The Total Chromaticity Diagram

Paul Centore

September 27, 2020

Abstract

The familiar chromaticity diagram consists of all the chromaticities that can be perceived in what might be called a state of neutral adaptation. In the von Kries model, different states of chromatic adaptation modify the retinal cones' responses, which in turn modify the perceived chromaticities of colour stimuli. This paper demonstrates mathematically that, in fact, each state of adaptation has its own chromaticity diagram (with significant overlap between the diagrams). The von Kries model is shown to predict that these new diagrams contain some colours that are outside the ordinary chromaticity diagram. This observation then necessitates enlarging the standard chromaticity diagram into a total chromaticity diagram, consisting of all chromaticities, that can occur in any adaptive state. The total diagram, which is calculated explicitly, is a superset of the standard diagram, and is approximately triangular in cone-fundamental space. As an application, the new diagram helps explain the high chromas sometimes seen in the Helson-Judd effect.

1 Introduction

In everyday life, chromatic adaptation insures that the prevailing illumination is perceived as a neutral, achromatic "white." As a side effect of neutralizing the illuminant, all colours are shifted somewhat. The von Kries model[1] is a longstanding hypothesis which attributes the shifts to three gain coefficients, operating independently on the three classes of retinal cones. Each coefficient adjusts a cone output by a multiplicative factor. The adjustments "discount the illuminant," so that an object's colour appears as it would when viewed in a neutral light, rather than being tinged with the hue of the actual illuminant. When lighting changes, the human visual system's state of adaptation changes with it, as the von Kries coefficients adjust to the overall lighting situation.

A colour can be decomposed mathematically into a luminance (specifying the colour's lightness or darkness) and a chromaticity (loosely speaking, the colour's hue and saturation). The familiar chromaticity diagram,[3] which we will denote $\mathcal{D}$, is a two-dimensional representation, based on the 1931 Standard Observer[4] defined by the Commission Internationale de l'Éclairage (CIE), of the set of chromaticities. This standard diagram is implicitly based on what might be called the "neutral" state of adaptation, which occurs, for instance in a standard colour-matching experiment, where an observer views small stimuli, typically subtending no more than a few degrees, against a black background. As this paper will show, however, many chromaticities occur only in non-neutral states, and so are left out of $\mathcal{D}$. To encompass all chromaticities, this paper will expand the standard diagram into a total chromaticity diagram $\mathcal{T}$. The total chromaticity diagram contains the standard chromaticity diagram as a subset, and also contains many new chromaticities outside it.

The foregoing statements will be proven mathematically, along with explicit calculations. A geometric setting[5] will allow easy visualization. The colours specified by the 1931 Standard Observer fill out a convex, asymmetric *spectrum cone* $\mathcal{C}$ in a real, three-dimensional vector space, sometimes called *colour space*. CIE XYZ coordinates provide a standard basis for this space. The standard chromaticity diagram is the two-dimensional section that results when the spectrum cone is cut by the plane $X + Y + Z = 1$. A set of CIE xy coordinates is conventionally used for this diagram.

The CIE's XYZ coordinates were standardized without reference to the three classes of retinal cones, whose properties were unknown in 1931. Today the retinal cones are much better studied. As a result, we can use an alternate set of LMS coordinates, where L, M, and S indicate the cones that respond predominantly to long-, medium-, and short-wavelength stimuli. A colour's LMS coordinates is a linear transformation of its XYZ coordinates. In fact, one could use the transformation to define an alternate basis for colour space that consisted of an L-vector, an M-vector, and an S-vector. The spectrum cone could be expressed equally well in LMS coordinates as in XYZ coordinates, and could be similarly sliced by the plane $L + M + S = 1$. lm coordinates can be assigned, analogous to xy coordinates, so an lm chromaticity diagram will be constructed. We will move as needed between the xy and lm diagrams.

Mathematically, the von Kries transform T for a given set of coefficients is a linear operator on three-dimensional colour space. In fact, the three coefficients, which are always positive, are eigenvalues for the transform, and the eigenvectors are the L-, M-, and S-vectors. The input of T is the XYZ vectors that would occur if a colour stimulus (a light beam that strikes the retina) were viewed in a neutral state of adaptation. T is restricted to real stimuli, so the domain of T (i.e. the set of all possible inputs) is therefore the XYZ spectrum cone $\mathcal{C}$. The output of T is a new vector in colour space, which corresponds to the colour the viewer perceives. The main mathematical insight of this paper is the observation that the image $T(\mathcal{C})$, given T's diagonal form, is another cone that overlaps its domain $\mathcal{C}$ significantly, but also (except in trivial cases) extends beyond $\mathcal{C}$. The colours that are in $T(\mathcal{C})$, but not in $\mathcal{C}$ itself, have no representation in the CIE system, yet are still genuine perceptions. The chromaticities of such colours are not contained in the standard chromaticity diagram.

This paper's main contribution is the construction of a total chromaticity diagram $\mathcal{T}$ that does contain them. To begin with, we will work in LMS and lm coordinates, where von Kries transforms have a more convenient form. In fact, we will find that the lm plane is usually sufficient, and its easy visualizability will aid understanding. Next we will consider the set of all possible von Kries transforms. Since a von Kries transform is induced by the chromaticity of the prevailing illumination, we can parametrize the set of von Kries transforms by the lm chromaticity diagram. By considering the transform resulting from each lm chromaticity, we account for all possible von Kries transforms. Now suppose that an lm chromaticity c, and thus a von Kries transform T_c, has been specified. The set of all possible chromaticities that T_c can produce is the image under T_c of the entire lm chromaticity diagram $\mathcal{L}$. The image $T_c(\mathcal{L})$ can be drawn in the same plane, $L + M + S = 1$. It largely overlaps $\mathcal{L}$, but (except in the case of the identity transform) there will also be some new chromaticities that are outside $\mathcal{L}$. To find $\mathcal{T}$, the set of all possible chromaticities, let c run over all possible chromaticites, and take the union of the resulting von Kries images:

$$\mathcal{T} = \bigcup_{c \in \mathcal{L}} T_c(\mathcal{L}). \tag{1}$$

This $\mathcal{T}$ is then the total chromaticity diagram.

