

Vision Models and Applications to Image and Video Processing

Christian J. Van Den Branden Lambrecht
(Editor)

Kluwer Academic Publisher
ISBN 0-7923-7422-3

August 1, 2001

Chapter 6

Karl-Heinz Bäuml
Xuemei Zhang
Brian Wandell

1. INTRODUCTION

Receptor catches “versus” color appearance: asymmetric color matching

The absorptions by the cone photoreceptors govern many of the important properties of color vision. Perhaps the most important consequence of the cone properties is that, even without reference to the spatial structure of the image, knowledge of the cone absorption properties can be used to predict when a pair of lights with different spectral power distributions will match in appearance. While the cone absorptions may tell us that two lights match, without knowing the spatial pattern of absorptions we do not know much about the appearance of the lights; do they appear light, or dark? Red or green?

The color appearance of a light depends on the viewing context in which the light is presented. There exist a number of dramatic demonstrations in the literature about how the color appearance of lights may vary with viewing context (Wandell, 1995; Kaiser and Boynton, 1996). Among them are the effects of adaptation and spatial pattern, which will be treated in some detail below.

Models of how context influences color appearance must be based on measurements that quantify these effects. The method of asymmetric color matching is a simple and effective method for quantifying the effect of context on color appearance (Wyszecki and Stiles, 1982). In this method an observer views some test light in a standard context and a matching field in a different, test context. The observer adjusts the matching field in the test context so that it has the same hue, saturation, and brightness, i. e. the same color appearance, as the test light in the standard context. Such matches are collected for a number of different

test lights and a number of context changes. Such measurements are the basic data used to investigate how changes in the viewing context transform the color appearance of lights.

The data from asymmetric matches can be used to examine the mapping between equivalent cone absorptions in the standard context and the test context. The overall form of the mapping between these two contexts is useful for predicting equivalent appearance across contexts. Also, as we shall see, the functional form may be instructive concerning the neural mechanisms that are responsible for governing color appearance. Once the functional form of the mapping is specified, it is also important to measure how the parameters of the transformation vary with the viewing context. Knowledge about the type of transformation and the dependence of the transformation's parameters on viewing context can provide us with very general tools to predict how color appearance varies with viewing context.

2. ADAPTATION

Effects of adaptation and the von Kries principle

One of the most striking effects of viewing context on the color appearance of a test light is adaptation to a colored uniform background field. The presence of a colored background can shift the color appearance of a neutral light to nearly any of a wide range of hues. The test light may look greenish for a reddish background, yellowish for a bluish background, or dark for a bright background. The method of asymmetric color matching allows us to quantify the effects of background color. Subjects view a test light presented against a standard background and then adjust the appearance of a matching field presented against a test background, so that test light and matching field have the same color appearance. If this procedure is repeated for a number of different test lights and a number of different background colors, the resulting data can be used to study how this adjustment process is organized.

A classic hypothesis, due to von Kries (1905), holds that the visual system adjusts to changes in the color of a background field by scaling the cone signals by an amount that depends on the photons absorbed from the background field. The test light viewed against the standard background and the matching field viewed against the test background then appear to match when the scaled cone signals are equal. Figure 6.1 shows the basic architecture of the model.

Over the years the von Kries hypothesis has been tested in several studies (Burnham et al., 1952; Hurvich and Jameson, 1958; Wassef, 1959; Walraven, 1976; Shevell, 1978; Werner and Walraven, 1982). The results

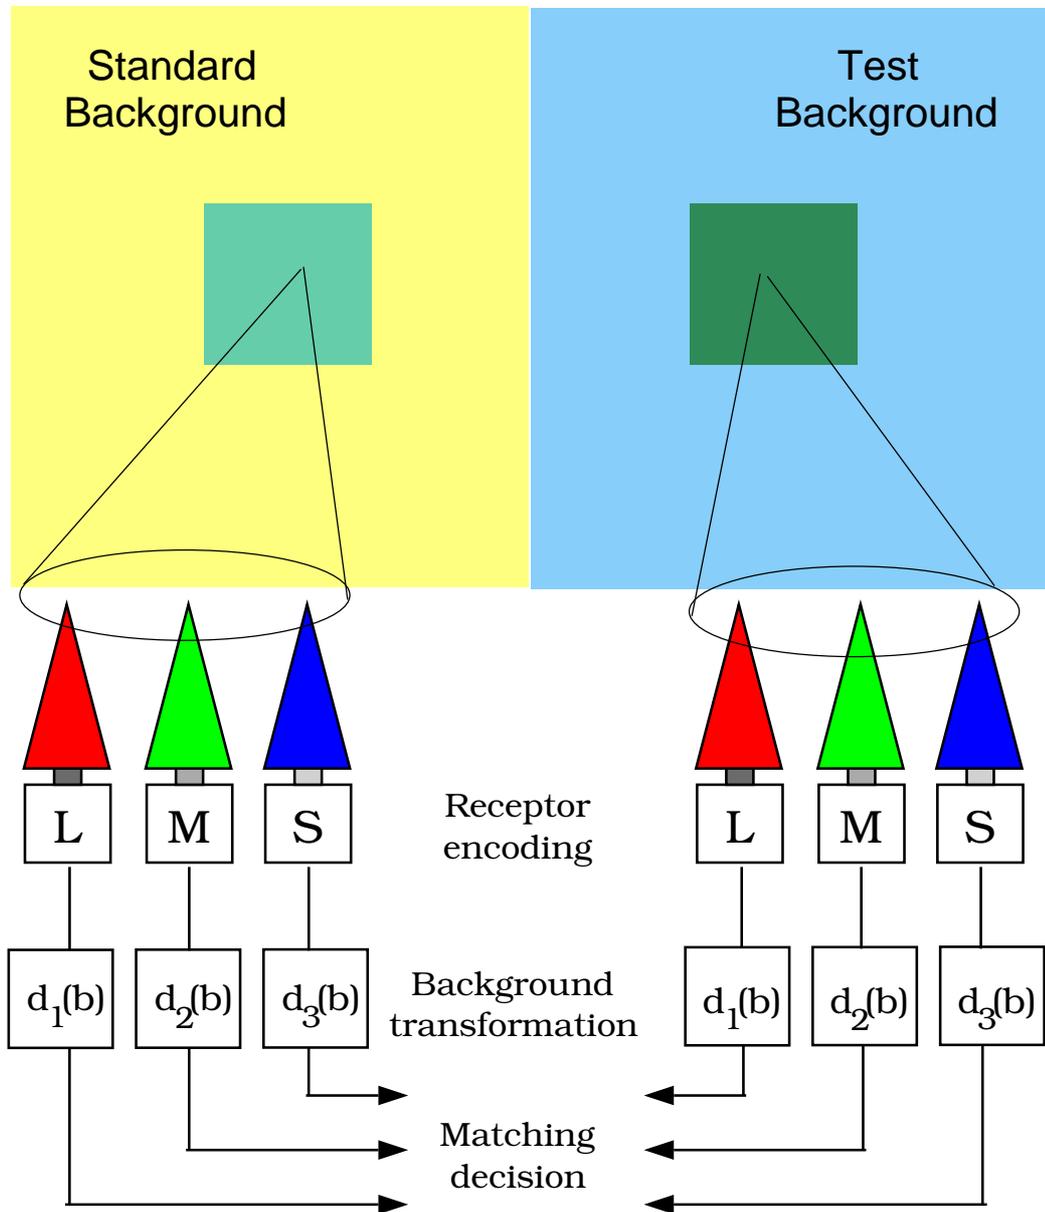


Figure 6.1 The von Kries principle. Subjects set color appearance matches between a test light presented against a standard background and a matching box presented against a test background. The mean rate of cone absorptions caused by the test light and the matching box are each scaled by an amount that depends on the color of the respective background (b). The test light and matching box match when the scaled representations are equal.

from these studies have been somewhat mixed. But, with advances in our understanding of the spectral absorption properties of the cones as well as simpler imaging systems, more direct and extensive tests have been achieved. There is now fairly convincing evidence that the von Kries principle is a good first-order account of how background color affects the color appearance of lights.

