

Color Histogram Specification by Histogram Warping

Mark Grundland* and Neil A. Dodgson

Computer Laboratory, University of Cambridge, 15 J. J. Thomson Ave., Cambridge CB3 0FD, UK

ABSTRACT

Histogram warping is a novel histogram specification technique for use in color image processing. As a general purpose tool for color correction, our technique constructs a global color mapping function in order to transform the colors of a source image to match a target color distribution to any desired degree of accuracy. To reduce the risk of color distortion, the transformation takes place in an image dependent color space, featuring perceptually uniform color axes with statistically independent chromatic components. Eliminating the coherence between the color axes enables the transformation to operate independently on each color axis. Deforming the source color distribution to reproduce the dominant color features of the target distribution, the histogram warping process is controlled by designating the color shifts and contrast adjustments for a set of key colors. Assisted by mode detection, matching quantiles establish the correspondence between the color distributions. Interpolation by monotonic splines serves to extend the mapping over the entire dynamic range without introducing artificial discontinuities into the resulting color density. We show how our method can be applied to color histogram equalization as well as color transfer from an example image or a color palette.

Keywords: Color image processing, color histogram specification, color histogram modification, color histogram equalization, color mapping function, color transfer, color palette, color space, mode detection, monotonic splines

1. INTRODUCTION

Histogram specification, sometimes called histogram matching or histogram normalization, is a basic signal processing operation. It is the transformation of one histogram into another by remapping the signal values to control the relative frequency of their occurrence. Histogram specification is the generalization of histogram equalization, which stipulates a uniform output histogram, making each signal value equally likely in order to maximize the signal's entropy and thus optimize the information content of the quantized signal. As a preprocessing step in signal analysis or a postprocessing step in signal synthesis, histogram specification is used to adjust the statistical profile of a signal's dynamic range while retaining the local structure and relative ordering of the signal values. For over thirty years, histogram specification has been the subject of parallel research with a wide variety of applications, including image synthesis, digital photography, video processing, remote sensing, medical imaging, computer vision, and speech recognition. In the context of image processing, histogram specification is primarily a technique of global contrast enhancement. In the context of pattern recognition, histogram specification is used to compensate for differences in data acquisition, improving the performance of a classifier by removing the influence of test conditions that were absent from the training environment. Despite a wide variety of available approaches, histogram specification remains an unresolved practical problem, especially for multidimensional data. The measure of success depends on the application. In the case of image enhancement for visual inspection, it may be subjective. As Gonzalez and Woods remark in their textbook¹, "histogram specification is, for the most part, a trial-and-error process". Clearly, there is still scope for improvement.

We present a novel histogram specification technique for use in color imaging[†]. Our histogram warping framework is designed to make color correction tasks easier to control. Given a source image and a target color distribution, our task is to construct a global color mapping function that transforms the colors of the source image so that the colors of the resulting image match the target color distribution to any desired degree of accuracy. Our approach to histogram modification encompasses a range of intermediate transformations, from histogram stretching and histogram moment



* Further author information: (Send correspondence to Mark Grundland)

Mark Grundland: Mark@eyemaginary.com

Neil A. Dodgson: Nad@cl.cam.ac.uk

[†] To obtain a version of this paper with color illustrations, please visit:
<http://www.eyemaginary.com/Portfolio/Publications.html>



Figure 1. Color transfer of lighting conditions using our histogram warping transformation in $La'b$ color space. The colors of the original images (top) are exchanged (bottom).

specification to exact histogram specification. Our approximate histogram specification technique is versatile enough to adapt to various applications. The target color distribution can be specified by a probability density function, an example image, or a color palette. For instance, the transfer of a color scheme from one image to another (Figure 1) can be useful in adjusting the illumination. In applications where speed is crucial, as the analysis time dominates the synthesis time in our algorithm, it is possible to perform the analysis on subsampled image data.

We treat histogram specification as a continuous transformation of a random variable expressing the distribution of colors along a color axis. To reliably apply this operation independently to each color axis, we advocate the use of a color space adapted to the statistical properties of the source color distribution, offering perceptually uniform color axes with statistically independent chromatic components. Eliminating the coherence between color axes reduces the risk of color aberration inherent in the decomposition of the 3D color histogram specification task into three computationally efficient 1D histogram transformations. The transformation of each color axis is defined by mapping a set of quantiles of the source distribution to the corresponding quantiles of the target distribution. Through the selection of these key colors, we can easily set the accuracy of the transformation. To better reproduce the target color distribution, we can guide the transformation according to the dominant colors of the target color distribution as reflected in the modes of its histogram. To smoothly extend the color mapping over the entire dynamic range of each color axis, we fit a monotonic interpolating spline to the matched color values of the corresponding quantiles. In effect, our continuously differentiable transformation ensures that color and contrast change in a continuous and predictable manner, and thereby reduces false contouring artifacts. By separately specifying the color shift and contrast adjustment for a set of key colors, our histogram warping technique can easily be configured to perform any monotonic transformation of a histogram.

Our previous work^{2,3} introduced histogram warping as a technique for interactive and automatic contrast enhancement of graylevel images. This paper expands our histogram warping framework to handle histogram specification for color

images. We start off with a survey of existing approaches for histogram specification and its extension to color processing. Subsequently, we describe the components of our algorithm: the color space conversion, the color mapping function, the color feature matching, and the color contrast adjustment. Finally, we demonstrate the applications of our technique, including color histogram equalization as well as color transfer from an example image or a color palette.