The paper is organized as follows. First, some background is given, consisting largely of a geometric description of colour space, the spectrum cone, and the standard chromaticity diagram; both CIE and cone coordinates are used. The von Kries model and some formalism for it are also presented. Second, the derivation of the total chromaticity diagram is detailed. The set of all von Kries transforms is shown to be indexed by $\mathcal{L}$, and the union of the images of $\mathcal{L}$ under all von Kries transforms is calculated, producing the total chromaticity diagram $\mathcal{T}$. Third, there is a brief description of how $\mathcal{T}$ helps explain the high chromas sometimes seen in the Helson-Judd effect. Finally, a short summary of the paper is given.

2 Chromatic Adaptation

2.1 Neutral Adaptation

In a traditional colour-matching experiment, a subject views two colours side-by-side, through an aperture. The colours themselves, which typically subtend a small angle of perhaps 2°, are viewed against a null background, which contains no visual stimulus and so appears black. Since the colours take up very little area, and since there is no other stimulus, the eye might be said to be in a state of neutral adaptation. The 1931 CIE Standard Observer[4] implicitly assumes such a neutral state.

Though they are very different physically, two side-by-side stimuli can still appear identical to an observer. In fact a *colour* can be defined as the maximal set of physical stimuli which appear identical in a standard colour-matching experiment. Further analysis[5] involving Grassmann's laws implies that the colours that the Standard Observer distinguishes fill an irregularly shaped convex cone, called the *spectrum cone*, in a real, three-dimensional vector space. The boundaries of the cone, as shown in Figure 1, are rays corresponding to monochromatic colours, that have power in only a single wavelength. The 1931 standard assigns the vector space a basis of three vectors, conventionally denoted X, Y, and Z, and any colour can be assigned XYZ coordinates. Many, in fact most, XYZ triples are mathematical fictions, or *imaginary colours*, which cannot be produced by any physical stimulus. The three axes themselves are imaginary: no physical stimulus, for instance, can produce only an X-response, without also producing a Y- or Z-response. X, Y, and Z were chosen so that all real colours would have non-negative XYZ coordinates, making the spectrum cone a subset of the positive octant.

While the spectrum cone itself is three-dimensional, an informative two-dimensional subset of it can be constructed, called the *chromaticity diagram* $\mathcal{D}$, by cutting the spectrum cone with the plane $X + Y + Z = 1$, as shown in Figure 1. Every ray through the origin in the spectrum cone intersects that plane in a unique point. The plane is given its own system

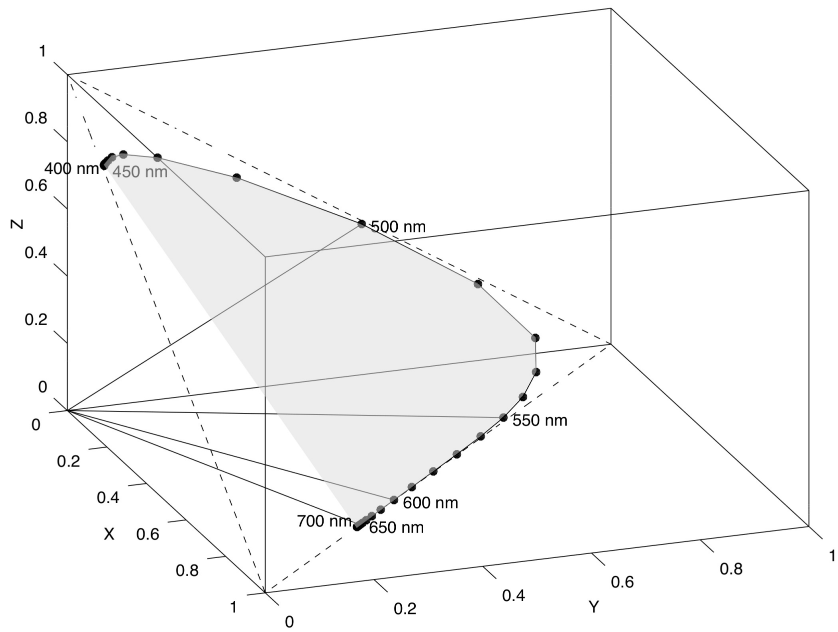

Figure 1: The XYZ Spectrum Cone, Cut by the Plane $X + Y + Z = 1$

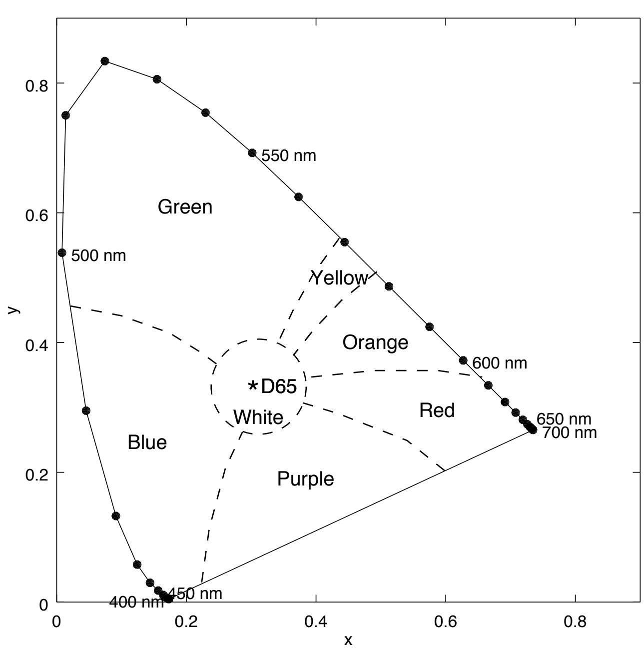

Figure 2: The xy Chromaticity Diagram

of chromaticity coordinates:

$$x = \frac{X}{X+Y+Z}, \tag{2}$$
$$y = \frac{Y}{X+Y+Z}, \tag{3}$$

so every colour has an xy chromaticity. Figure 2 shows the chromaticity diagram as it is typically presented, with its x and y axes transformed to form a right angle. The boundary of the chromaticity diagram, called the *spectrum locus*, intersects all the monochromatic rays in the visible spectrum, from 400 to 700 nm, as shown in the figure.