One of the most extensive tests of the von Kries principle comes from a study by Chichilnisky and Wandell (1995) who collected asymmetric color matches for a wide range of different test lights and background colors. Chichilnisky and Wandell tested the von Kries hypothesis in two steps. First, they considered the general form of the mapping between the cone contrasts measured in the two different viewing contexts. The von Kries hypothesis predicts that the transformation describing the asymmetric color matches should have the properties of a linear system. Consistent with the linear property of homogeneity, Chichilnisky and Wandell found that multiplying the cone absorptions of the test light resulted in observers scaling the cone absorptions of the matching field by an equal amount. Also, consistent with superposition they found that the match to the superposition of two test lights was roughly identical to the sum of the matches to the two individual test lights.¹

After accepting linearity as a first approximation, Chichilnisky and Wandell then tested whether the subjects' matches were in agreement with the von Kries principle. They found that the description of their data did not deteriorate substantially when matches were predicted by the von Kries model, in which the linear transformation is only a scaling of the cone absorptions, compared to the more general linear model. They also found that the assumption of a scaling of the cone signals provided a better account of the asymmetric color matches than did the assumption of a scaling of three opponent-colors systems.

These results generalize to viewing contexts in which comparisons are made between a collection of surfaces viewed under two different illuminants (Brainard and Wandell, 1992; Arend, 1993; Bäuml, 1995, 1999, Brainard et al., 1997, Brainard, 1998). This more complex visual display is a closer model to natural viewing. In this circumstance, color appearance of cone absorptions is strongly influenced by changes in the illumination. Once again, these adjustments can be quantified using asymmetric color matching: Subjects set a matching field under the test illuminant so that it has the same color appearance as a test light under the standard illuminant. Matches are set for a number of different test lights and test illuminants. The asymmetric matches in this type of situation are in broad agreement with those found when the color of a uniform background field is changed. The matches are well described

by linear transformations and are roughly consistent with the von Kries principle. Also, the von Kries principle provides a substantially better account of the matches than does the assumption of a scaling in an opponent-colors space (Bäumel, 1995, 1999).

The main limitation concerning these asymmetric color matches concerns differences between the behavior of increments and decrements. In making asymmetric matches, a clear principle is that increments in one context are never matched by decrements seen in a second context. Mausfeld and Niederee (1993) suggested that the analysis of color transformations should be broken up into sub-regions corresponding to the incremental and decremental signals from the different sensor classes. Chichilnisky and Wandell (1995) overlooked this effect in their original matching paper (see Mausfeld and Niederee, 1993; Chichilnisky and Wandell, 1996). In follow-up studies, they discovered several interesting differences between the way cone signals are scaled for incremental and decremental test lights (Chichilnisky and Wandell, 1996, 1999). They show that some improvements in the von Kries scaling predictions can be obtained by separating the scale factors for incremental and decremental cone signals (see also De Valois et al., 1997). Such differences between increments and decrements are not restricted to displays with uniform background fields but are present in displays that consist of a collection of illuminated surfaces as well (?).

In summary, effects of adaptation can be described by means of a model in which the cone signals are scaled as a function of changes in the viewing context, changes in background color or illumination. Scaling in cone space accounts for the data better than, for instance, the assumption of a scaling in the opponent-colors space.

In the next section we investigate the role of spatial pattern on the color appearance of lights. As we will see, in this case the opponent-colors space provides a better explanation of the effects of viewing context than does the cone space.

3. PATTERN EFFECTS

Effects of spatial pattern and a pattern-color separable model

A large number of studies have shown that both contrast sensitivity and color appearance vary as a function of spatial pattern. Many of the experiments studying these effects have used simple periodic patterns, such as sinewaves or squarewaves, as experimental stimuli. These contrast patterns are presented on a uniform background field, and sensitivity or appearance is measured while systematically varying the spatial

frequency of the pattern (Knowles et al., 1949, Ingling Jr. and Scheibner, 1970; Granger and Heurtley, 1973; Mullen, 1985; Elsner et al., 1987). These studies all show that spatial frequency has strong effects on the appearance of yellow-blue patterns, less but still considerable effects on red-green patterns, and relatively moderate effects on achromatic patterns.

The question arises how we can describe this effect of spatial frequency theoretically. Can we account for it by assuming the existence of three color pathways – for instance, the three cone systems or three opponent-colors systems – whose values are simply scaled as a function of spatial frequency, or do we need more complex models with a multiplicity of color mechanisms to describe the effect? Poirson and Wandell (1993) have addressed this issue in detail. Using an asymmetric color matching procedure, they measured the color appearance of the bars in squarewaves patterns. Their stimuli included contrast patterns with various colors, contrasts, and spatial frequencies (1, 2, 4, and 8 cpd). To match the color appearance of the bars in the squarewave, subjects adjusted the appearance of a uniform matching box to match. Consistent with prior research, Poirson and Wandell found a considerable effect of spatial frequency. Qualitatively the effect was that the appearance of the bars became more and more desaturated as the spatial frequency of the patterns increased.

In order to account for these effects of spatial pattern Poirson and Wandell (1993) proposed a pattern-color separable model. This model assumes that, in the first stage, the mean rate of cone absorptions from the squarewave's bar and the uniform matching box undergo a linear transformation into an intermediate color representation. In the second stage, the intermediate color representation values are scaled by an amount that depends on the local spatial pattern. The squarewave bar and the uniform box appear to match when the final, scaled representations are equal. Figure 6.2 shows the basic architecture of the model. Just as the von Kries model, this model is a special case of a linear model.

For each spatial frequency Poirson and Wandell studied the extent to which the subjects' matches to patterns of varying contrast and color direction were in agreement with the properties of a linear system. Consistent with homogeneity they found that, as they scaled the color contrast of the pattern, subjects scaled the contrast of the matching box equally. Consistent with superposition they found that the match to the superposition of two bars was identical to the sum of the matches to the two individual bars. These properties held true to within a tol-