2. SURVEY OF RELATED WORK

2.1. Algorithms for Histogram Specification

Histogram specification^{1,4-6} traces its roots to histogram equalization^{1,7}, one of the earliest image processing algorithms which is still widely used for contrast enhancement. Its classical implementation is based on the fact that transforming a random variable by its cumulative distribution function results in a uniform distribution. Histogram specification^{1,4-6} is performed by using the source cumulative distribution function to map the source histogram to an equalized histogram and then using the inverse of the target cumulative distribution function to reshape the equalized histogram to reproduce the target histogram. Typically it is used to specify the distribution of colors in an image. It can also be applied to render spatial structures⁸ by specifying the histograms of the subband coefficients of an image decomposition formed by resolution and orientation selective filters, such as a steerable image pyramid. For antialiasing, the source image values may need to be appropriately preprocessed^{4,5,9}. Histogram specification can be performed more directly¹⁰ by constructing a lookup table to provide a single-valued mapping between the source and target histograms. Histogram specification by sorting¹¹ performs a multi-valued mapping between the order statistics, where each element of the source distribution is mapped to the correspondingly ranked element of the target distribution. Under various measures of discrepancy between distributions, the optimization of the histogram specification transformation can be framed in terms of tree search^{12,13}, integer linear programming¹⁴, dynamic programming¹⁵, or a histogram distance metric¹⁶. It is possible to rigorously derive an intermediate mapping that balances the discrepancy with respect to multiple target distributions¹⁷. All of these exact histogram specification techniques generate a discrete mapping designed to reduce quantization error, typically without any account for the distortion of the source histogram that the mapping entails. Forcing the source image to conform as exactly as possible to an arbitrary target color distribution necessitates abrupt changes in color and contrast. By relaxing this constraint, our approximate histogram specification technique results in fewer image artifacts.

The simplest method for approximate histogram specification is a linear transformation, commonly used in color processing to specify the range¹⁸ or the mean and standard deviation¹⁹ of the resulting color distribution. Specification of higher order statistical moments²⁰ is more difficult, requiring the numerical optimization of a nonlinear transformation. A more adaptable histogram modification method models the histogram as a tree of nested intervals²¹, allowing for a hierarchy of linear transformations. Piecewise linear mappings^{1,22-24} can approximate histogram equalization and histogram specification. However, they lack the flexibility for separate adjustment of color and contrast. As previously proposed for use in histogram modification, piecewise exponential splines can fail to be continuously differentiable²⁵ while piecewise cubic splines can fail to be monotonic⁶. Our histogram warping transformation is purposefully designed to resolve these deficiencies of previous piecewise defined methods for histogram modification.

2.2. Color Processing for Histogram Modification

Traditionally, color spaces have been developed for the convenience of physical measurement (XYZ), physiological measurement (LMS), perceptual measurement (Lab and Luv), user control (HSV , HSB , and HLS), or hardware design (RGB , YIQ , $YCbCr$, and CMY). By default, these standard color spaces share a common weakness. Since their fixed color axes disregard the distribution of colors in the image, they can not eliminate the statistical dependence between the color axes. Removing the correlation between the color axes prior to independently modifying the histogram of each color axis has proven useful in texture synthesis⁸ and remote sensing¹⁸. We improve on this approach by selecting color axes which exhibit minimal statistical dependence. Rather than adapting the color space to the image, it is always possible to devise a color space that reflects the statistical properties of a representative ensemble of images. For such a decorrelation strategy to be effective, the color space should only be applied to color distributions that are similar to the training set. For instance, recent color transfer experiments¹⁹ have tried using the $l\alpha\beta$ color space. Defined on the basis of a dozen wilderness scenes, the $l\alpha\beta$ color space features logarithmic axes to model visual perception. Instead, our color space is founded on Lab . However, even the standard CIE color spaces, Lab and Luv , are only designed to be perceptually uniform for just noticeable color differences and they still suffer from clearly visible distortions for some hues²⁶. While

research is continuing on newer perceptually uniform color spaces²⁷, such as *DIN99c*, their effectiveness in practical applications remains to be seen. The large color differences commonly encountered by our algorithm are better modeled by color order systems, such as *Munsell*. As they are usually defined through a collection of printed color samples, the need for configuring 3D lookup tables²⁶ discourages their use. Sophisticated color appearance models, such as *CIECAM97s*, remain difficult to apply in practice because they require correct assumptions about the viewing conditions.

Independent processing of appropriately chosen color axes has the decisive advantage of speed and simplicity in many practical applications of color adjustment. Color axes that correspond to perceptually meaningful attributes make color processing easier to control. Relying on the color axes to impose a sensible ordering of colors, standard 1D histogram modification techniques do not readily extend to 3D histogram transformations. Without a natural ordering relation to arrange colors in sequence, it is harder to derive a well founded matching between color distributions²⁸. However, 3D color processing has the potential to perform more powerful and flexible transformations. Not every transformation of the joint color distribution can be decomposed into monotonic transformations of the marginal color distributions. For instance, a 3D transformation is necessary when an arbitrary correspondence between a source set of input colors and a target set of output colors is used to define the mapping of the entire color space. By taking into account the shape of the available color space, a 3D transformation can avoid mapping colors out of gamut. Yet, these methods remain unpopular because their complex implementations invariably demand much greater processor or memory resources. Existing methods for multidimensional histogram specification by a multivariate transformation require either an iterative solution to a partial differential equation²⁹ or a numerical solution to a nonlinear minimization problem³⁰.

An alternative approach employs conditional color distributions, where a profusion of color axes induces an appropriate partition of the color space. When used in color image enhancement to improve visibility of image features without altering the interpretation of image content, the color space decomposition is often designed to ensure that the hue values are left unchanged^{31,32}. A general purpose color space decomposition³³ can consist of the luminance histogram, the hue histograms conditioned on luminance values, and the saturation histograms conditioned on hue and luminance values. In the histogram explosion approach³⁴, the color axes are rays emanating from a common origin so that each histogram is conditioned on its spherical angles in the color space. The histogram decimation approach³⁵ recursively subdivides the color space into regions of equal volume while suitably displacing the colors of each region. These computationally intensive approaches are prone to producing false contouring artifacts whenever they result in apparent discontinuities between the independently obtained color mappings of adjacent regions of the color space. The same problem affects attempts to discretize the color histogram specification process by performing color palette mapping³⁶.

3. COLOR SPACE

When mapping the colors of the source image to match the target color distribution, our algorithm performs a continuous deformation of the color space by warping each color axis independently. In our experiments, the choice of color space is observed to have three substantial effects on the outcome. The sequence of color tones along each color axis influences the color correspondences inferred by the algorithm. The nonlinearity of the color axes governs the perceptual uniformity of the color transitions, resulting in the smoothness of the color gradients. The statistical independence of the chromatic axes is required to ensure that separate transformations can be applied independently to each axis without disturbing any statistical relationship between them. For instance, the direct use of *RGB* color channels gives unpredictable results since these axes are both highly correlated and perceptually nonuniform. Independently altering the statistically dependent color primaries in an additive color space tends to produce unexpected hues. Our color histogram specification technique is carefully formulated to avoid introducing colors that are absent from both the source and target color schemes.