Chromaticity is often broken down further into a hue and a saturation component. Figure 2 shows the approximate locations of various hues in the chromaticity diagram. A central core of whites, or colours without discernible hues, is surrounded by a rainbow of basic hues like red, orange, yellow, etc. The colours near the spectrum locus are very vivid and saturated, and the saturation decreases steadily as one moves inward. The "center" of the diagram, though not defined exactly, will be taken to be $(x, y) = (0.31, 0.33)$, the chromaticity of Illuminant D65.

The familiar chromaticity diagram $\mathcal{D}$ contains the chromaticity of every colour that can be perceived in the context of standard colour-matching experiments. Those experiments, however, implicitly impose a neutral adaptation. We will soon see that other states of adaptation allow further chromaticities, outside $\mathcal{D}$.

2.2 Cone Fundamentals

The human retina contains three kinds of cones that respond physically to incoming light: L cones that respond mainly to longer wavelengths, M cones that respond to middle wavelengths, and S cones that respond to shorter wavelengths. The relative magnitude of the cone responses depends on the chromaticity of the light they are responding to. While LMS responses have yet to be measured directly, one can show mathematically that the LMS responses are a linear transformation of the XYZ responses. Many numerical forms of the transformation have been proposed; this paper will use the Hunt-Pointer-Estevez[2] (HPE) matrix:

$$\begin{bmatrix} L \\ M \\ S \end{bmatrix} = \begin{bmatrix} 0.3897 & 0.6890 & -0.0787 \\ -0.2298 & 1.1834 & 0.0464 \\ 0.0000 & 0.0000 & 1.0000 \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \end{bmatrix}. \tag{4}$$

L, M, and S can be viewed as vectors, defined by Equation (4), in the same three-dimensional space as X, Y, and Z. Mathematically, in fact, LMS and XYZ are just two different but interchangeable bases for three-dimensional colour space. The spectrum cone can be drawn in LMS coordinates just as readily as in XYZ coordinates. Chromaticity coordinates l and m can be defined in LMS space by cutting the spectrum cone with the plane $L+M+S=1$ (see Figure 3), just as x and y are defined in XYZ space:

$$l = \frac{L}{L+M+S}, \tag{5}$$

$$m = \frac{M}{L+M+S}. \tag{6}$$

Figure 4 shows the resulting lm chromaticity diagram, which we will denote $\mathcal{L}$.

2.3 The von Kries Model

Colour-matching experiments typically display a stimulus against a black background, inducing what we have called a neutral adaptation. The same stimulus, however, can appear different when viewed against another colour, or in a more complicated setting, and it has been concluded that the visual system adapts to the viewing conditions. When viewing a realistic scene, the eye receives a multitude of stimuli, most of which are influenced by the prevailing illumination. By a mechanism that is not yet understood, the visual system extracts sufficient information about the lighting that it can automatically adjust its colour perceptions, in a process called *chromatic adaptation*.

At the start of the twentieth century, Johannes von Kries[1] proposed a model, sometimes called the '*von Kries hypothesis*, for chromatic adaptation. The model relates two sets of responses:

1. (L_n, M_n, S_n) : the initial cone responses, due solely to physical causes, without any adjustment, and
2. (L_a, M_a, S_a) : the post-adaptation cone responses.

They are related by simple coefficients:

$$L_a = k_L L_n, \tag{7}$$

$$M_a = k_M M_n, \tag{8}$$

$$S_a = k_S S_n. \tag{9}$$

The relationships can also be written as a linear transformation:

$$\begin{bmatrix} L_a \\ M_a \\ S_a \end{bmatrix} = \begin{bmatrix} k_L & 0 & 0 \\ 0 & k_M & 0 \\ 0 & 0 & k_S \end{bmatrix} \begin{bmatrix} L_n \\ M_n \\ S_n \end{bmatrix}. \tag{10}$$

The linear transformation in Equation (10), called a *von Kries transform*, has a special form: not only is it diagonal, but all its eigenvalues are positive, and the three eigenvectors are always just the L, M, and S axes.

These relationships can result in colour constancy. Suppose a subject is viewing an object under some prevailing illumination I, where I is a power function over the visible spectrum. I, for example, could result from incandescent lighting, fluorescent lighting, indirect daylight, etc. The object has a physically invariant reflectance spectrum ρ, which is also a function (between 0 and 100%) over the visible spectrum. The stimulus that reaches the eye is then