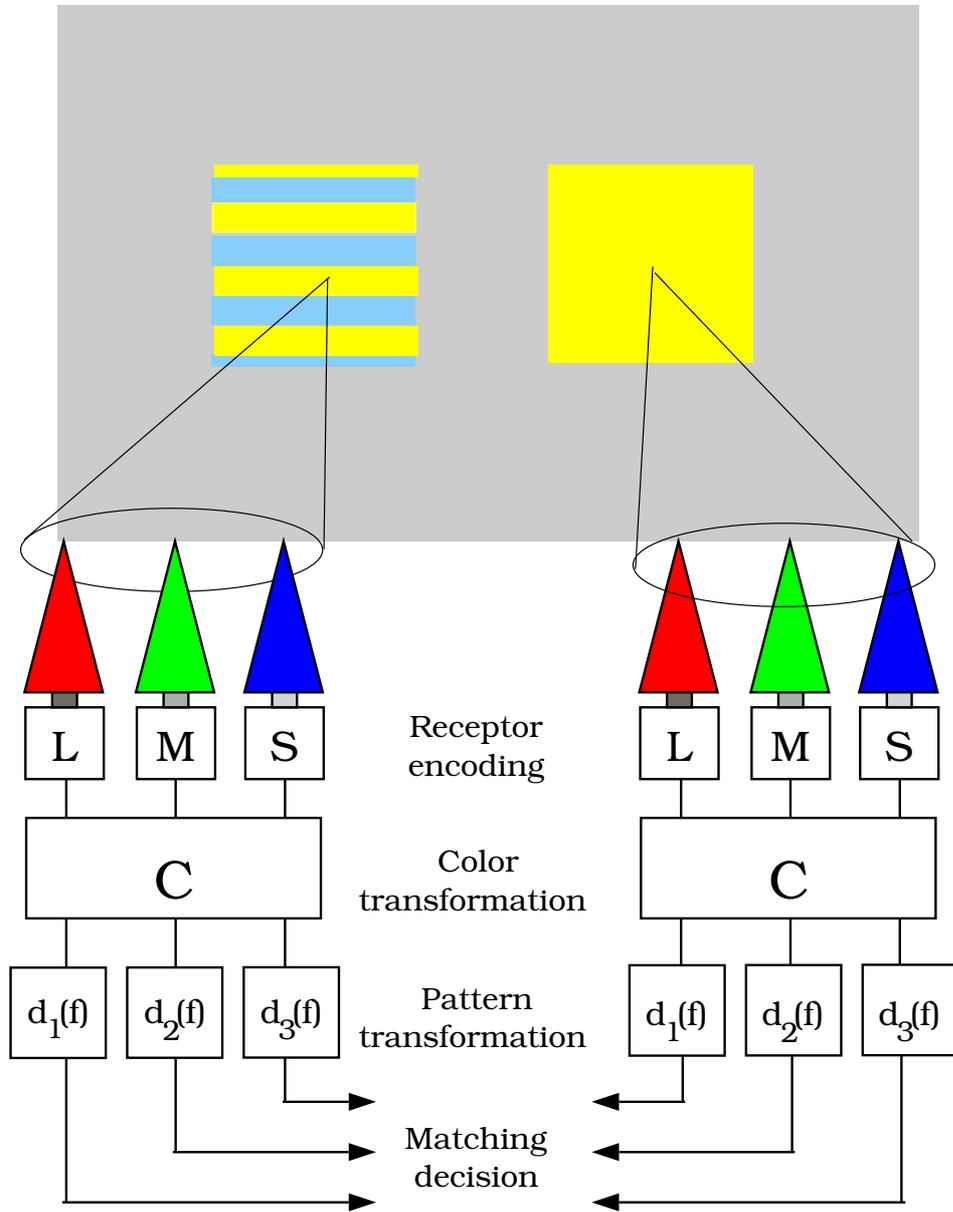


Figure 6.2 The pattern-color separable model. Subjects set color appearance matches between a bar in the squarewave gratings and a uniform field. In the first stage of the model, the mean rate of cone absorptions caused by the bar and the uniform matching box undergo a linear color transformation into an intermediate color representation. In the second stage, the intermediate color representation values are each scaled by an amount that depends on the local spatial pattern (f). The squarewave bar and the uniform box match when the final, scaled representations are equal.

erance of twice the precision of repeated matches. So, at least to first approximation, the matches could be described using linear models.

On the basis of this linearity finding, Poirson and Wandell next tested whether the subjects' matches could be described using the pattern-color separable model. They found that the description of their subjects' matches did not deteriorate substantially when using this model rather than a general linear model which does not impose separability, thus supporting the hypothesis of pattern-color separability.

From their subjects' matches Poirson and Wandell estimated the color and pattern responsivity functions of the three putative pathways of the intermediate color representation. They found that the effects of spatial pattern were mediated by two color-opponent pathways, one red-green and one yellow-blue pathway, and one all-positive black-white pathway (see Figure 6.3).² Consistent with prior research, they found that the black-white pathway was not much affected by spatial frequency, while the two opponent pathways showed a strong loss of contrast information as the frequency was increased (see Figure 6.4). Poirson and Wandell (1996) formulated a variant of the model to capture the effect of the pattern on sensitivity to lights as well. They found that the model describes the data to a degree comparable to the degree to which it describes the data from the color appearance experiments.

The spatial contrast sensitivity function to red-green has been measured in several different labs, and there is good agreement. Figure 6.5 shows measurements using three very different methods from three separate laboratories. The measurements from Poirson and Wandell (1993) are based on appearance, while those from Mullen (1985) and Sekiguchi et al. (1993) are based on contrast sensitivity measurements. It is satisfying to find that appearance and sensitivity measurements agree as well as they do.

Generalization of the model to mixture gratings

For the pattern-color separable model to be useful for capturing the effect of spatial pattern, it must be generalized to more complex patterns. Indeed, the appearance of more complex patterns must be predicted by knowing the effect of pattern on simple sinewave or squarewave patterns. This issue was addressed in a study by Bäuml and Wandell (1996).

Using an asymmetric color matching procedure Bäuml and Wandell (1996) measured the effect of spatial pattern for isolated squarewaves and sums of these squarewaves, superimposed in equal phase and orientation. Low spatial frequency patterns were used (1, 2, and 4 cpd). We applied the same test logic as in Poirson and Wandell (1993) by testing

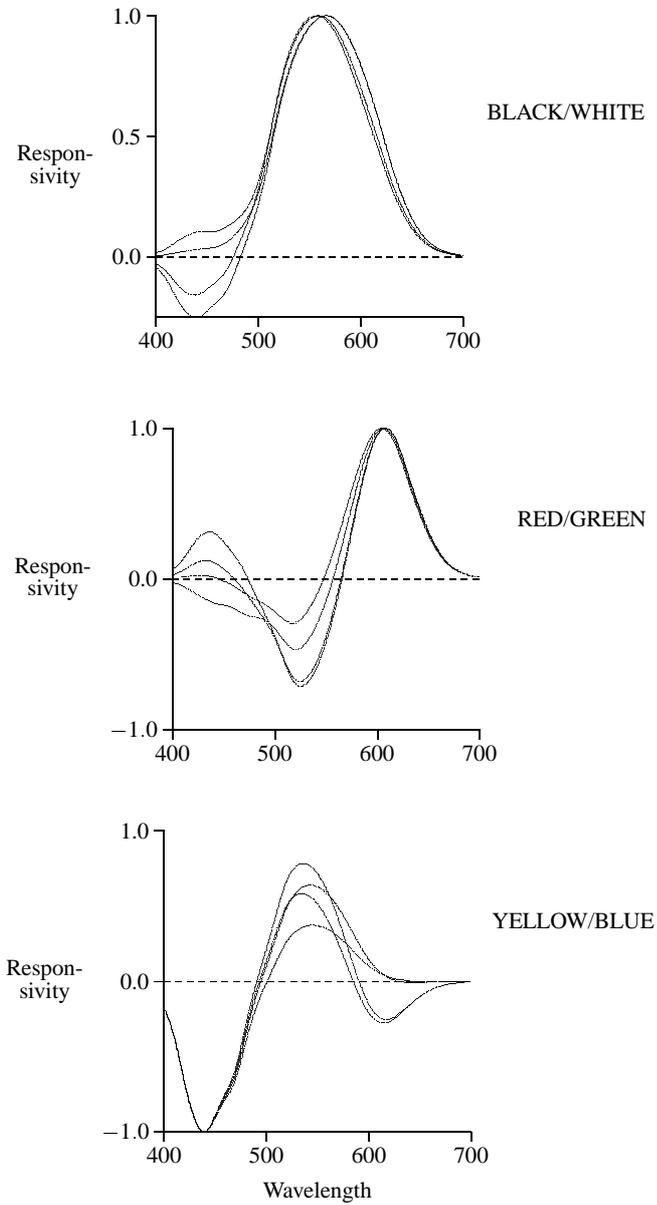


Figure 6.3 Color responsivity function estimates of the four observers from the studies of Poirson and Wandell (1993) and Bäuml and Wandell (1996). The estimates are based on fitting the linear pattern-color separable model to each observer's whole data set. The functions can be classified a black-white mechanism, a red-green mechanism, and a yellow-blue mechanism.