For statistically independent and perceptually uniform color axes, we advocate the use of a color space that adapts to the source image. We start with the perceptually uniform *Lab* color space. Observe that any linear combination of the *ab* chromatic axes remains a perceptually uniform color axis. By translation and scaling, we standardize the chromatic axes to have zero mean and unit variance. Rotating the axes to remove the correlation between them, we perform principal component³⁷ analysis *a'b'* using a singular value decomposition of the covariance matrix. While decorrelating the chromatic axes is adequate for many cases, particularly difficult images can benefit from additional processing because decorrelation generally does not produce statistically independent chromatic axes. To further alleviate the side-effects of operating on each chromatic axis separately, we apply independent component³⁸ analysis *a''b''* to find a linear transformation that minimizes an estimate of the mutual information between the axes, a measure of their statistical

dependence. In effect, the resulting axes are apt to reflect the distinguishing characteristics of chromatic distribution since their projection is chosen to make the histogram of each axis appear as dissimilar from a Gaussian distribution as possible. Observe that independent component analysis can only improve on principal component analysis as long as the independent components are not Gaussian since uncorrelated Gaussian random variables are already independent.

Using a set of 16 color transfer experiments on stock photographs, we tested a wide variety of color spaces (including *XYZ*, *LMS*, $l\alpha\beta$, *HSV*, *RGB*, *R'G'B'*, *R"G"B"*, *Lab*, *L'a'b'*, *L"a"b"*, *La'b'*, and *La"b"*). As expected, a color space that produces reasonable results for one image can not be guaranteed to be equally effective for another. Yet, some color spaces appear much more prone to color distortion than others. For color histogram warping, we selected the *La"b"* as our color space because, under our subjective visual inspection, it gave the most consistently reliable results, especially in the presence of large color shifts. By separating the brightness of the luminance channel *L* from the hues and tones of the independent chromatic components *a"b"*, our algorithm permits novel shades and tints to occur while reducing the possibility that unforeseen hues arise. Such an approach to global color adjustment is in agreement with human color perception, where a change in brightness intensity signals a readily discounted change of illumination while a change of hue category signals a more significant change of scenery.

When our algorithm results in an out of gamut color, we project it in the direction of its original color in the source image and clip it against the nearest face of the *RGB* color cube. Recent studies³⁹ have found that viewers tend to favor gamut clipping over gamut compression. With the independent component analysis resulting in nonorthogonal chromatic axes, occasionally a significant portion of the color distribution can be mapped out of gamut. In practice, should such a problem be detected, the algorithm could be instructed to fall back from *La"b"* to *Lab* or *RGB*. Aliasing artifacts may also arise whenever the mapping stretches the contrast of quantized color values. Because our histogram transformation operates on continuous color values rather than discrete color levels, we need to augment each quantized source color with a suitable fractional part. To mask aliasing artifacts, we slightly perturb the source color values with white noise⁹. For higher quality rendering, one may use a blue noise dithering mask⁴⁰.

4. COLOR TRANSFORMATION

For each color axis, the histogram warping transformation is controlled by defining its effect on a set of key colors, with its displacement determining the color shift and its slope determining the contrast adjustment. Thus, the transformation $y = T(x)$ is specified by the mapping of corresponding key colors values $b_k = T(a_k)$ along with their contrast adjustments $d_k = T'(a_k)$. In this way, we have separate and independent control over how the color histogram is shifted $a_k \neq b_k$, stretched $d_k > 1$ or compressed $0 \leq d_k < 1$. We use a piecewise defined, interpolating spline to ensure that these parameters have local effect. Our only constraint is that the parameters must delineate a valid monotonic function, with $a_{k-1} < a_k$, $b_{k-1} \leq b_k$ and $0 \leq d_k < \infty$. In addition to being potentially invertible, a monotonic transformation serves to preserve the natural order of colors and thus avoids the banding artifacts caused by local reversals of image polarity. We require a continuous transformation in order to avoid abrupt changes in color that can create false contours in homogeneous image regions. We also require its derivative to be to be continuous in order to avoid abrupt changes in contrast that can create false contours in smooth image gradients. In effect, a continuously differentiable transformation refrains from introducing artificial discontinuities into the resulting color density. To satisfy these requirements, our histogram warping transformation relies on a piecewise rational quadratic interpolating spline^{41,42}:

$$T(x) = b_{k-1} + \frac{r_k t^2 + d_{k-1} (1-t)t}{r_k + (d_k + d_{k-1} - 2r_k)(1-t)t} (b_k - b_{k-1}) \quad \text{where } r_k = \frac{b_k - b_{k-1}}{a_k - a_{k-1}} \text{ and } t = \frac{x - a_{k-1}}{a_k - a_{k-1}} \text{ for } x \in [a_{k-1}, a_k]. \quad (1)$$

How a transformation stretches and compresses the histogram determines its impact on image contrast. There can be either noise and aliasing artifacts where the histogram is overly stretched or detail loss where it is overly compressed. Such defects manifest themselves as spurious holes and spikes in the resulting histogram. Since the dynamic range has a fixed span of colors, color shift makes contrast adjustment inevitable. Furthermore, raising the contrast in one region of the histogram necessitates lowering the contrast in another, and vice versa. The tradeoff is governed by the choice of transformation function. For more precise control over how the transformation stretches and compresses the color transformation, it is possible to use a piecewise rational cubic interpolating spline with adjustable shape parameters⁴².