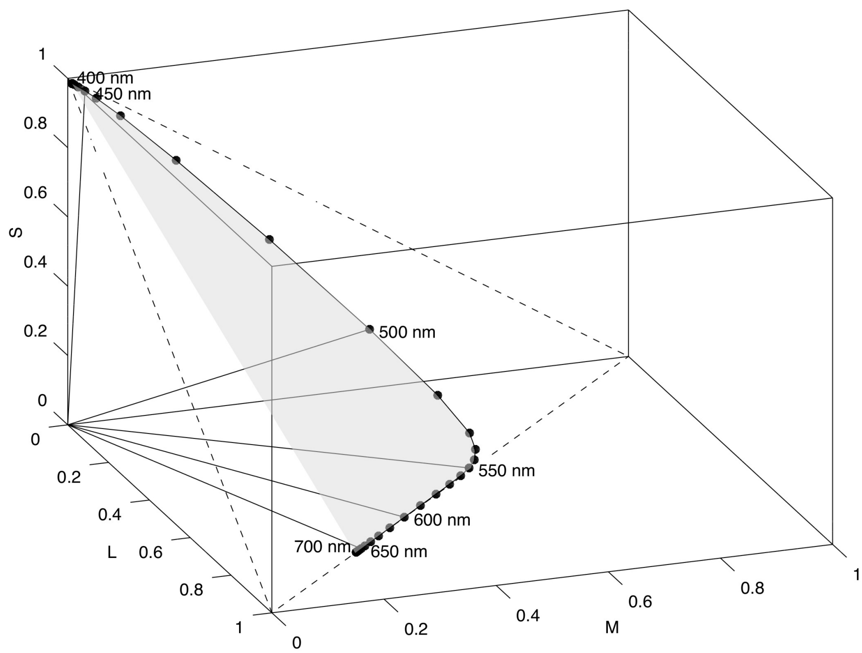

Figure 3: The LMS Spectrum Cone, Cut by the Plane $L + M + S = 1$

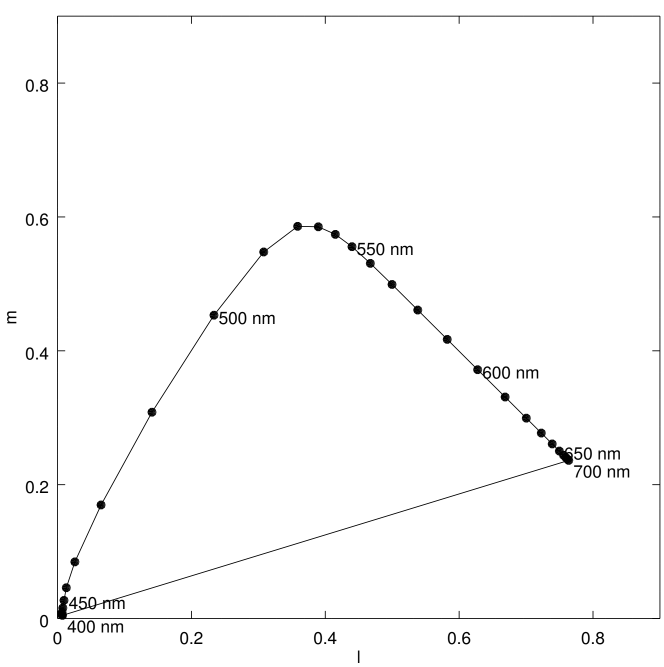

Figure 4: The lm Chromaticity Diagram

$I\rho$, and this stimulus varies greatly as I varies. Colour constancy, which has been observed empirically, means that the object's colour appears the same to the observer, even as lighting changes from incandescent to fluorescent to natural, and so on.

Chromatic adaptation is needed for colour constancy. The idea is that the visual system adapts to the illumination by heightening or reducing the sensitivity of each kind of cone. The coefficients k_L, k_M, and k_S quantify the adjustment. Without any adjustment, a stimulus would produce some cone responses L_n, M_n and S_n in an unadapted observer. The visual system, applying Equations (7) through (9), adapts by multiplying the raw responses by k_L, k_M, and k_S. The result is L_a, M_a, and S_a, which is just the coordinates the object would produce if the object were neutrally illuminated. The observer then perceives the object's colour as the colour it would have in a neutral setting, so the illumination has been corrected for.

As an example, suppose an adapted viewer observes a blank white canvas hanging on a wall, in a windowless room illuminated with a greenish light. In a colour-matching experiment, against a null background, that light has CIE coordinates $XYZ = (42.0, 100.0, 25.0)$, which implies $L_nM_nS_n = (83.3, 109.9, 25.0)$. The greenish light reflects off the canvas; the reflected light is the same shade of green. The light reflected from the canvas enters the viewer's eye and produces those values as its unadjusted physical cone response: $L_nM_nS_n = (83.3, 109.9, 25.0)$. The viewer, however, is adapted to the greenish illuminant, and discounts it with coefficients $k_L = 11.7, k_M = 9.2, k_S = 43.5$, producing adapted signals

$$L_a = k_L L_n = 11.7 \cdot 83.3 = 975 \tag{11}$$

$$M_a = k_M M_n = 9.2 \cdot 109.9 = 1011 \tag{12}$$

$$S_a = k_S S_n = 43.5 \cdot 25.0 = 1083. \tag{13}$$

The adapted $L_aM_aS_a$ values can be converted to XYZ values of $(957, 998, 1083)$, which are consistent with Illuminant D65, a white light. The viewer will therefore perceive the white canvas as white, even though the light the canvas sends to his eye is green.

A numerical issue arises when using the von Kries model for long-wavelength monochromatic stimuli. The Z values for wavelengths of 650 nm or greater is 0, and the HPE matrix, whose bottom row is $[0, 0, 1]$, sends the Z value directly to the S value. Then both sides of Equation (9) are 0, so that k_S is undefined, and the von Kries model cannot be applied. To handle this situation, no calculations were made on stimuli that only took on positive values between 650 and 700 nm inclusive. When a calculation was needed at 700 nm, the 700 nm point was replaced with a point we denote ∼700 nm, which is one-hundredth of the way along the line joining 700 nm to 400 nm. These adjustments chop a barely visible notch off the bottom right corner of the lm diagram, and do not affect the calculations significantly.