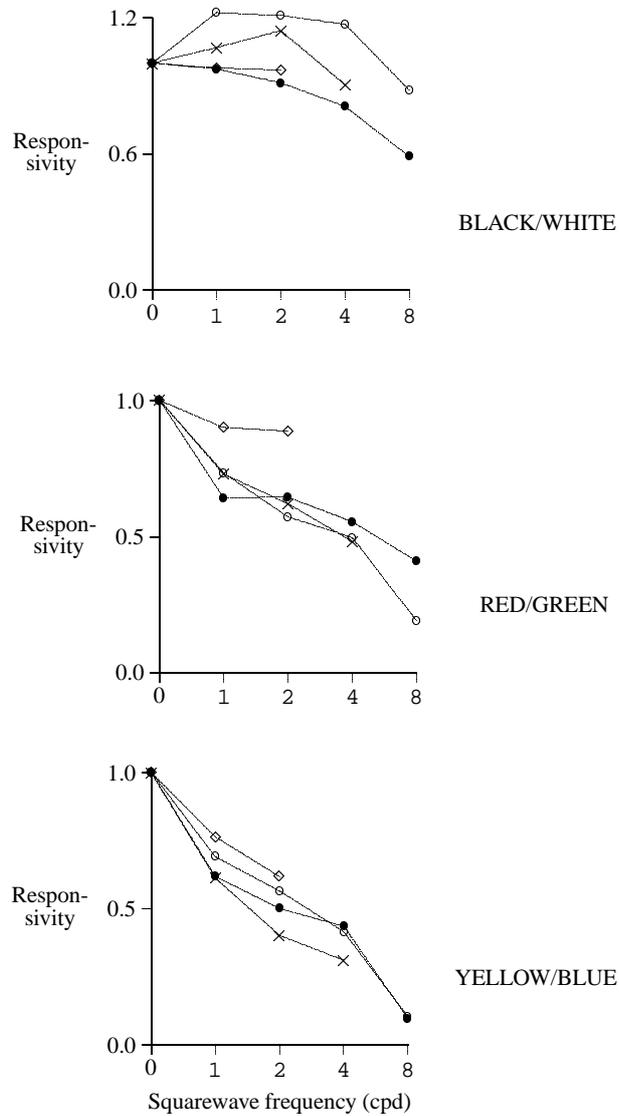


Figure 6.4 Pattern responsivity function estimates of the four observers from the studies of Poirson and Wandell (1993) and Bäuml and Wandell (1996). The estimates are based on fitting the linear pattern-color separable model to each observer's whole data set. The functions show the scalings for the putative black-white, the putative red-green, and the putative yellow-blue mechanism (compare also Figure 6.3).

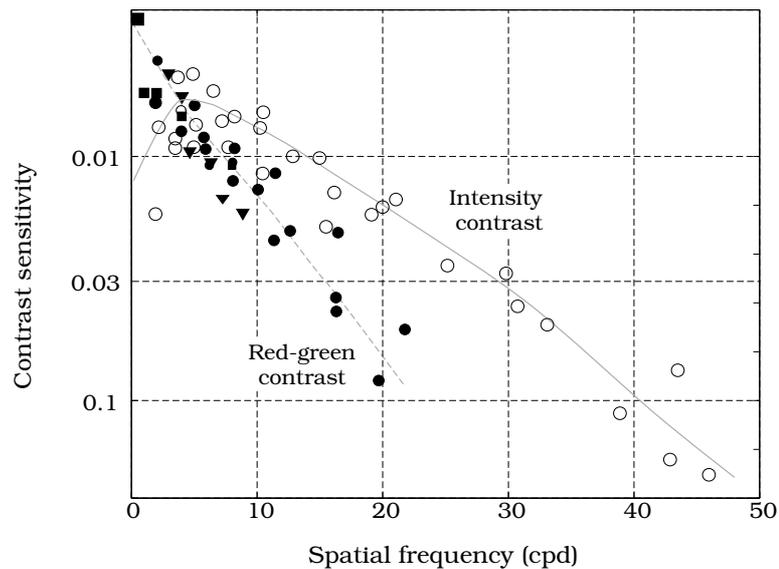


Figure 6.5 Comparison of red-green and intensity spatial contrast sensitivity functions. The filled symbols show contrast sensitivity measurements made using stimuli visible mainly to the red-green opponent-colors mechanisms. The data from Anderson et al. (1991; inverted triangles) and Sekiguchi et al. (1993; circles) show absolute contrast. The data from Poirson and Wandell (1993; squares) were obtained using a method that does not permit estimation of the absolute sensitivity. These data have been shifted vertically and should be compared only with respect to the fall-off in sensitivity. Measurements by Sekiguchi et al. and Anderson et al. compensated for the chromatic aberration of the eye. The Poirson and Wandell measurements did not, but at the relatively low spatial frequencies (below 7 cpd) where the three data sets overlap red-green chromatic aberration is not a large factor (Marimont and Wandell, 1994). The open circles show contrast sensitivity measurements from three observers to light-dark patterns (Sekiguchi et al., 1993). The shaded curves were drawn by hand to emphasize the higher spatial resolution of the light-dark variation measurements.

whether matches were consistent with homogeneity. In agreement with this principle, we found that if a bar within a pattern formed by the mixture of squarewaves of different spatial frequencies was matched by a uniform matching box and we scaled the contrast of the test pattern, then the contrast of the matching box was equally scaled (see Figure 6.6). Similarly, we tested whether the matches satisfied superposition. Suppose that one of the two bars comprising a squarewave was matched by a uniform matching box, and that a bar in a second squarewave of different spatial frequency was matched by a second uniform matching box. Now, when we formed the superposition of the squarewaves such that the two bars superimposed, we found that subjects matched the superposition of the bars by the sum of the matching boxes (see Figure 6.6). Homogeneity and superposition did not hold perfectly. Rather, as in the Poirson and Wandell study, we found the matches to be consistent with the properties of a linear system to within a tolerance of twice the precision of repeated matches.

Bäumel and Wandell also tested to what extent the description of the data deteriorated if pattern-color separability was imposed onto the model. We found that imposing separability onto the linear model did not substantially reduce the fit of the model. There was only a slight deterioration in fit comparable in degree to the one found in the Poirson and Wandell study. Moreover, there was good agreement between the color and pattern responsivity functions estimated for our two observers and those estimated for the two observers in the Poirson and Wandell study (compare Figure 6.3 and Figure 6.4). These results indicate that the pattern-color separable model can be generalized to more complex patterns and that it is indeed useful to capture the effect of pattern.

In summary, effects of spatial pattern can be well described by means of a model which assumes that the cone signals are transformed into three opponent-colors signals and that these opponent-colors signals are scaled as a function of the spatial frequency content in the image. Such a model can account for the color appearance of isolated squarewave patterns and the color appearance of mixture gratings that consist of the sum of isolated squarewaves. As opposed to the effects of adaptation, which are well described as a scaling in the cone space, the effects of spatial pattern thus are well described as a scaling in an opponent-colors space. Different color spaces are appropriate to account for different effects of viewing context.

