues listed above in the overwhelming majority of practical cases.



Figure 2. An illustration of the five-images-at-a-time user-choice methodology, as applied to a hierarchy of brightness and contrast intervals [1, 2], where each of the brightness and contrast choices are typically part of a three-stage hierarchy varying from coarse, through medium to fine (a single JAVD).

2. 2 User Interface

Even allowing for these satisfactory practical results, the challenge remained of how to present the UI to the consumer in order to ensure ease of use and to be free from ambiguity. We recall that for brightness and contrast, allowing the user to reach the preferred state of personal choice becomes one of incorporating a straightforward hierarchical procedure for which an information-theoretic optimization procedure can be applied [1,2]. In the case of color, however, the obvious problem arises of the number of dimensions involved. Three separate five-image scales relating to the separate RGB increments is scarcely practical, and would prove prohibitively complex and confusing to the unskilled user. However, again by resorting to our large database of typical consumer images, a solution was rapidly found to this dilemma, based on several practical conclusions.

The first of these was the observation that, for the small changes typical of color correction, and as validated by our practical observations, described above, then in our notation, any two small changes in the color vectors are equivalent to a defined change in the third. For example, in general terms, PlusA and PlusB are equivalent to MinusC. Thus any two color components may be chosen as variables for our methodology, and the question then becoming that of choosing the most appropriate of these, leading naturally to the second observation. This was an aspect of the large validation image-database itself, and the statistics of the brightness, contrast and color-balance changes that were found to be appropriate. From these it was observed that more than ninety-nine per cent of all images needing color correction could be conveniently accomplished by using the RB variables. This observation was later substantiated by a considerable body of independent statistical data [3], as shown below, and based on consumer response to the characteristics of around three million digital images, as summarized in Figure 3.

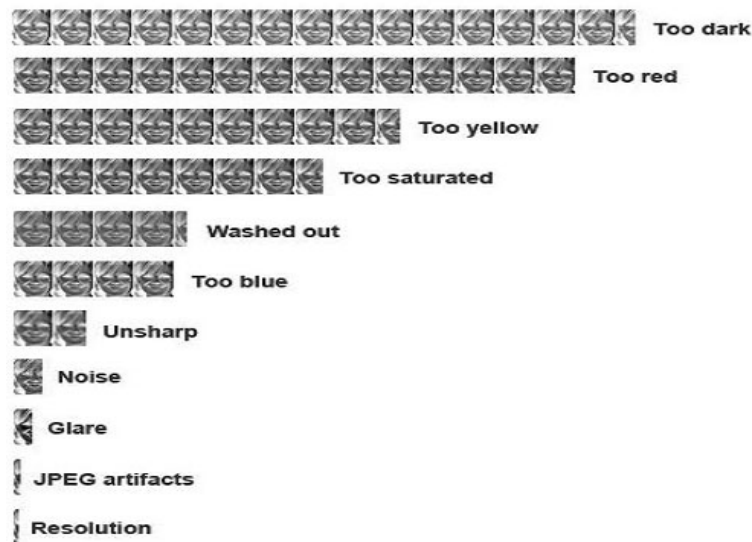


Figure 3. Statistics taken from a large database showing the reasons consumers considered digital prints to be unsatisfactory [3].

This independent consumer data showed a close correlation with and provided a broad verification of our own statistics, and confirmed the RB variables as being most appropriate for our purpose. In the broadest empirical terms, hindsight renders this conclusion obvious. Manufacturers of digital devices have little difficulty in providing the appropriate spectral response in the mid-region of the visible spectrum, and both the finesse and the resulting approximations and limitations typically occur at the extremes of the distribution. Thus providing a color-correcting mechanism that fine-tunes the color balance at the outer region provides the maximum leverage and control to the user.

Given the choice of color controls, the question remains of organizing this within the existing methodology in a similar user-friendly manner, such that its use will be unambiguous and lead to preferred color choice in the most direct manner. After considerable experimentation and consumer testing, the scheme illustrated in Figure 4 was selected (it should be noted that the actual color versions of images shown here and throughout in gray-scale can be inspected at www.whiterosedigital.com). Although the five-image choice scheme differed in concept from those for brightness and contrast, and the choice hierarchy operated on the basis of a re-set and re-calculate procedure, it was found that users rapidly adapted to the choices, and described the five-image choice scale as providing 'warm' to 'cold' versions of their originals, as was the intent. The 'choose-and-reset' hierarchy, as described previously [2], is illustrated in Figure 5.



Figure 4. The user-choice images provided for image color correction, allowing an appropriate color change to be compared to the original (center).