3 Derivations and Calculations

Now that the von Kries model has been introduced, we can use it to calculate the total chromaticity diagram, which consists of all possible chromaticities, occurring over all possible states of adaptation. Each physically possible illuminant takes on a chromaticity in the lm chromaticity diagram shown in Figure 4, so there is one set of von Kries coefficients (up to a multiplicative factor) for each lm chromaticity. The lm chromaticities thus serve to index the set of von Kries transforms. Equation (10) allows the image of a von Kries transform to be calculated. Any transform's domain can be taken to be the set of lm chromaticities. Most of its image will consist of lm chromaticities from Figure 4, but some points in the image will be outside that diagram. The union of all the images of all possible transforms is the total chromaticity diagram $\mathcal{T}$.

The following subsections work out the steps in the above development. It will be seen that the derivations could take place equally well in two-dimensional chromaticity space (using chromaticity diagrams) or three-dimensional colour space (using cones). We will focus predominantly on the two-dimensional case, where helpful pictures can be drawn easily. Convexity results, which take advantage of the special form of von Kries matrices, will significantly simplify calculations.

3.1 The Set of All von Kries Transforms

Equation (10) writes a von Kries transforms as a diagonal matrix, whose diagonal entries define the transform. While the three entries can vary, physical requirements imply that not all sets of three numbers can occur as entries. To see why, rearrange Equations (7) through (9) to get

$$k_L = L_a/L_n, \tag{14}$$
$$k_M = M_a/M_n, \tag{15}$$
$$k_S = S_a/S_n. \tag{16}$$

The von Kries coefficients on the left depend on two sets of LMS coordinates, one of which is an adapted version of the other. Suppose we specify just one physically possible input, and its adjusted output. Then those vectors define the coefficients, and therefore define the transform, so we could calculate the image of any other input.

Empirical investigations show that a white or neutral, which we will take to be D65, is always a possible perception, so on empirical grounds there is no loss of generality in requiring that $L_aM_aS_a = (97.4, 101.6, 108.9)$, which are (up to multiplication by a constant) the coordinates for D65:

$$k_L = 97.4/L_n \tag{17}$$
$$k_M = 101.6/M_n \tag{18}$$
$$k_S = 108.9/S_n, \tag{19}$$

Assigning an input vector that produces D65 as its output is then sufficient to define a von Kries transform, and every von Kries transform can be specified that way. Thus the set of von Kries transforms is in one-to-one correspondence with the set of colours in the LMS spectrum cone.

Furthermore, since Equations (14) through (16) are invariant if the two LMS vectors are multiplied by the same factor, and chromaticities do not change if the three von Kries coefficients are all multiplied by the same factor, we can simplify matters by working in two-dimensional chromaticity space instead of three-dimensional colour space. The lm coordinates of D65 are $(0.32, 0.33)$, so we can construct a von Kries transform by specifying which

lm pair is mapped to $(0.32, 0.33)$. A physical interpretation is that a scene is illuminated by that lm pair, and chromatic adaptation uses a von Kries adjustment to insure that lm pair is mapped to the chromaticity for D65. The set of von Kries transforms, restricted solely to chromaticities, is therefore in one-to-one correspondence with the set of chromaticities in the lm diagram shown in Figure 4.

This fact will be used in calculating the total chromaticity diagram. We will let lm vary over all possible points in Figure 4, and find the von Kries transform corresponding to those points. By the foregoing argument, this method will find all possible transforms. The union of the images of all those transforms will then be the total chromaticity diagram $\mathcal{T}$.

3.2 The Image of a von Kries Transform

This section uses a concrete example to illustrate the central observation of this paper: the image of a von Kries transform is partly outside the lm chromaticity diagram $\mathcal{L}$, and the points that are outside represent new chromaticities that have not yet been accounted for; they will take their places in the total chromaticity diagram.

In a previous example, a white canvas was viewed in a room illuminated by a greenish light I, with LMS coordinates $(83.3, 109.9, 25.0)$ and lm coordinates $(0.38, 0.50)$. Chromatic adaptation shifts I to Illuminant D65, whose lm coordinates are $(0.32, 0.33)$ Both points are shown in the lm chromaticity diagram in Figure 5.

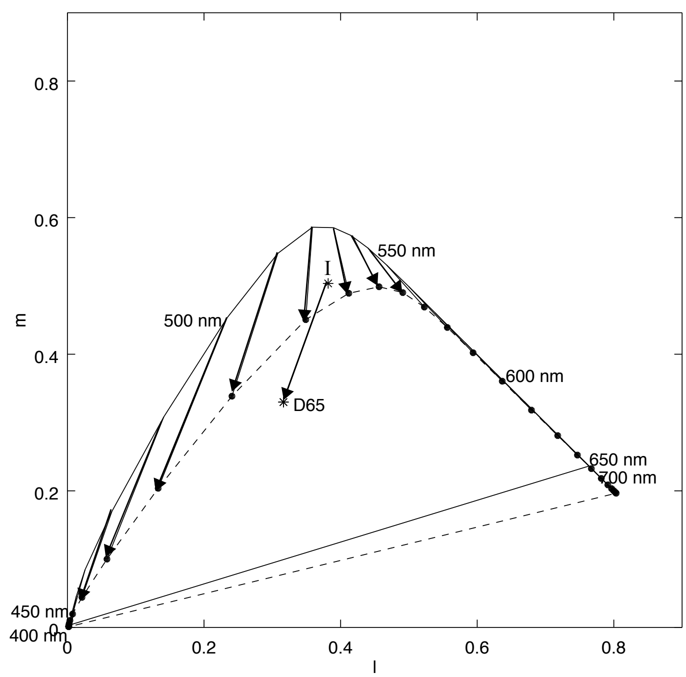

Figure 5: The Image of lm Chromaticities after Adaptation to I

By the discussion of the previous section, I leads to a unique von Kries transform T_I, completely determined by the condition that $T_I(I) = \text{D65}$. The range of T_I is the set

$$\mathcal{S} = \{T_I(l,m)\big|(l,m) \text{ is in the } lm \text{ chromaticity diagram}\}. \tag{20}$$

Figure 5 indicates with arrows how T_I moves I and the monochromatic spectrum locus chromaticities to their images. Of course, all the other chromaticities are also moved. The image of T_I is the area that is bounded by the dotted line.