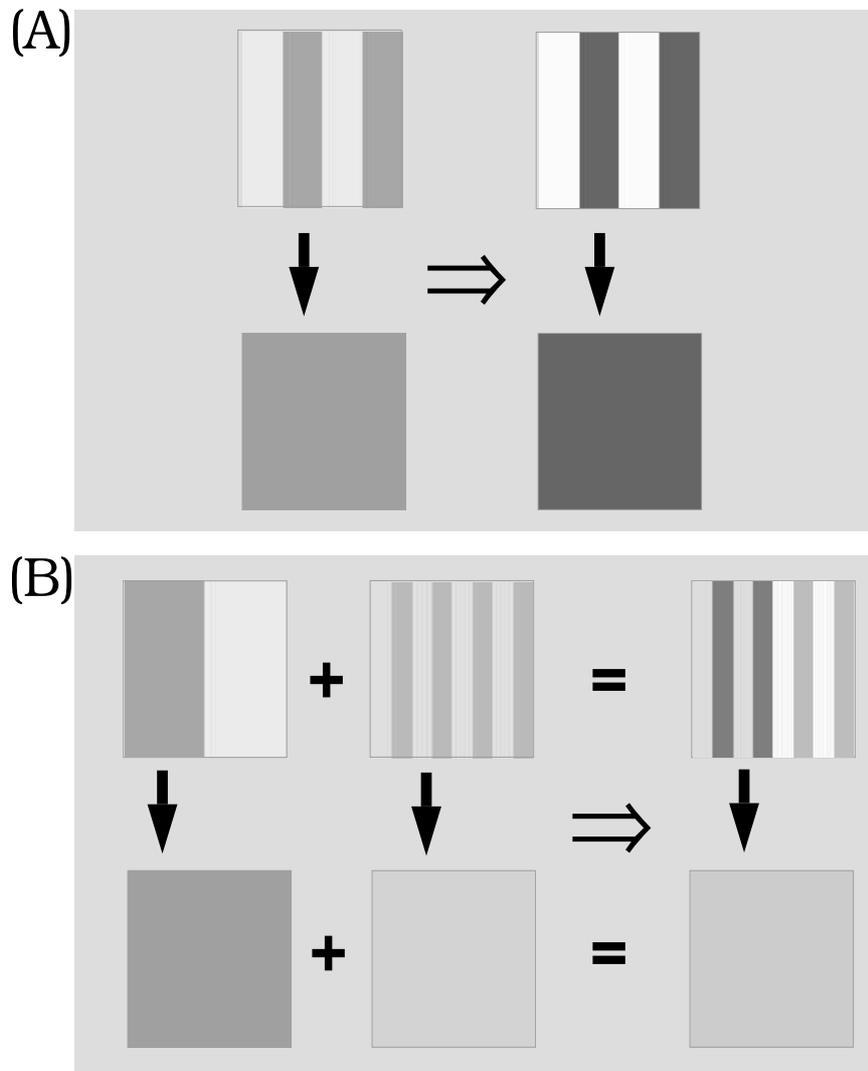


Figure 6.6 Tests of homogeneity (A) and additivity (B) are shown in graphical form. (A) Suppose that a bar within a squarewave pattern is matched by a uniform matching box. Will doubling the contrast of the test pattern also double the contrast of the matching box (homogeneity)? (B) Suppose that one of the two bars comprising a squarewave is matched by a uniform matching box; and that a bar in a second squarewave is matched by a second uniform matching box. Will the superposition of these bars be matched by the sum of the uniform matching fields (additivity)?

4. APPLICATIONS IN COLOR METRICS

In the previous sections, we discussed how some context effects of color can be modeled using linear scaling operations in different color spaces. In particular, illuminant adaptation effects can be modeled with scaling in the cone space, and spatial context effects can be modeled with spatial scaling in an opponent color space. These findings are useful in engineering applications when we need to evaluate accuracy of color rendering under different illuminant conditions, and in images as part of complex spatial patterns.

Metrics for predicting the visibility of color changes of large uniform targets under a fixed viewing condition have been used widely in industry to describe tolerances for color reproduction of large samples. The CIELAB color difference metric (CIE, 1971) is one important standard that has been used successfully for many years. In the CIELAB color difference calculation, XYZ values of the two colors to be compared are first normalized by the XYZ values of a "white point", to partially account for the effect of adaptation to different illuminants. The normalized XYZ values are then converted to $L^* a^* b^*$ values through a non-linear transformation. A perceived difference between the two samples is calculated as the Euclidean distance of the two colors in this $L^* a^* b^*$ space. In this calculation, the white point normalization step is done in the XYZ space instead of the cone space, as psychophysical findings indicated. There is also no spatial processing to account for the spatial context effect. The linear scaling models for adaptation effects and spatial effects can be, and have been, incorporated into extensions of the CIELAB color difference metric.

RLAB

In one extension to the CIELAB metric, Fairchild and Berns (1993) have proposed an extension, called RLAB, with an improved correction for light adaptation. Given the XYZ value of a surface under one illuminant, RLAB calculates a new set of XYZ values that approximates an appearance match of the original color under a different illuminant. The main calculation used for this correction is an independent scaling in the cone space (von Kries scaling) to model illuminant adaptation. This model is consistent with the findings in previous asymmetric color matching experiments, thus more accurate than the XYZ normalization step in the CIELAB calculation. After scaling in cone coordinates, another 3×3 transformation is performed on the LMS values to account for some luminance-dependent effects on appearance (increase of perceived colorfulness and contrast with increasing luminance). The LMS

values are transformed back to XYZ values in preparation for the regular CIELAB calculation.

The RLAB extension to CIELAB can be very useful in applications where color reproduction has to be done across media and for different viewing conditions. To test how well the RLAB extension reduces errors in color reproduction, Fairchild and Berns experimentally compared CRT reproductions of hardcopy color samples. They showed three different CRT reproductions for each hardcopy color sample, one of them an XYZ match, one a CIELAB match, and one a RLAB match. They found that observers most often (about 2/3 of the times) picked the RLAB match as the best reproduction. This confirms that applying the linear scaling operation in the correct color space improves accuracy of chromatic adaptation predictions by color metrics.

The work on RLAB is a clear and straightforward application of the von Kries principle. The same principle can be found in other engineering applications on the topic of color appearance (Hunt, 1982, 1987, 1991, 1994, 1995, Hunt and Pointer, 1985; Nayatani et al., 1981; Luo et al., 1996, CIECAM97).

S-CIELAB

The spatial effects on color appearance have been used in a second proposed extension to CIELAB, called S-CIELAB (Zhang and Wandell, 1997). This extension adds a spatial adjustment to the metric to correct for the loss of color sensitivity at high spatial frequencies. Consistent with the data described above, the calculation is performed in a pattern-color separable opponent color space.

The pattern-color separable opponent color space in the S-CIELAB color metric is motivated by concerns of generalization and computational efficiency. In a pattern-color separable color representation, the spatial sensitivity to the pattern in one color channel stays the same regardless of the patterns in the other color channels. In many spatial color metrics (e.g., Watson, 1993; van den Branden Lambrecht and Farrell, 1996), the assumption of pattern-color separability is made implicitly. Therefore, we can perform the spatial filtering independently in each color channel, and then combine the results. If the spatial sensitivity of one color channel is not independent of the other color channels, then one would have to figure out a different spatial sensitivity curve every time the patterns in the other color channels change. To use such pattern sensitivity data, an enormous lookup table has to be built that contains the pattern sensitivity data for each possible color contrast direction. This method is apparently very inefficient, and almost impos-

sible to build a practical metric with. Therefore, we have the practical need to find a pattern-color separable color space for a color difference metric that incorporates spatial effects.

The pattern-color separability assumption needs to be experimentally validated, however. If chromatic aberration is the sole factor leading to color-dependent pattern sensitivity, then pattern-color sensitivity is not separable in a 3-dimensional color representation, i.e. three spatial sensitivity curves are not enough to characterize the whole wavelength-dependent modulation transfer function (Marimont and Wandell, 1994). There have been other measurements of spatial effects done in different color spaces (e.g. Kelly, 1983).

The Poirson and Wandell studies (Poirson and Wandell, 1993, 1996), discussed in detail earlier in this chapter, tested the separability hypothesis explicitly. They concluded that a pattern-color separable model describes spatial effects on color appearance and discrimination to a first-order of accuracy. These results suggest that there is further neural blurring on the input images after the chromatic aberration of the optics of the eye, which is confirmed in a few other studies (Mullen, 1985; Sekiguchi et al., 1993). Neural processing combined with chromatic aberration yields a system that is approximately pattern-color separable.

The Poirson and Wandell measurements were made on simple square-wave and sinusoidal patterns. Can we generalize the pattern-color separable results to real images? Bäuml and Wandell (1996) showed that simple combination of squarewave patterns support the same pattern-color separable appearance model (discussed in the earlier sections). The opponent color directions and spatial tuning curves derived from their data are consistent with the Poirson and Wandell results. This confirms a first step of generalization toward complex spatial patterns.