Figure 5. Illustration of the color choice hierarchy: any desired color component may be selected in units of individual JAVDs, and the image recomputed, reset and available for subsequent choice of color balance.

2. 3 User Evaluation

The above color-correction methodology and user-interface have been tested extensively on several thousands of consumer images requiring a very wide range of practical changes in color balance. The results from this experience have confirmed the suppositions outlined above, and especially that the overwhelming majority of consumer images can be restored to preferred color balance without difficulty or ambiguity by the unskilled user. The difficulties that were encountered in a very proportion of all images (less than one percent) included atypical and extreme color misbalance, and images having insufficient captured signal-to-noise ratio in one or more of the color channels: the latter class are of course resistant to any form of color correction, no matter how elaborate. However, there remains a small but stubborn segment of the consumer market that resists all but the elusive 'one-click-fix' button, and we now review problems and progress in this field.

3. AUTOMATIC IMAGE ENHANCEMENT

3. 1 Practical Approach

The advantages of the basic enhancement procedure described above make it readily accessible to approaches based the use of artificial intelligence (AI). Properties such as the absolute nature and reversibility of each step in the image optimization procedure render it ideal for the parallel recording of each individual image enhancement procedure, and the correlation between the most appropriate measures of the initial and final image statistics. Already promising results have been obtained relating to the specific characteristics of image-acquisition devices, nature of scene-types, the specific lighting conditions and the display-technology defining the viewing environment, right through to individual preference.

To date, we have successfully developed automatic methods for the correction of images deficient in brightness or contrast. In these cases the incoming image statistics offer sufficient clues of their deficiency, and thus offer a possible automatic source of data for the remedy. The color histograms offer such clues, and we have already developed a method of mapping directly from the histogram statistics to the functions defining the brightness and contrast vectors of our basic methodology. Illustrative examples of these are shown below in Figure 6. The nature of the respective before and after histograms is as might be expected in both these cases. In the brightness case, clustering at low pixel values is replaced by a broadening across the entire pixel range, while in the contrast case a clustering in the mid-pixel range is replaced by a similar broadening across the range.

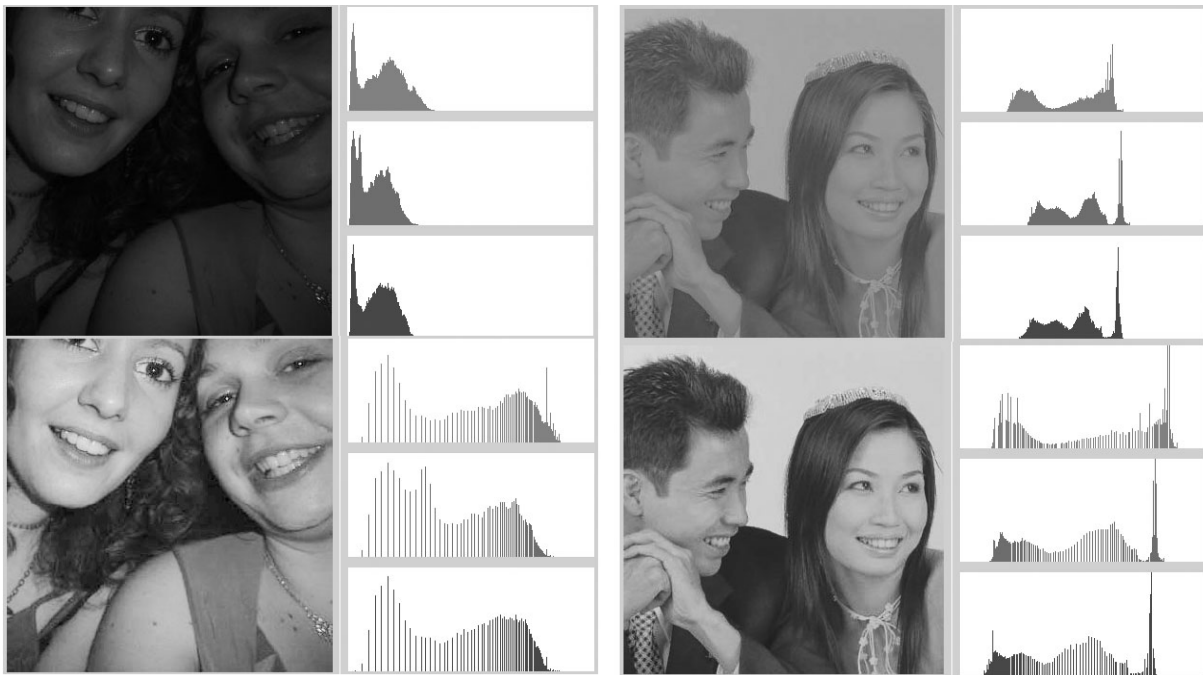


Figure 6. Left: Before and after versions of an image requiring brightness correction, along with the corresponding before and after RGB histograms (in descending order). Right: Before and after versions of an image requiring contrast correction, also with the before and after histograms.