The most notable feature of $\mathcal{S}$ is that it contains points that are outside the original chromaticity diagram: one can see a wedge just below the diagram $\mathcal{L}$ which is still inside $\mathcal{S}$. The colours corresponding to these outside points cannot be perceived in the normal state of adaptation, but can be perceived under the green illuminant. By multiplying L, M, and S by various coefficients, the von Kries transform has, for practical purposes, changed the cone response functions, so that previously impossible colours can be produced. These new colours will appear in the total chromaticity diagram $\mathcal{T}$.

While the idea of new colours might seem implausible at first, the following thought experiment should make it more believable. Suppose that there was some way to activate the long-wavelength cones without simultaneously activating the medium- or short-wavelength cones. Then one would expect that the colour perception produced in this way would be outside , completely novel, and have no precedent in any colours experienced previously. Such a condition would correspond to von Kries coefficients of $k_L = 1,$, $k_M = 0$, and $k_S = 0$. Even though Sect. 3.1 shows that those coefficients cannot be produced by any real illuminant, they can still be considered as a theoretically extreme von Kries case. The von Kries coefficients resulting from illuminants like the green light are not as extreme as the coefficients 1, 0, and 0, but they are still extreme enough to produce colours outside the diagram, and those colours could not be produced in normal illumination. The colours would therefore be new to a viewer who had only experienced neutral lighting conditions.

3.2.1 Simplifying Calculations via Convexity

While we have worked so far in the easily visualized two-dimensional chromaticity plane, similar derivations produce new colours in three-dimensional colour space. Linear transformations like von Kries transforms preserve convexity relations, so a von Kries image of the spectrum cone $\mathcal{C}$ must be another convex cone. Like $\mathcal{S}$, the image cone will mostly overlap the original cone, but still introduce some points outside.

In the three-dimensional case these facts about the image cone can be inferred, at least heuristically, from the form of the von Kries matrix in Equation (10). Apart from trivial examples (like multiples of the identity matrix), no diagonal transformation would be expected to send a cone, especially a cone with an irregular profile, into a subset of itself. It is not too surprising, then, to find that some of the image points are outside the original set of colours in three dimensions, as well as in two.

Convexity arguments can also streamline calculations in the chromaticity plane. Since the spectrum cone is convex, the corresponding chromaticity diagram is also convex: it is the intersection of two convex sets (the spectrum cone and the plane $L + M + S = 1$), and the intersection of two convex sets is again convex (Theorem 2.19 of Ref. 6). Now suppose we start with a convex subset such as the chromaticity diagram $\mathcal{D}$, of the plane $L + M + S = 1$. Then the cone $\mathcal{C}_\mathcal{D}$ generated by $\mathcal{D}$ is also convex, and the image of $\mathcal{C}_\mathcal{D}$ under any von Kries transformation (indeed, under any linear transformation) is itself a convex cone. The intersection of that convex cone with $L + M + S = 1$ is then convex, too.

In the current case, we applied T_I directly to the convex chromaticity diagram in Figure 5, and produce $\mathcal{S}$. Instead of applying T_I directly, however, we could have taken the more roundabout route of extending the diagram to a convex cone, applying the three-dimensional von Kries transform to produce a new convex cone that was somewhat outside the first cone, and then intersecting the new cone with $L + M + S = 1$. The result would still be the set $\mathcal{S}$, but the roundabout route shows that $\mathcal{S}$ must be convex. In general, we can see that a von Kries transform preserves convex structures in the two-dimensional chromaticity plane as well as in three-dimensional colour space.

(A von Kries transform in two dimensions is a projective rather than a linear transformation, and projective transformations do not always preserve convex structures. In this case, though, the von Kries matrix is diagonal, and all its eigenvalues are positive. This form, combined with the fact that we are working only in the positive LMS octant, guarantees that convexity is preserved.)

Since convex structure is preserved, we can calculate $\mathcal{S}$ from only a few points. The lm chromaticity diagram is the convex hull of its vertices, which are just the monochromatic points on the spectrum locus. Since T_I preserves convexity, the image of the diagram under T_I is just the convex hull of the images of its vertices. This argument justifies the form of $\mathcal{S}$ seen in Figure 5, which is just a polygon joining $T_I(400 \text{ nm})$ to $T_I(410 \text{ nm})$ to $T_I(420 \text{ nm})$, and so on through the visible wavelengths. This technique of evaluating convex sets only at their vertices simplifies calculations for von Kries transforms.

3.3 The Total Chromaticity Diagram

The example in the previous section necessitated adding new chromaticities to the diagram $\mathcal{L}$. While that example considered just one von Kries transform, this section will consider the set of all von Kries transforms simultaneously. The result will be a total chromaticity diagram $\mathcal{T}$ that contains any chromaticity that could result from any state of chromatic adapatation.

Such a diagram is the union of of all the sets like $\mathcal{S}$, where the illuminant that generates $\mathcal{S}$ is free to vary over all of $\mathcal{D}$. Since the original chromaticity diagram is the convex hull of the spectrum locus, and since the von Kries transform preserves convex structures, it follows that $\mathcal{T}$ is the convex hull of the von Kries images of the chromaticities on the spectrum locus. Figure 6 shows the images for some selected wavelengths on the spectrum locus, while Figure 7 shows the convex hull of all the images. The latter figure is the total chromaticity diagram $\mathcal{T}$.