From the combination of all these studies, we conclude that a pattern-color separable architecture is suitable for a spatial color difference metric. It allows the use of three independent spatial scaling functions for three color directions in an opponent color space, which is both computationally easy, and gives us some assurance that the results can generalize to a variety of color images to the first order of approximation.

Implementation of S-CIELAB. The implementation of S-CIELAB includes a spatial scaling step in an opponent color space taken from the Poirson and Wandell (1993) color appearance experiment. This step is prior to the usual CIELAB color difference calculation. The scaling of different spatial frequencies is a spatial filtering operation. The spatial

filters are also derived from the spatial tuning curves given in the Poirson and Wandell (1993) study.

Figure 6.7 illustrates how to calculate the S-CIELAB representation. The color image data are transformed into an opponent-colors space. Each opponent-colors image is convolved with a filter whose shape is determined by the visual spatial sensitivity to that color dimension; the area under each of these kernels integrates to one, so that large uniform fields will pass through the filters unchanged. Then, the filtered representation is transformed to a CIE XYZ representation, and this representation is transformed to CIELAB space. The resulting S-CIELAB representation includes both the spatial filtering and the CIELAB processing.

For color difference calculation between an original image and its reproduction, the CIELAB ΔE value is computed, point by point, between the two S-CIELAB representations. This produces a perceptual color difference image that describes the difference between the two images.

Experimental evaluation of S-CIELAB. Does the addition of a spatial processing step to CIELAB improve the metric's accuracy in predicting perceptual color differences? Several experiments have been conducted to evaluate S-CIELAB on color patterns and images.

In an experiment on halftone texture visibility (Zhang et al., 1997c), S-CIELAB predicted the visibility of printed color halftone patterns more accurately than regular CIELAB, as shown in the data plots of Figure 6.8. In another experiment (Zhang et al., 1997b; Zhang and Wandell, 1998), human subjects were asked to mark on computer-displayed halftone or JPEG images where they see distortions from the original image. This experiment provided subjective data in the form of "image distortion maps" that indicate probability of visible distortions for the halftone or JPEG images. Figure 6.9 plots the ROC curves for CIELAB, S-CIELAB, and RMSE (root mean square error on device RGB values) metrics, which shows how well each metric predicts subjects' responses. A metric that accurately predicts the subjective decisions should have hit rates (metric sees distortion when subjects see it) as large as possible, and false alarm rates (metric sees distortion when subjects do not see it) as small as possible, i.e. an ROC curve that bows away from the diagonal line as much as possible. The ROC curve labeled **Data** is generated using mean image distortion map as the predictor. This marks (approximately) the best a metric can do in terms of prediction accuracy. The S-CIELAB metric predicted where subjects see halftone distortions more accurately than CIELAB, indicating an advantage in adding a spatial processing step in the CIELAB metric.

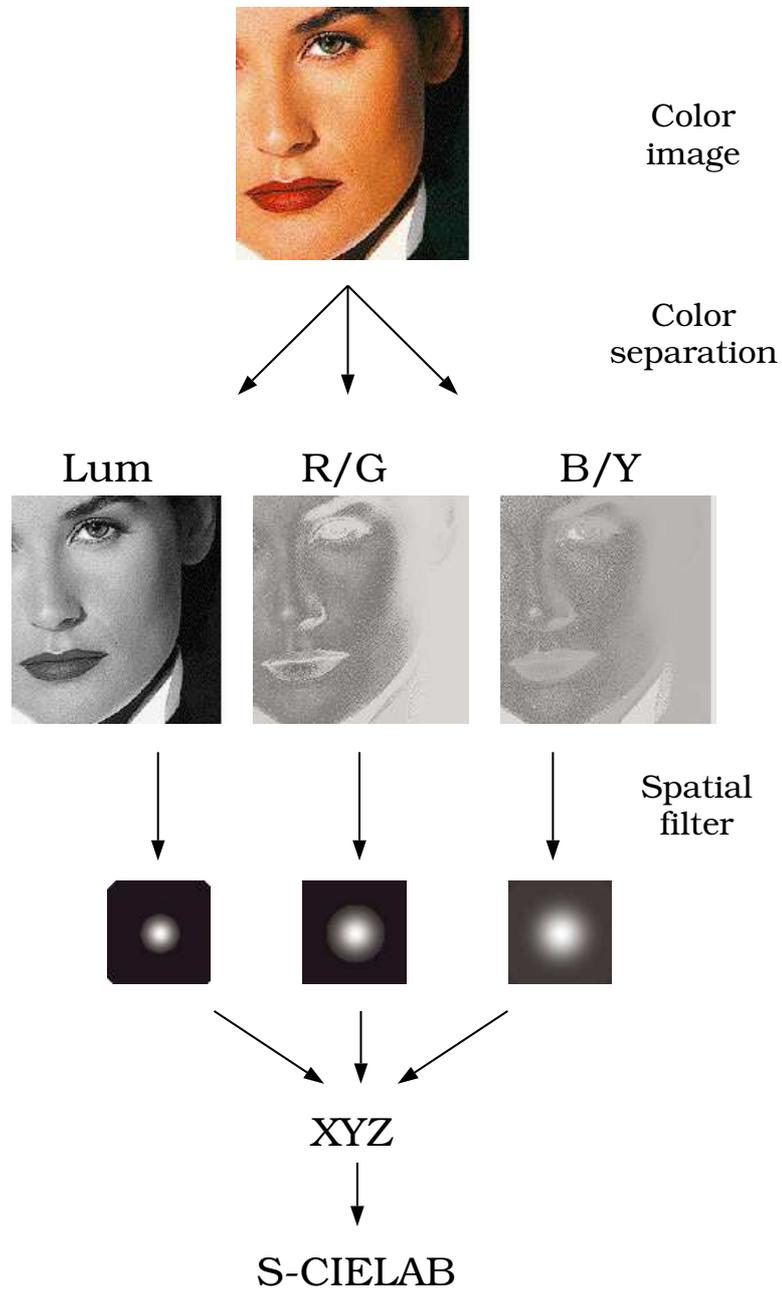


Figure 6.7 A flow-chart showing how to compute S-CIELAB.

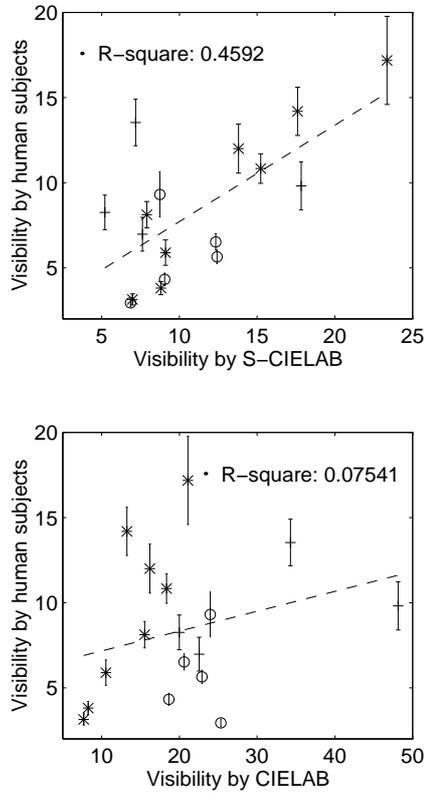


Figure 6.8 Perceptual texture visibility of printed halftone patterns plotted against (a) visibility measured by S-CIELAB (top) and (b) visibility measured by CIELAB (bottom). Data points shown as asterisk symbols correspond to patterns halftoned with black and white dots; circles correspond to patterns halftoned with red and green dots; pluses correspond to patterns halftoned with blue and yellow dots. The data points should fall on a straight line if the model's predictions correlate perfectly with the experimental data. The R^2 values of the best-fitting straight lines are shown in the plots.