5. COLOR MATCHING

Histogram specification seeks to reshape the source color distribution to take on the features of the target color distribution. For the source A and target B colors, we examine their probability densities f , cumulative distributions F , and quantiles F^{-1} . The intended effect of applying the transformation $y = T(x)$ to the color distribution $F_A(x)$ of the source image is for the outcome, $F_{T(A)}(y) = F_A(T^{-1}(y))$ with $f_{T(A)}(y) = f_A(T^{-1}(y)) \left(T'(T^{-1}(y)) \right)^{-1}$, to approximate the target color distribution $F_B(x)$. Exact histogram specification, using $S(x) = F_B^{-1}(F_A(x))$, precisely reproduces the target distribution $F_{S(A)}(y) = F_A(S^{-1}(y)) = F_A(F_A^{-1}(F_B(y))) = F_B(y)$. Its effect on contrast $S'(x) = f_A(x)/f_B(S(x))$ is to stretch the more probable regions and compress the less probable regions of the source histogram as compared with the corresponding regions of the target histogram. Yet, for many practical applications, exact histogram specification is a needlessly rigid directive, causing unwarranted contrast distortion of the source image. Typically, the goal is to retain the color scheme rather than necessarily reproduce the color histogram, as the color balance is reflected in the overall shape of the distribution while minor discrepancies between histograms can often be below the threshold of visual perception.

An approximate transformation can prove more effective by ensuring that color and contrast change in a gradual and controlled manner. Therefore, our algorithm only enforces the correspondence $F_{T(A)}(b_k) = F_B(b_k)$ of the resulting distribution to the target distribution at a few representative key colors $b_k = T(a_k) = F_B^{-1}(F_A(a_k))$ selected to capture the relevant features of the color distributions. It maps a set q_k of quantiles $a_k = F_A^{-1}(q_k)$ of the source distribution to the corresponding quantiles $b_k = F_B^{-1}(q_k)$ of the target distribution. Thus the source color range $[a_0, a_K] = [F_A^{-1}(0), F_A^{-1}(1)]$ is transformed to match the target color range $[b_0, b_K] = [F_B^{-1}(0), F_B^{-1}(1)]$. The total number $K + 1$ of matching quantiles used in the color mapping determines the accuracy of the transformation. As it increases, our transformation $T(x)$ converges to exact histogram specification $S(x)$, since $|T(x) - S(x)| \leq b_k - b_{k-1}$ for $x \in [a_{k-1}, a_k]$. The key colors can be chosen in various ways. Evenly spaced quantiles can be used to convey the overall shape of the target color distribution, while its modes can be used to discern its predominant colors. The key colors can also be designated by a color palette.

6. COLOR CONTRAST

A key color's contrast adjustment controls the local deformation of its region of the source color histogram. The contrast adjustments $d_k = T'(a_k)$ serve both to improve the accuracy of histogram specification and to moderate the distortion of histogram transformation. To combat artifacts, our implementation places limits on contrast adjustments $\lambda^{-1} \leq d_k \leq \lambda$ (our experiments set $\lambda = 5$). Observe that preserving the contrast of the key colors, by setting every $d_k = 1$, merely pushes the brunt of the distortion to the intermediate colors between them. But performing exact histogram specification at the key colors, by setting contrast adjustments $d_k = f_A(a_k)/f_B(b_k)$ so that $f_{T(A)}(b_k) = f_A(a_k) \left(T'(a_k) \right)^{-1} = f_B(b_k)$, also risks excessive distortion of the intermediate colors. We avoid the direct use of the ratio between the source and target histograms since the raw histograms that are derived from image data can exhibit substantial short range variations whereas the impact of the contrast adjustments extends from the key colors to their intermediate colors. In this context, nearest neighbor estimation⁴³ offers a simple, adaptive way to more smoothly approximate probability densities. We use it to relate the contrast adjustments $d_k = T'(a_k)$ to the slopes of the color mapping $r_k = (b_k - b_{k-1}) / (a_k - a_{k-1})$.

For a data set of size n , given the distance $\Delta(x|\eta) = |x - \xi(x|\eta)|$ from x to its η nearest neighbor $\xi(x|\eta)$, the nearest neighbor density estimator is $f^*(x|\eta) = \eta / (2n\Delta(x|\eta)) \approx f(x)$. Each span $[x - \Delta(x|\eta), x + \Delta(x|\eta)]$ holds the same probability mass $\rho = \eta/n \approx 2|F(x) - F(\xi(x|\eta))|$. We write the left nearest neighbor density estimator as $f^L(x|\rho) = \rho / \Delta^L(x|\rho) \approx f(x)$, where $\Delta^L(x|\rho) = x - \xi^L(x|\rho)$ is the distance from x to its left neighbor $\xi^L(x|\rho) = F^{-1}(F(x) - \rho)$ such that the span $[\xi^L(x|\rho), x]$ has probability mass ρ . Since the key colors are defined by matching quantiles, their spans are equiprobable $\rho_k = F_B(b_k) - F_B(b_{k-1}) = F_A(a_k) - F_A(a_{k-1})$ so that $\xi_B^L(b_k|\rho_k) = b_{k-1}$ and $\xi_A^L(a_k|\rho_k) = a_{k-1}$. Hence, as $f_A^L(a_k|\rho_k) / f_B^L(b_k|\rho_k) = \Delta_B^L(b_k|\rho_k) / \Delta_A^L(a_k|\rho_k) = (b_k - b_{k-1}) / (a_k - a_{k-1}) = r_k$, the left slope of the color mapping provides a left sided estimate of the contrast adjustment. For the right nearest neighbor density estimator $f^R(x|\rho) = \rho / \Delta^R(x|\rho) \approx f(x)$ with $\Delta^R(x|\rho) = \xi^R(x|\rho) - x$ and $\xi^R(x|\rho) = F^{-1}(F(x) + \rho)$, by similar argument $f_A^R(a_k|\rho_{k+1}) / f_B^R(b_k|\rho_{k+1}) = r_{k+1}$, the right slope of the color mapping provides a right sided estimate of the contrast adjustment. To better regulate how the transformation affects intermediate colors, we use the local median colors, α_k and β_k , to define the local slopes that act as the left and right contrast adjustment estimates, d_k^L and d_k^R :

$$\begin{aligned} \alpha_k &= \xi_A^L(a_k | \frac{1}{2}\rho_k) = \xi_A^R(a_{k-1} | \frac{1}{2}\rho_k) = F_A^{-1}(\frac{1}{2}F_A(a_{k-1}) + \frac{1}{2}F_A(a_k)) \\ \beta_k &= \xi_B^L(b_k | \frac{1}{2}\rho_k) = \xi_B^R(b_{k-1} | \frac{1}{2}\rho_k) = F_B^{-1}(\frac{1}{2}F_B(b_{k-1}) + \frac{1}{2}F_B(b_k)) \end{aligned} \quad (2)$$

$$d_k^L = \frac{f_A^L(a_k | \frac{1}{2}\rho_k)}{f_B^L(b_k | \frac{1}{2}\rho_k)} = \frac{b_k - \beta_k}{a_k - \alpha_k} \quad \text{and} \quad d_k^R = \frac{f_A^R(a_k | \frac{1}{2}\rho_{k+1})}{f_B^R(b_k | \frac{1}{2}\rho_{k+1})} = \frac{\beta_{k+1} - b_k}{\alpha_{k+1} - a_k}. \quad (3)$$