In Figure 7, the standard chromaticity diagram $\mathcal{L}$ is shaded a dark grey, while the total chromaticity diagram $\mathcal{T}$ is shaded a light grey. The total chromaticity diagram contains the standard chromaticity diagram as a subset, and extends considerably beyond it. While the dotted line marks the boundary we calculated for $\mathcal{T}$, the previous discussion shows that that boundary is not known with much precision. A simple triangle, with vertices $(0.00, 0.00)$, $(0.90, 0.10)$, and $(0.15, 0.85)$, can therefore be taken as a handy representation of $\mathcal{T}$.

A natural step at this juncture is to convert Figure 7 from LMS to XYZ coordinates, and view $\mathcal{T}$ as a superset of the familiar xy chromaticity diagram of Figure 2. Unfortunately,

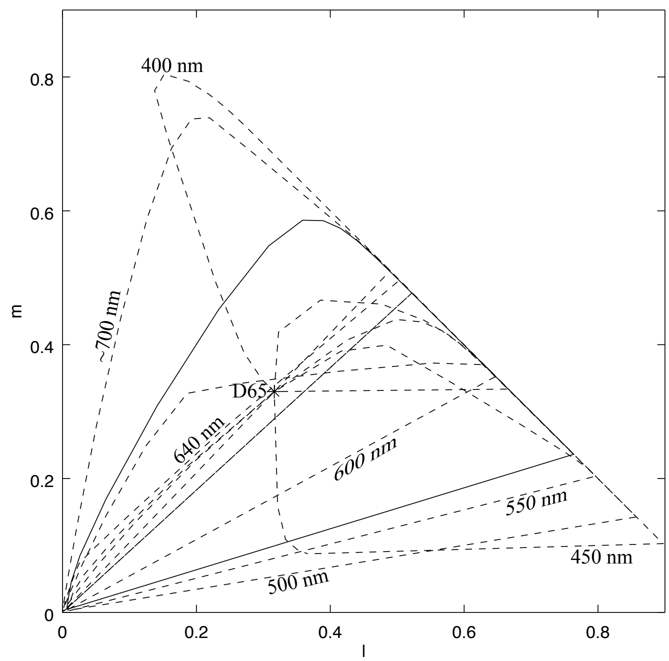

Figure 6: Images of the von Kries Transform for Various Wavelengths

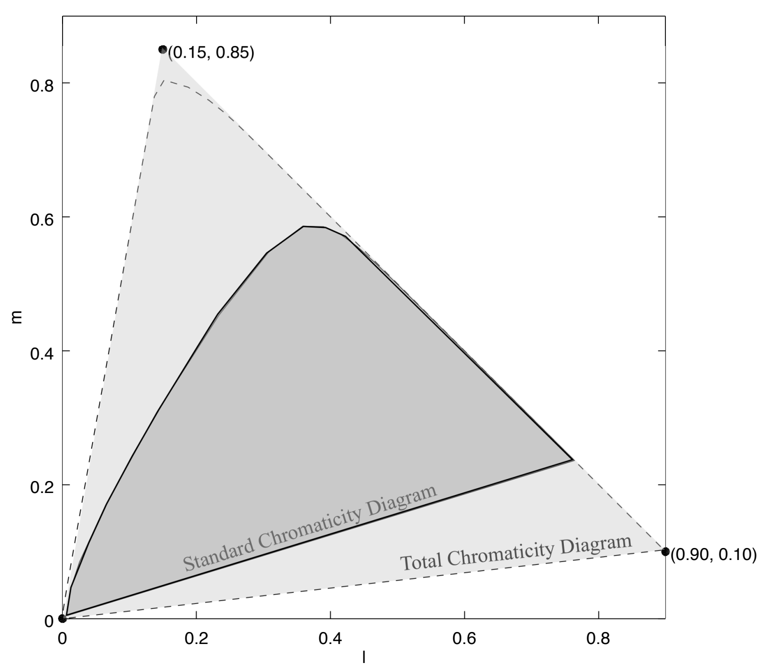

Figure 7: Image of the von Kries Transform of the Standard Chromaticity Diagram

the resulting figure is not well-defined. To see why, convert the vertex $(0.15, 0.85)$ from lm to xy coordinates. First, find a three-dimensional LMS vector with chromaticity coordinates $(0.15, 0.85)$. Any one will give the same conversion, so choose $(0.18, 1.00, 0.00)$. Using the HPE matrix to convert $(0.18, 1.00, 0.00)$ to XYZ coordinates gives $(-0.78, 0.69, 0.00)$. At this point, the problem becomes apparent: the X value is negative, which means that the vector as a whole is outside the positive octant, and the (non-negative) ray generated by that vector will not intersect the chromaticity plane $X + Y + Z = 1$. Recall that one reason for choosing the plane $X + Y + Z = 1$ was that it crosses every ray of the XYZ spectrum cone. Now that chromaticity space has been expanded, the cone has also been expanded, so much so that some rays do not cross $X + Y + Z = 1$, and therefore do not define any xy chromaticity coordinates. Thus $\mathcal{T}$ must stay in lm coordinates.

3.4 Caveats

The derivations just presented draw conclusions in a valid manner from their premises, but it should be remembered that the premises themselves have some uncertainty, inherited from their empirical underpinnings. Revising these premises could modify our calculations and produce a slightly different $\mathcal{T}$, Likely these changes would be insignificant, especially since the boundaries of the diagram are rarely reached in practice, but they should be noted for further investigations.

The first such premise is the use of D65 as our "neutral" illuminant. While a reasonable choice, one could just as reasonably choose many nearby illuminants, which might give somewhat different diagrams.

The second premise is the von Kries model itself. While this model has been used and investigated for over a century, it still is not clear how accurate it is (see Chap. 9 of Ref. 7). At the least, it usually provides a good approximation for commonly encountered cases. The cases in this paper, however, extend well beyond that region to highly artificial cases such as monochromatic illuminants, and the model might break down at these extremes. Apart from the von Kries hypothesis, many other chromatic adaptation models have been proposed. This paper's calculations could easily be repeated with another model substituted for the von Kries model, and their total diagrams compared.