3. 2 Color Corrections

The practical situation becomes more complicated when attempting to design automatic color-correction controls. As discussed previously, when offered the choice due to a simple enhancement facility, the correction of faces and flesh tones becomes a dominant aim among the majority of consumers. Furthermore it can readily be established that a majority of consumers prefer ‘warm’ flesh tones as opposed to the ‘cold’ blue tints often found, even with the most advanced digital cameras, and we have had successes in establishing a recipe sequence of the color corrections illustrated in Figure 4 to this class of images. However a large proportion of faces also require corrections in brightness, and a lesser proportion in contrast, prior to a change of color balance.

Figure 7, left, shows a facial image first requiring an overall increase in brightness prior to color change, and the histograms yield clues to these obvious changes. However in Figure 7, right, an image is shown requiring only a color rebalance, and the before and after versions are then at approximately equal luminance. In this case the changes in the RGB histograms are subtle, and that of the original yields no obvious clue as to the need for color correction, although the user has no difficulty in both seeing the need for and making a significant shift.

It should be noted that there is significant contemporary activity in face recognition, especially at the acquisition stage, and then subjecting this part of the image to various degrees of favorable enhancement prior to display or printing. In this case the enhancement is typically that of image segmentation, where *conditional* pixel-mapping is used, and unless carefully applied usually may introduce the usual artifacts associated with such techniques, and such artifacts are especially unacceptable to consumers in the reproduction of faces.

Our own initial to the problem, where there is *a priori* segmentation of facial areas, is to use modest automatic relative color shifts towards the warmer flesh tones preferred by the majority of consumers. But it should be stressed that the solution provided by such an approach may often be considerably less satisfactory to the user than those provided by solutions for the entire scene. In fact it is well known but little appreciated that enhancement solutions for any region of a scene depend intimately on context. Thus a chosen color balance solution for an isolated face may look worse than the original when replaced in original context. Our own solutions to this problem, taking all such factors into account, are part of our ongoing investigations into automatic solutions, and will be reported at future conferences in this series.




Figure 7. Left: Before and after versions of an image requiring brightness, contrast, and color corrections, along with the corresponding before and after RGB histograms (in descending order). Right: Before and after versions of an image requiring only color correction, also with the before and after histograms.

SUMMARY AND CONCLUSIONS

Individual preference for image quality is tied closely to the perception of faces and flesh-tones, and in turn, this defines the preferred color balance. In this context we have described a simple color-balance choice methodology, as part of an overall user-friendly enhancement procedure for digital images. This color methodology was designed from first principles using the same approach as that employed for the correction of the brightness and contrast attributes of consumer digital images. By virtue of its ease of operation and the real-time error-free nature of the color-correction procedure, it can readily be inserted in appropriate software form anywhere a consumer interacts with a digital image, from camera, printer, or scanner, to web or photo-kiosk.

The overall enhancement methodology was based on describing the image characteristics (brightness, contrast, color) by a framework of orthogonal vectors covering the image quality volume appropriate for a large consumer database of typical digital images. It was found that vectors describing color balance could be defined within this overall set, and in such a way that small absolute increments could be translated into user control, enabling rapid choice of preferred color balance. Thus consumers could readily achieve their preferred facial color, typically several absolute increments away from that acquired by the digital device. From a large number of digital images corrected in this way, it was also noted that scenes in general, and facial characteristics in particular, could not be controlled by preferred changes in color balance alone, but usually also involved the balancing of the brightness and contrast image quality components.

Progress towards automatic scene and facial correction were also reported. Due to its absolute nature, the basic methodology is ideally suited for this purpose, and lends itself naturally to the collection of a large data base of preferred user reproduction, that in turn can be related to the statistical properties of the original image. Correlation between the two then permits an automatic link. We have described some of the problems associated with this approach, and shown some of the earliest results in this direction.

REFERENCES

1. R. Shaw and P. Johnson, *Ubiquitous Image processing: A Novel Image-Enhancement Facility for Consumers*, Proceedings IS&T/SPIE Electronic Imaging, Color Imaging XI, SPIE Vol. 6058 (2006).
2. R. Shaw, *A User-friendly Digital Image Processing Procedure: Technical Implementation*, Proceedings IS&T/SPIE Electronic Imaging, Color Imaging XII, SPIE Vol. 6493 (2007).
3. See data provided at: www.smugmug.com.
4. See color images at: www.whiterosedigital.com.