For the slope of the transformation to reflect a compromise between its left and right estimates, it is reasonable to apply the geometric mean $d_k = (d_k^L d_k^R)^{1/2}$, which has the advantage of preserving flat slopes, so $d_k = 0$ when $d_k^L = 0$ or $d_k^R = 0$, as well as canceling inverse slopes, so $d_k = 1$ when $d_k^L = (d_k^R)^{-1}$. However, unless the key colors are defined by evenly spaced quantiles, the left and right contrast adjustment estimates generally do not carry equal weight $\frac{1}{2}\rho_k \neq \frac{1}{2}\rho_{k+1}$. Therefore, our contrast adjustments are defined using the geometric mean of the slopes of the local median colors weighted by the relative probability mass of their spans:

$$d_0 = \frac{\beta_1 - b_0}{\alpha_1 - a_0}, \quad \dots, \quad d_k = \left(\frac{b_k - \beta_k}{a_k - \alpha_k} \right)^{\frac{F_A(a_k) - F_A(a_{k-1})}{F_A(a_{k+1}) - F_A(a_{k-1})}} \left(\frac{\beta_{k+1} - b_k}{\alpha_{k+1} - a_k} \right)^{\frac{F_A(a_{k+1}) - F_A(a_k)}{F_A(a_{k+1}) - F_A(a_{k-1})}}, \quad \dots, \quad d_K = \frac{b_K - \beta_K}{a_K - \alpha_K}. \quad (4)$$

In our experiments, this contrast adjustment formula has proven effective at moderating the distortion of the histogram transformation. In fact, it is related to a derivative formula proposed for monotonic splines⁴². However, outside of the key colors, the result of this transformation is not guaranteed to coincide with exact histogram specification. In situations where accuracy needs to take precedence over distortion, it is possible to ensure that an intermediate color $\gamma_k = F_A^{-1}(F_B(\frac{1}{2}b_k + \frac{1}{2}b_{k-1}))$ of the source color distribution is correctly mapped to the corresponding local midpoint color $\frac{1}{2}b_k + \frac{1}{2}b_{k-1} = T(\gamma_k)$ of the target color distribution by appropriately specifying the contrast adjustments:

$$d_k = d_{k-1} + \phi_k = d_0 + \sum_{j=1}^k \phi_j \quad \text{where} \quad \phi_k = \frac{(2\gamma_k - a_k - a_{k-1})(b_k - b_{k-1})}{(a_k - \gamma_k)(\gamma_k - a_{k-1})}. \quad (5)$$

This sequence of contrast adjustments can be initialized in various ways. It is always possible to select d_0 to ensure a monotonic transformation, with every $d_k \geq 0$. If feasible, the contrast of one of the key colors can also be preserved, so that $\text{median}(d_k) = 1$. Yet, when bounds are imposed on all the contrast adjustments $\lambda^{-1} \leq d_k \leq \lambda$, they may no longer uphold the mapping of the local midpoint colors $\frac{1}{2}b_k + \frac{1}{2}b_{k-1} = T(\gamma_k)$, and thus we decided against using this approach.

7. COLOR APPLICATIONS

7.1. Histogram specification by a color distribution function

Our histogram warping technique can be used to alter the color histogram of an image to resemble a specified probability distribution. In particular, by designating a uniform target distribution, our method offers a quick way to equalize the color histograms of the color axes (Figure 2). In applications such as remote sensing, color histogram equalization serves to make image features easier to detect, interpret and compare by standardizing the contrast and colorfulness of the image. Our approach is more flexible than exact color equalization, which dictates that each possible color value be assigned to an equal number of image pixels. By adjusting the accuracy of our transformation, it is easy to control the tradeoff between enforcing a flat histogram and retaining the features of the image's original color distribution.




Figure 2. Color histogram equalization (from left to right): the original image, our median equalization ($K=2$), our quartile equalization ($K=4$), and exact histogram equalization. This example uses the standard RGB color space. In remote sensing applications, colors are often artificially exaggerated to improve the contrast between image features.

Table 1. Robust descriptive statistics, based on quantile analysis, evaluated for the standard Gaussian and uniform distributions.

Order	K	Property	Statistical Measure	Quantile Formula	Uniform	Gaussian
0	$K = 1$	Range	Minimum and Maximum	$[F^{-1}(0), F^{-1}(1)]$	$[0, 1]$	$(-\infty, \infty)$
1	$K = 2$	Location	Median	$F^{-1}(\frac{1}{2})$	0.5	0
2	$K = 4$	Dispersion	Interquartile Range	$F^{-1}(\frac{3}{4}) - F^{-1}(\frac{1}{4})$	0.5	1.34898
3	$K = 4$	Skewness	Bowley Skewness	$\frac{F^{-1}(\frac{3}{4}) - 2F^{-1}(\frac{2}{4}) + F^{-1}(\frac{1}{4})}{F^{-1}(\frac{3}{4}) - F^{-1}(\frac{1}{4})}$	0	0
4	$K = 8$	Kurtosis	Moors Kurtosis	$\frac{F^{-1}(\frac{7}{8}) - F^{-1}(\frac{5}{8}) + F^{-1}(\frac{3}{8}) - F^{-1}(\frac{1}{8})}{F^{-1}(\frac{6}{8}) - F^{-1}(\frac{2}{8})}$	1	1.23310