The third premise is the use of the HPE matrix to convert between XYZ and LMS coordinates. Many matrices have been suggested for this conversion. Another natural candidate is M_{CAT02} (see Sect. 9.7 of Ref. 2). While M_{CAT02} has been thoroughly vetted and selected by the CIE for its latest colour appearance model, it has the disadvantage of occasionally producing negative LMS coordinates. On the other hand, it has the advantage of not producing any coordinates of value 0, which the HPE matrix occasionally does; we saw earlier that the von Kries model could not be applied to these cases. Overall, no unequivocally correct conversion is known, which inevitably leads to some uncertainty in the calculation of $\mathcal{T}$.

If the premises mentioned here are some day replaced with more certain ones, then a more accurate $\mathcal{T}$ could be calculated, but until then Figure 7 should be a good practical approximation.

4 High Chromas in the Helson-Judd Effect

The Helson-Judd effect[2,8,9] occurs in a room which is illuminated solely by monochromatic light. After adapting to that environment, subjects view a set of neutral Munsell samples, whose reflectance spectra are non-selective (i.e. fairly flat), against a simple background. Under these conditions, many subjects report that dark neutral samples take on a colour complementary to the illuminant's colour, sometimes with chromas that seem impossibly high, as if a glowing light were superimposed on them.

The total chromaticity diagram might explain these apparently impossible perceptions. Suppose, for instance, that the illumination has wavelength 450 nm. Then the von Kries model predicts that all chromaticities will appear in the irregular region labeled 450 nm in the bottom right of Figure 6. While some of this region overlaps with the standard chromaticity diagram, a fair bit of it is also outside. Saturation, which is often used as a rough approximation for perceptual chroma, is zero for practical purposes near white illuminants like D65, and increases steadily as one moves towards the spectrum locus. Since the region resulting from 450 nm extends considerably beyond the spectrum locus in some directions, the saturations in those directions could plausibly increase to levels beyond those encountered in normal illumination. In fact, they could represent colour perceptions that the subjects had never experienced before. Almost by definition they would be beyond the limits found in colour systems like Munsell or NCS, which are standardized on neutral illuminants. As shown in Figure 6 of Ref. 10, an unnaturally saturated shadow colour will sometimes seem to glow or emit light, which could explain the glow that the subjects report.

While many other aspects of the Helson-Judd effect remain mysterious, the total chromaticity diagram, and the von Kries model on which it is based, might suffice to explain the high chromas that are sometimes reported. A monochromatic source will produce as extreme a von Kries transform as possible. The boundary of the von Kries image of a monochromatic source goes through the neutral region, and the image is limited to one side of $\mathcal{L}$—the side whose hues are complementary to the source. As a result, only complementary hues can be seen. The neutrals themselves, such as the Munsell greys that observers view, are shifted towards and maybe beyond the spectrum locus; they thus become rather saturated, perhaps beyond any level that observers have viewed previously.

5 Summary

This paper has derived a surprising conclusion: what we call *the* chromaticity diagram is really just *a* chromaticity diagram, that corresponds to a neutral state of adaptation. Every state of adaptation has its own chromaticity diagram, which is a convex subset of the chromaticity plane $L + M + S = 1$. An illuminant defines a state of adaptation via a von Kries transform, and that state's chromaticity diagram is the image of the standard lm chromaticity diagram under that transform. The union of all these diagrams, over the set of all possible illuminants, defines the total chromaticity diagram $\mathcal{T}$. To a good approximation, $\mathcal{T}$ is a triangle that contains the standard diagram and extends considerably beyond it. The colours of $\mathcal{T}$ outside the standard diagram cannot be produced under neutral illumination. Some have probably never been seen by most observers, and would likely appear impossibly saturated if they were seen. In the case of purely monochromatic illumination, the Helson-Judd effect, which can be modeled as an extreme case of von Kries adaptation, sometimes produces such unusually saturated colours.

References

1. Johannes von Kries, "Chromatic Adaptation," Festschrift der Albrecht-Ludwig-Universität, pp. 145-158, 1902, translated and reproduced in David L. MacAdam, *Sources of Colour Science*, MIT Press, Cambridge, MA, 1970.
2. Mark D. Fairchild. *Color Appearance Models, 2nd ed.,* John Wiley & Sons, Ltd, 2005.
3. G. Wyszecki and W. S. Stiles, *Color Science: Concepts and Methods, Quantitative Data and Formulae, 2nd ed.*, John Wiley & Sons, 1982.
4. Deane B. Judd, "The 1931 I. C. I. Standard Observer and Coordinate System for Colorimetry," JOSA, Vol. 23, pp. 359-374, October 1933.
5. Paul Centore, *The Geometry of Colour*, 221 pages, 2017.
6. S. R. Lay, *Convex Sets and Their Applications*, Dover Publications, 2007.
7. Brian A. Wandell, *Foundations of Vision*, Sinauer Associates, Sunderland MA, 1995.
8. Harry Helson, "Fundamental Problems in Color Vision. I. The Principle Governing Changes in Hue, Saturation, and Lightness of Non-Selective Samples in Chromatic Illumination," Journal of Experimental Psychology, Vol. 23, No. 5, November 1938.
9. Deane B. Judd, "Hue, Saturation, and Lightness of Surface Colors with Chromatic Illumination," Journal of Research of the National Bureau of Standards, Vol. 24, Research Paper RP1285, March 1940. (This paper also appeared in the Journal of the Optical Society of America, Vol. 30, Issue 1, pp. 2-32, 1940.)
10. Paul Centore, "Shadow Colours for Painters," Journal of the International Colour Association, Vol. 11, pp. 18-27, 2013.