In addition, in computational simulations of multi-level color halftoning (Zhang et al., 1997a), S-CIELAB predicted that adding more ink levels to the Magenta-color ink yields the largest reduction in visibility of halftone artifacts (Figure 6.10), whereas adding more levels on the yellow ink does not improve the quality. This prediction is consistent with current industry practices on uses of multi-level color halftoning. Unmodified CIELAB, however, did not make the correct prediction in this case. The addition of a spatial scaling process in opponent color space in S-CIELAB allows it to make qualitatively useful predictions for color imaging applications.

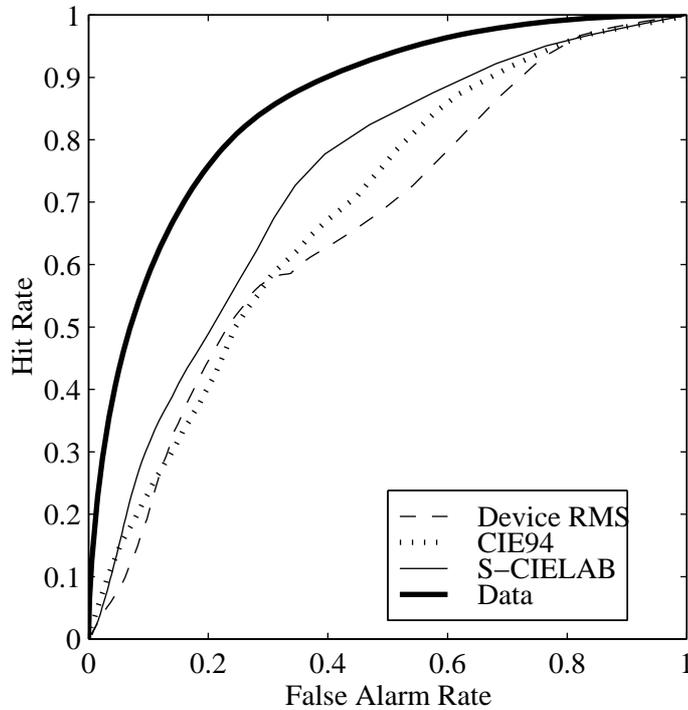


Figure 6.9 ROC curves for RMSE, CIELAB ΔE_{94} , and S-CIELAB ΔE_{94} , using hit rates and false alarm rates computed from data on the halftone image pairs, at many threshold levels for each metric.

The simple spatial scaling in opponent color space used in S-CIELAB only accounts for first-order spatial effects on color appearance and discrimination. For image color distortions that are more structured, such as JPEG distortions, more complicated mechanisms (e.g. contrast masking effects, Legge and Foley, 1980; Wilson et al., 1983; De Valois and Switkes, 1983) have to be introduced to accurately model the perceptual effect. In the "image distortion map" experiment, neither CIELAB nor S-CIELAB did a good job predicting the visibility of JPEG distortions. Other color image metrics exist that take into account contrast masking effects, such as DCTune (Watson, 1993). In DCTune, spatial context effect is accounted for by scaling in the DCT domain. In addition, contrast masking is modeled by adjustment of gain according to the background contrast in each DCT channel. The DCTune metric made accurate threshold predictions on the visibility of JPEG compression artifacts (Watson et al., 1997).

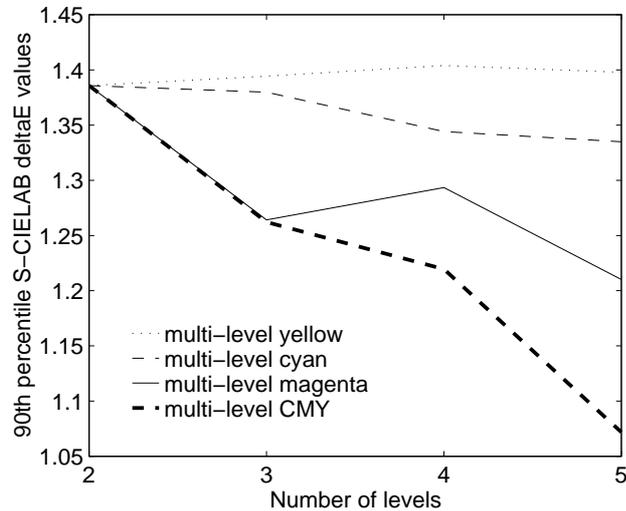


Figure 6.10 Comparison of using multiple ink levels on different inks. The thick line represents simulation results using multiple ink levels on all 3 inks. The thin lines represents simulations using multiple ink levels on one of the inks (with binary levels on the other two inks).

The experimental data discussed above showed that our understanding of the spatial effect on color can have an impact in real world applications. The performance of a standard color metric is improved when we added frequency-dependent spatial scaling in an opponent color space.

SUMMARY OF APPLICATIONS

From the RLAB metric and S-CIELAB metric, we see that modeling color and spatial context effects in the right color space has practical benefits in color imaging applications. Both the RLAB calculation and the S-CIELAB spatial filtering calculation are relatively independent pre-processing modules for the CIELAB color difference calculation. Therefore, it is easy to combine the two operations to create a CIELAB color difference metric that incorporates both the effect of spatial pattern and the effect of illuminant adaptation. Such a combined metric can be tested when experimental data are available on color reproduction of images under different viewing conditions.

Notes

1. The cone absorptions of the test and matching lights were represented as differences from the color coordinates of the background field, so that the experiments included measurements of both increments and decrements. Representing the test lights in terms of contrast is a common practice in visual psychophysics. Using this principle specifically for measuring asymmetric color matches was suggested by Walraven's (1976) discounting-the-background principle.

2. The rows of the color transformation matrix C (see Figure RESP PATTERN) define the relative contributions of the L, M, and S cone responses to each pathway. Each pathway's color responsivity can be plotted as a function of wavelength. These wavelength responsivity curves are calculated from the entries of the matrix C as follows. The i^{th} row contains the relative contributions of L, M, and S cones for the i^{th} pathway. Hence, the spectral responsivity of the i^{th} pathway is $c_{i1}L(\lambda) + c_{i2}M(\lambda) + c_{i3}S(\lambda)$, where $L(\lambda)$, $M(\lambda)$, and $S(\lambda)$ are the spectral responsivity of L, M, and S cones respectively.

References

- Anderson, S., Mullen, K., and Hess, R. (1991). Human peripheral spatial resolution for achromatic and chromatic stimuli: Limits imposed by optical and retinal factors. *Journal of Physiology*, 442:47–64.
- Arend, L. (1993). How much does illuminant color affect unattributed colors? *Journal of the Optical Society of America A*, 10:2134–2147.
- Bäumel, K.-H. (1995). Illuminant changes under different surface collections: examining some principles of color appearance. *Journal of the Optical Society of America A*, 12:261–271.
- Bäumel, K.-H. (1999). Simultaneous color constancy: how surface color perception varies with the illuminant. *Vision Research*, 39:1531–1550.
- Bäumel, K.-H. (2001). Increments and decrements in color constancy. *Journal of the Optical Society of America A* (in press).
- Bäumel, K.-H. and Wandell, B. A. (1996). The color appearance of mixture gratings. *Vision Research*, 36:2849–2864.
- Brainard, D. H. (1998). Color constancy in the nearly natural image. 2. Achromatic loci. *Journal of the Optical Society of America A*, 14:307–325.
- Brainard, D. H., Brunt, W. A., and Speigle, J. M. (1997). Color constancy in the nearly natural image. 1. Asymmetric matches. *Journal of the Optical Society of America A*, 14:2091–2110.
- Brainard, D. H. and Wandell, B. A. (1992). Asymmetric color matching: how color appearance depends on the illuminant. *Journal of the Optical Society of America*, 9:1433–1448.
- Burnham, R. W., Evans, R. M., and Newell, S. M. (1952). Influence of color perception of adaptation to illumination. *Journal of the Optical Society of America*, 42:597.
- Chichilnisky, E.-J. and Wandell, B. A. (1995). Photoreceptor sensitivity changes explain color appearance shifts induced by large uniform backgrounds in dichoptic matching. *Vision Research*, 35(2):239–254.