For approximate color histogram specification, our transformation $b_k = T(a_k)$ decides on the key colors where the resulting distribution will coincide with the target distribution. It establishes a mapping between sets of evenly spaced, matching quantiles of the source $a_k = F_A^{-1}(\frac{k}{K})$ and target $b_k = F_B^{-1}(\frac{k}{K})$ color distributions. Quantiles readily offer a comprehensive characterization (Table 1) of a distribution’s shape that is more resilient to outliers and noise than the traditionally used central moments. The number of matched quantiles $0 \leq k \leq K$ sets the accuracy of the transformation. As long as the median is not too far away from the mode, only a few evenly spaced quantiles suffice to describe a typical unimodal distribution. For instance, a rough but robust $K = 2$ transformation maps the minimum, the median, and the maximum color values of the source image to the corresponding quantiles of the target color distribution. A more accurate $K = 4$ transformation uses the quartiles to model color contrast, thus resulting in a color distribution with the same range, location, dispersion, and skewness as the target color distribution. Higher order features, such as kurtosis captured with a $K = 8$ transformation, often tend to have a less apparent influence on the resulting image²⁰.

7.2. Histogram specification by an example image

Our histogram warping technique can perform color transfer (Figure 1 and 3), transforming the colors of a source image to resemble the colors of a target image. This type of method is useful for various color correction tasks: preprocessing for image comparison, rendering artificial scenes with natural colors, altering the illumination of depicted scenes, synthesizing textures and photomosaics, removing shadows from aerial images, eliminating striping artifacts from multisensor images, reducing flicker between video frames, as well as calibrating 2D image slices for 3D volumetric reconstruction. Typically, the challenge is to apply the minimal color distortion necessary to the source image so that the resulting image reflects the dominant color characteristics of the target image. Previous approaches commonly relied on exact histogram specification or a linear histogram transformation in a predetermined color space. Our technique substantially improves the reliability of the transformation by using a color space adapted to the statistics of the source image and a color feature set adapted to the statistics of the target image. However, global histogram specification can never be entirely foolproof since by looking at the distribution of colors only in the dynamic range and not in the spatial domain, it is impossible to tell how changing the relative proportions of colors will change the semantics of the picture. Color histogram specification that adapts intelligently to local image structure is a worthy subject for future research.

To capture dominant colors of the target color distribution, we use the modes of its histogram as the key colors that guide our transformation and determine its accuracy. We establish a mapping $b_k = T(a_k)$ between the corresponding set of quantiles $a_k = F_A^{-1}(F_B(b_k))$ of the source distribution and the modes b_k of the target color distribution. When the target color distribution is derived from an image, the raw histogram of the color values can exhibit spurious variations due to quantization effects. Mode detection for empirical distributions requires carefully calibrated smoothing of the data. Our implementation relies on kernel density estimation⁴³, a robust nonparametric statistical method. We approximate the target density $f_B(y)$ as a sum of Gaussians centered on each of the n_B color values y_i of the target image:

$$f_B(y) = \frac{1}{n_B} \sum_{i=1}^{n_B} N(y|y_i, h_B^2) \quad \text{and} \quad f'_B(y) = \frac{1}{n_B h_B^2} \sum_{i=1}^{n_B} (y_i - y) N(y|y_i, h_B^2) \quad \text{with} \quad N(x|\mu, \sigma^2) = \frac{e^{-(x-\mu)^2/2\sigma^2}}{\sqrt{2\pi}\sigma}. \quad (6)$$



Figure 3. Color transfer of scene ambience using our histogram warping transformation in $La''b''$ color space. The colors of the original images (top) are exchanged (bottom).

The bandwidth h_B sets the resolution of the analysis, the minimal scale for detectable histogram features. There are many approaches for automatic scale selection in density estimation⁴⁴. Rather than a truncated kernel with finite support, we selected the Gaussian kernel to ensure that the bandwidth acts as a proper smoothing parameter. In effect, increasing bandwidth causes the density estimate to become smoother and its modes fewer, which increases the estimate's bias while decreasing its variance. Minimizing the asymptotic mean integrated squared error of the derivative estimate results in an optimized bandwidth $h_B = O(n_B^{-1/7})$ for determining the sensitivity of the mode detection process to fluctuations in the data. A conservative bandwidth ensures no more modes are detected than would be observed in an asymptotically optimal estimate. Using this maximal smoothing principle⁴⁵, we calculate the largest degree of smoothing compatible with the interquartile range, $h_B = 0.7816774 \left(F_B^{-1}(\frac{3}{4}) - F_B^{-1}(\frac{1}{4}) \right) \left(n_B^{-1/7} \right)$. We define a legitimate mode b_k , spanning a minimal probability mass 2δ (our experiments set $\delta = 2\%$), to be a critical point $f'_B(b_k) = 0$ such that every interval $[b_k^L, b_k^R]$, where $b_k \in [b_k^L, b_k^R]$ and $\delta \geq F_B(b_k^R) - F_B(b_k^L)$, satisfies $f'_B(b_k^L) \geq 0 \geq f'_B(b_k^R)$. The mode can be quickly located by repeated bisection: compute the midpoint $b_k \leftarrow \frac{1}{2}b_k^L + \frac{1}{2}b_k^R$ and move the left limit $b_k^L \leftarrow b_k$ when $f'_B(b_k) > 0$ or move the right limit $b_k^R \leftarrow b_k$ when $f'_B(b_k) < 0$. To avoid detecting ripples on an inflection or a plateau, it is better to use a method that validates the probability mass of a detected mode. Hence, we use the mean shift procedure⁴⁶ which converges to the mode by monotonically ascending its gradient in steps adapted to the density estimate:

$$b_k \leftarrow \left(\sum_{i=1}^{n_B} y_i N(b_k | y_i, h_B^2) \right) / \left(\sum_{i=1}^{n_B} N(b_k | y_i, h_B^2) \right). \quad (7)$$



Figure 4. Recoloring of image colors by applying a color palette using our histogram warping transformation in $L^a b^b$ color space. The color palette designates which colors of the original image (top right) need to be emphasized.

7.3. Histogram specification by a color palette

Our histogram warping transformation can render an image according to a specified palette of colors (Figure 4). This method can serve as an interactive tool for color emphasis. In graphic design, it is desirable for illustrations to share a unified color scheme. A color scheme can be selected to convey a particular mood or theme. It can be used underline desirable aspects of an image by making more pronounced those image colors which accord with the given palette. We selectively apply a color palette to an image, emphasizing those palette colors that are already present in the image. We use the palette colors to define a target color distribution that is adapted to the source image. For each color axis, the target color distribution is expressed as an equally weighted sum of Gaussians centered on the palette colors. For each palette color, we need to estimate the desired spread of its component in the sum. When a palette color is prominent in the source image, it is further emphasized by concentrating the spread of its influence. On the other hand, when a palette color is absent from the source image, it will attract less attention if the spread of its influence is increased.

Treating the palette colors as the key colors $b_k = T(a_k)$, we map the corresponding set of quantiles $a_k = F_A^{-1}(F_B(b_k))$ of the source distribution to the target palette's colors b_k . For the target color density $f_B(y) = \frac{1}{K-1} \sum_{k=1}^{K-1} N(y | b_k, \sigma_k^2)$, its cumulative distribution is truncated to fit the color range and its quantiles are evaluated numerically. We use expectation maximization⁴⁷ to estimate the parameters σ_k^2 of the Gaussian components so that the target color distribution accords with the n_A color values x_i subsampled from the source image. At each iteration, the parameters σ_k^2 are updated to maximize the likelihood of observing the source color values given an estimate of the target color components:

$$\sigma_k^2 \leftarrow \left(\sum_{i=1}^{n_A} (x_i - b_k)^2 p_k(x_i) \right) / \left(\sum_{i=1}^{n_A} p_k(x_i) \right) \quad \text{for} \quad p_k(x_i) = \left(N(x_i | b_k, \sigma_k^2) \right) / \left(\sum_{j=1}^{K-1} N(x_i | b_j, \sigma_j^2) \right). \quad (8)$$

8. CONCLUSION

Our work presents original contributions to the color space, color transformation, and color features used in histogram specification. Our image dependent La^*b^* color space relies on independent component analysis to make the algorithm more reliable by being less prone to hue artifacts. Our histogram warping transformation applies continuously differentiable, monotonic splines to make the algorithm easier to control through the separate specification of color shifts and contrast adjustments for a set of key colors. Our color features use mode detection by kernel density estimation to make the algorithm more adaptable to the structure of color distribution it seeks to reproduce. For the restricted case of unimodal distributions, color features based on evenly spaced quantiles can offer a robust description of distribution shape akin to traditional statistical moments. We are now exploring the practical applications for our histogram warping framework, and the initial results appear encouraging. So far we have investigated color specification by a probability distribution, color transfer between images, and color emphasis according to a color palette, in addition to our earlier experiments^{2,3} in interactive and automatic contrast enhancement. In future work, we wish to apply histogram warping to adaptive contrast enhancement as well as interactive color correction.

REFERENCES

1. R. C. Gonzalez and R. E. Woods, *Digital Image Processing*, 2 ed, Prentice Hall, Upper Saddle River, 2002.
2. M. Grundland and N. A. Dodgson, "Interactive Contrast Enhancement by Histogram Warping," in *International Conference on Computer Vision and Graphics*, 2004. To appear.
3. M. Grundland and N. A. Dodgson, "Automatic Contrast Enhancement by Histogram Warping," in *International Conference on Computer Vision and Graphics*, 2004. To appear.
4. R. A. Hummel, "Histogram Modification Techniques," *Computer Graphics & Image Processing*, **4**, 3, pp. 209-224, 1975.
5. R. C. Gonzalez and B. A. Fittes, "Gray-Level Transformations for Interactive Image Enhancement," *Mechanism & Machine Theory*, **12**, 1, pp. 111-122, 1977.
6. L. O'Gorman and L. S. Brotman, "Entropy-Constant Image Enhancement by Histogram Transformation," in *Applications of Digital Image Processing VIII*, Proceedings of SPIE, **575**, pp. 106-113, 1985.
7. E. L. Hall, "Almost Uniform Distributions for Computer Image Enhancement," *IEEE Transactions on Computers*, **23**, 2, pp. 207-208, 1974.
8. D. J. Heeger and J. R. Bergen, "Pyramid-Based Texture Analysis/Synthesis," in *Proceedings of SIGGRAPH*, Computer Graphics, pp. 229-238, 1995.
9. H. B. Bidasaria, "A Method for Almost Exact Histogram Matching for Two Digitized Images," *Computer Vision, Graphics, & Image Processing*, **34**, 1, pp. 93-98, 1986.
10. X. D. Yang, Q. Xiao, and H. Raafat, "Direct Mapping between Histograms: An Improved Interactive Image Enhancement Method," in *IEEE International Conference on Systems, Man, and Cybernetics*, **1**, pp. 243-247, 1991.
11. J. P. Rolland, V. Vo, B. Bloss, and C. K. Abbey, "Fast Algorithms for Histogram Matching: Application to Texture Synthesis," *Journal of Electronic Imaging*, **9**, 1, pp. 39-45, 2000.
12. S.-K. Chang and Y.-W. Wong, "Optimal Histogram Matching by Monotone Gray Level Transformation," *Communications of the ACM*, **21**, 10, pp. 835-840, 1978.
13. S.-K. Chang and Y.-W. Wong, "Ln Norm Optimal Histogram Matching and Application to Similarity Retrieval," *Computer Graphics & Image Processing*, **13**, 4, pp. 361-371, 1980.
14. M. Mese and P. P. Vaidyanathan, "Optimal Histogram Modification with Mse Metric," in *IEEE International Conference on Acoustics, Speech, and Signal Processing*, **3**, pp. 1665-1668, 2001.
15. I. J. Cox, S. Roy, and S. L. Hingorani, "Dynamic Histogram Warping of Image Pairs for Constant Image Brightness," in *Proceedings of the International Conference on Image Processing*, **2**, pp. 366-369, 1995.
16. J. Morovic, J. Shaw, and S. Pei-Li, "A Fast, Non-Iterative and Exact Histogram Matching Algorithm," *Pattern Recognition Letters*, **23**, 1-3, pp. 127-135, 2002.
17. J. Kautsky, N. K. Nichols, and D. L. B. Jupp, "Smoothed Histogram Modification for Image Processing," *Computer Vision, Graphics, & Image Processing*, **26**, 3, pp. 271-291, 1984.
18. A. R. Gillespie, A. B. Kahle, and R. E. Walker, "Color Enhancement of Highly Correlated Images. I. Decorrelation and Hsi Contrast Stretches," *Remote Sensing of Environment*, **20**, 3, pp. 209-235, 1986.