- Chichilnisky, E.-J. and Wandell, B. A. (1996). Seeing gray through the on and off pathways. *Visual Neuroscience*, 13(3):591–596.
- Chichilnisky, E.-J. and Wandell, B. A. (1999). Trichromatic opponent colors classification. *Vision Research*, 39(20):3444–3458.
- CIE (1971). *Recommendations on uniform color spaces, color difference equations, psychometric color terms*. Publication CIE 15 (E.-1.3.1), Supplement No.2. Bureau Central de la CIE, Vienna.
- De Valois, K. K. and Switkes, E. (1983). Simultaneous masking interactions between chromatic and luminance gratings. *Journal of the Optical Society of America*, 73(1):11–18.
- De Valois, R. L., De Valois, K. K., Switkes, E., and Mahon, L. (1997). Hue scaling of isoluminant and cone-specific lights. *Vision Research*, 37(7):885–897.
- Elsner, A. E., Berns, S. A., and Pokorny, J. (1987). Changes in constant-hue loci with spatial frequency. *Color Research and its Applications*, 12:42–50.
- Fairchild, M. D. and Berns, R. S. (1993). Image color-appearance specification through extension of cielab. *Color Research and Application*, 18(3):178–190.
- Granger, E. M. and Heurtley, J. C. (1973). Visual chromaticity — modulation transfer function. *Journal of the Optical Society of America*, 63:1173–1174.
- Hunt, R. W. G. (1982). A model of colour vision for predicting colour appearance. *Color Research and Applications*, 7:95–112.
- Hunt, R. W. G. (1987). A model of colour vision for predicting colour appearance in various viewing conditions. *Color Research and Applications*, 12:297–314.
- Hunt, R. W. G. (1991). Revised colour-appearance model for related and unrelated colours. *Color Research and Applications*, 16:146–165.
- Hunt, R. W. G. (1994). An improved predictor of colourfulness in a model of colour vision. *Color Research and Applications*, 19:23–26.
- Hunt, R. W. G. (1995). *The Reproduction of Colour*. Fountain Press, England, 5th edition.
- Hunt, R. W. G. and Pointer, M. R. (1985). A colour-appearance transform for the cie 1931 standard colorimetric observer. *Color Research and Applications*, 10:165–179.
- Hurvich, L. M. and Jameson, D. (1958). Further development of a quantified opponent-colours theory. In Teddington, editor, *England National Physical Laboratory Symposium No.8: Visual Problems of Color*, London: H.M.S.O.
- Ingling Jr., C. R. and Schneibner, H. M. O. (1970). Color naming of small foveal fields. *Vision Research*, 10:501–511.

- Kaiser, P. K. and Boynton, R. M. (1996). *Human color vision*. Optical Society of America, Washington, D.C., 2nd edition.
- Kelly, D. H. (1983). Spatiotemporal variation of chromatic and achromatic contrast thresholds. *Journal of the Optical Society of America*, 73(6):742–750.
- Knowles Middleton, W. E. and Holmes, M. C. (1949). The apparent colors of surfaces of small subtense - a preliminary report. *Journal of the Optical Society of America*, 39(7):582–592.
- Legge, G. E. and Foley, J. M. (1980). Contrast masking in human vision. *Journal of the Optical Society of America*, 70:1458–1471.
- Luo, M. R., Lo, M.-C., and Kuo, W.-G. (1996). The LLAB(l:c) colour model. *Color Research and Applications*, 21:412–429.
- Marimont, D. L. and Wandell, B. A. (1994). Matching color images: the effects of axial chromatic aberration. *Journal of the Optical Society of America*, 11(12):1–11.
- Mausfeld, R. and Niederee, R. (1993). An inquiry into relational concepts of colour, based on incremental principles of colour coding for minimal relational stimuli. *Perception*, 22(4):427–462.
- Mullen, K. T. (1985). The contrast sensitivity of human colour vision to red-green and blue-yellow chromatic gratings. *Journal of Physiology*, 359:381–400.
- Nayatani, Y., Takahama, K., and Sobagaki, H. (1981). Formulation of a nonlinear model of chromatic adaptation. *Color Research and Application*, 6(3):161–171.
- Poirson, A. B. and Wandell, B. A. (1993). Appearance of colored patterns: pattern-color separability. *Journal of the Optical Society of America*, 10(12):2458–2470.
- Poirson, A. B. and Wandell, B. A. (1996). Pattern-color separable pathways predict sensitivity to simple colored patterns. *Vision Research*, 36(4):515–526.
- Sekiguchi, N., Williams, D. R., and Brainard, D. H. (1993). Efficiency in detection of isoluminant and isochromatic interference fringes. *Journal of the Optical Society of America*, 10(10):2118–2133.
- Shevell, S. K. (1978). The dual role of chromatic backgrounds in color perception. *Vision Research*, 18:1649–1661.
- van den Branden Lambrecht, C. J. and Farrell, J. E. (1996). Perceptual quality metric for digitally coded color images. In *Proceedings of the European Signal Processing Conference*, pages 1175–1178, Trieste, Italy. Available on http://1tswwww.epfl.ch/pub_files/vdb/.
- von Kries, J. (1905). Die Gesichtsempfindungen. In W. Nagel (Ed.), *Handbuch der Physiologie des Menschen* (Vol. 3, pp. 109-279). Braunschweig: Vieweg.

- Walraven, J. (1976). Discounting the background - the missing link in the explanation of chromatic adaptation. *Vision Research*, 16:289–295.
- Wandell, B. (1995). *Foundations of Vision*. Sinauer Press, Sunderland, MA.
- Wassef, E. G. T. (1959). Linearity of the relationship between the tristimulus values of corresponding colours seen under different conditions of chromatic adaptation. *Optica Acta*, 6:378.
- Watson, A. B. (1993). DCT quantization matrices visually optimized for individual images. In *SPIE Proceedings*.
- Watson, A. B., Borthwick, R., and Taylor, M. (1997). Image quality and entropy masking. In *SPIE Proceedings*, volume 3016.
- Werner, J. and Walraven, J. (1982). Effect of chromatic adaptation on the achromatic locus: The role of contrast, luminance and background color. *Vision Research*, 22:929–943.
- Wilson, H. R., McFarlane, D. K., and Phillips, G. C. (1983). Spatial frequency tuning of orientation selective units estimated by oblique masking. *Vision Research*, 23:873–882.
- Wyszecki, G. and Stiles, W. S. (1982). *Color science: concepts and methods, quantitative data and formulae*. Wiley, New York.
- Zhang, X., Farrell, J. E., and Wandell, B. A. (1997a). Applications of a spatial extension to CIELAB. In *Proceedings of the IS&T/SPIE 9th Annual Symposium on Electronic Imaging*, volume 3025, pages 154–157.
- Zhang, X., Setiawan, E., and Wandell, B. A. (1997b). Image distortion maps. In *Final Program and Proceedings of the Fifth IS&T/SID Color Imaging Conference. Color Science, Systems and Applications*, pages 120–125, Scottsdale, AZ, USA. IS&T, SID.
- Zhang, X., Silverstein, D. A., Farrell, J. E., and Wandell, B. A. (1997c). Color image quality metric S-CIELAB and its application on halftone texture visibility. In *COMPCON97 Digest of Papers*, pages 44–48. IEEE.
- Zhang, X. and Wandell, B. A. (1997). A spatial extension of CIELAB for digital color image reproduction. *Journal of the SID*, 5(1):61–63.
- Zhang, X. and Wandell, B. A. (1998). Color image fidelity metrics evaluated using image distortion maps. *Signal Processing*, 70:201–214.