19. E. Reinhard, M. Adhikhmin, B. Gooch, and P. Shirley, "Color Transfer between Images," *IEEE Computer Graphics & Applications*, **21**, 5, pp. 34-41, 2001.
20. D. D. Thompson and R. C. Gonzalez, "Image Enhancement by Moment Specification," in *Proceedings of the 15th Southeastern Symposium on System Theory*, pp. 134-137, 1983.
21. X. D. Yang, "Adaptive Representation of Histogram Using Interval Tree," in *IEE International Conference on Image Processing and its Applications*, pp. 490-493, 1992.
22. K. Sang-Yeon, H. Dongil, C. Seung-Jong, and P. Jong-Seok, "Image Contrast Enhancement Based on the Piecewise-Linear Approximation of Cdf," *IEEE Transactions on Consumer Electronics*, **45**, 3, pp. 828-834, 1999.
23. F. Hilger and H. Ney, "Quantile Based Histogram Equalization for Noise Robust Speech Recognition," in *Proceedings of EUROSPEECH*, **2**, pp. 1135-1138, 2001.
24. J. C. Segura, C. Benitez, A. de la Torre, A. J. Rubio, and J. Ramirez, "Cepstral Domain Segmental Nonlinear Feature Transformations for Robust Speech Recognition," *IEEE Signal Processing Letters*, **11**, 5, pp. 517-520, 2004.
25. A. Raji, A. Thaibaoui, E. Petit, P. Bunel, and G. Mimoun, "A Gray-Level Transformation-Based Method for Image Enhancement," *Pattern Recognition Letters*, **19**, 13, pp. 1207-1212, 1998.
26. J. J. McCann, "Color Spaces for Color-Gamut Mapping," *Journal of Electronic Imaging*, **8**, 4, pp. 354-364, 1999.
27. G. Cui, M. R. Luo, B. Rigg, G. Roesler, and K. Witt, "Uniform Colour Spaces Based on the Din99 Colour-Difference Formula," *Color Research & Application*, **27**, 4, pp. 282-290, 2002.
28. P. E. Trahanias and A. N. Venetsanopoulos, "Color Image Enhancement through 3-D Histogram Equalization," in *Proceedings of 11th International Conference on Pattern Recognition*, **3**, pp. 545-548, 1992.
29. E. Pichon, M. Niethammer, and G. Sapiro, "Color Histogram Equalization through Mesh Deformation," in *Proceedings of the International Conference on Image Processing*, **2**, pp. 117-120, 2003.
30. L. Lucchese and S. K. Mitra, "A New Method for Color Image Equalization," in *Proceedings of the International Conference on Image Processing*, **1**, pp. 133-136, 2001.
31. I. Pitas and P. Kiniklis, "Multichannel Techniques in Color Image Enhancement and Modeling," *IEEE Transactions of Image Processing*, **5**, 1, pp. 168-171, 1996.
32. I. M. Bockstein, "Color Equalization Method and Its Application to Color Image Processing," *Journal of the Optical Society of America A*, **3**, 5, pp. 735-737, 1986.
33. A. R. Weeks, L. J. Sartor, and H. R. Myler, "Histogram Specification of 24-Bit Color Images in the Color Difference (C-Y) Color Space," in *Nonlinear Image Processing X*, Proceedings of SPIE, **3646**, pp. 319-329, 1999.
34. P. A. Mlsna and J. J. Rodriguez, "A Multivariate Contrast Enhancement Technique for Multispectral Images," *IEEE Transactions on Geoscience & Remote Sensing*, **33**, 1, pp. 212-216, 1995.
35. P. A. Mlsna, Z. Qiang, and J. J. Rodriguez, "3-D Histogram Modification of Color Images," in *Proceedings of the International Conference on Image Processing*, **3**, pp. 1015-1018, 1996.
36. J. Morovic and P. L. Sun, "Transforming 3d Colour Histograms of Images," in *Proceedings of 1rst European Conference on Color in Graphics, Imaging and Vision*, pp. 104-108, 2002.
37. I. T. Jolliffe, *Principal Component Analysis*, Springer-Verlag, New York, 1986.
38. A. Hyvärinen, J. Karhunen, and E. Oja, *Independent Component Analysis*, Wiley, New York, 2001.
39. J. Morovic and M. R. Luo, "The Fundamentals of Gamut Mapping: A Survey," *Journal of Imaging Science & Technology*, **45**, 3, pp. 283-290, 2001.
40. T. Mitsa and K. J. Parker, "Digital Halftoning Technique Using a Blue-Noise Mask," *Journal of the Optical Society of America A*, **9**, 11, pp. 1920-1929, 1992.
41. J. A. Gregory and R. Delbourgo, "Piecewise Rational Quadratic Interpolation to Monotonic Data," *IMA Journal of Numerical Analysis*, **2**, pp. 123-130, 1982.
42. M. Sarfraz, M. Al-Mulhem, and F. Ashraf, "Preserving Monotonic Shape of the Data Using Piecewise Rational Cubic Functions," *Computers & Graphics*, **21**, 1, pp. 5-14, 1997.
43. D. W. Scott, *Multivariate Density Estimation: Theory, Practice, and Visualization*, Wiley, New York, 1992.
44. L. Devroye, "Universal Smoothing Factor Selection in Density Estimation: Theory and Practice," *Test*, **6**, pp. 223-320, 1997.
45. G. R. Terrell, "The Maximal Smoothing Principle in Density Estimation," *Journal of the American Statistical Association*, **85**, 410, pp. 470-477, 1990.
46. Y. Cheng, "Mean Shift, Mode Seeking, and Clustering," *IEEE Transactions on Pattern Analysis & Machine Intelligence*, **17**, 8, pp. 790-799, 1995.
47. D. M. Titterton, A. F. M. Smith, and U. E. Makov, *Statistical Analysis of Finite Mixture Distributions*, Wiley, Chichester, 1985.