Local Linear Models for Improved von Kries Adaptation

G. D. Finlayson, A. Alsam, and S. D. Hordley
School of Information Systems
University of East Anglia
Norwich NR4 7TJ
United Kingdom

graham, steve@sys.uea.ac.uk

Abstract

In this paper we explore the conditions under which the von Kries model of colour constancy is exact, an investigation motivated by the fact that in practice the model has been shown to work well for a wide range of imaging devices [1] despite the fact that existing theory [2] predicts that it should perform poorly. We present a modified theory which reconciles this apparent contradiction and which is based on the observation that von Kries adaptation treats sensor responses independently of one another. Starting from this point we show how to recover, for a single sensor, set of surfaces, and reference illuminant, the set of *von Kries illuminants*: all lights for which von Kries adaptation is a perfect model of illumination change. To help us in this task we use a linear model of surface reflectance, but importantly, we use a *local* model: that is, a model derived by examining reflectance only in the region to which the sensor of interest is sensitive. Adopting such a model and treating sensors independently of one another we show that our new theory accurately predicts the good practical performance of the von Kries model.

1. Introduction

Visual perception begins when light entering the eye is focused by the lens onto the surface of the retina. On the retina's surface are a great many light sensitive *cone* cells, each classified into one of three types, differentiated by the fact that they respond more or less strongly to light energy at different wavelengths of the electromagnetic spectrum. Typically the three types are referred to as long- medium- and short- wavelength sensitive cells and their responses to a given light stimulus are denoted as a triplet (l,m,s). It is this triplet of cone responses which forms the basis of our colour perception.

Since the response of these cells depends on the relative energy in the light incident upon them, it is clear that a change in illumination will result in a change in the cone cell responses. For example, when we look at a white piece of paper indoors under a tungsten light, or outdoors under bluesky daylight the relative energy of the sources is different and by implication so too are the cone responses to it. And yet, given two different stimuli and differing cone responses we perceive the paper to be a constant white in both cases. This phenomenon is known as *colour constancy* and it is the subject of an ongoing debate as to whether (and if so, to what extent) our own visual system is really colour constant. In addition, the question as to how a general visual system (our own, or computer) might achieve colour constancy is also the subject of much research.

One of the most widely adopted models for colour constancy is the so called *von Kries* model of adaptation in which it is proposed that constancy is achieved by an independent scaling of the three different types of cone responses such that the scale factors account for changes in the nature of the prevailing illumination. Let (l_1,m_1,s_1) represent the responses to a surface under one light and (l_2,m_2,s_2), the corresponding responses under a second light. Then in the *von* Kries model the two are related by:

$$l_1 = \alpha l_2,\ \ m_1 = \beta m_2,\ \ s_1 = \gamma s_2 \tag{1}$$

The model was first proposed by *von* Kries [3] whose work on asymmetric colour matching led him to propose that such an adaptation mechanism exists in our own visual processing.

The extent to which our visual system is colour constant and the question as to whether or not the *von* Kries model is the mechanism by which this constancy is achieved is not the focus of this paper. Rather, we are interested in exploring the theoretical limits of the model itself. Specifically we set out to reconcile an apparent contradiction which has arisen with respect to the model. West and Brill [2] have previously investigated the conditions under which *von* Kries adaptation supports perfect colour constancy and their analysis predicts that the model will be a poor one for the type of illuminants encountered in the real world. Yet the poor performance predicted by their analysis is at odds with our experience of using the model in practice. Part of the explanation for this can be found in the work of Finlayson *et al* who have shown [1, 4] that an important factor affecting the performance of the model is the properties of the sensors to which it is applied. In particular if sensors are narrowband (they respond to light energy in only a restricted region of the visible spectrum) then von Kries models illumination change well. Moreover, they showed that even when sensors are not narrowband it is almost always possible to find a linear transform of them which results in a set of sensors which are more narrowband and for which the von Kries model is in turn a good model of illumination change.

But while Finlayson *et al*'s work provides empirical support for the model it does not provide a theory to explain why the model should work. In this paper we provide a new theoretical analysis which bridges the gap between the pessimistic predictions of West and Brill's work and the good performance which can be obtained in practice. Our analysis is based on a re-formulation of West and Brill's original work in which we pose, and answer the following question. Given a triplet of sensor sensitivities, a set of surface reflectance functions and an illuminant spectral power distribution, what are the corresponding set of illuminants for which the *von* Kries model affords perfect colour constancy? We solve for this set of illuminants and show that, as West and Brill's original work predicts, for typical situations, von Kries' model is a poor one. To reconcile the mismatch between this theory and the good practical performance of the model we make a further important modification to West and Brill's analysis. Specifically, we solve not for the set of von Kries invariant illuminants with respect to a triplet of sensor responses, but instead, we solve for a set of invariant lights with respect to a single sensor. That is, we consider sensors independent of one another (as von Kries' model itself does) and we show that this modification leads to a much improved theoretical performance which matches well the good performance we achieve in practice.

We begin (§2) by defining a simple model of image formation and we explore the issue of illuminant change in the context of this model. We then present (§3) a modified form of West and Brill's analysis original analysis to determine a set of *von* Kries illuminants for a trichromatic device. In § 4 we show that this analysis predicts that *von* Kries adaptation is a poor model of illumination change, a fact which motivates our work. Our modified analysis, for a *monochromatic* sensor, is presented in § 4 and the practical import of this analysis is explored in § 5.

2. Image Formation and Illumination Change

We adopt a simple model of image formation founded on the interaction of three different factors: light, surface and sensor. A light is characterised by $E(\lambda)$, its spectral power distribution (SPD) which defines how much energy it emits at each wavelength. A surface is represented by $S(\lambda)$, its surface reflectance function defining what proportion of light incident upon it is reflected on a per-wavelength basis. Finally the ith sensor of a device or visual system is denoted by its spectral sensitivity function $Q_i(\lambda)$ which characterises what fraction of the light incident upon it the sensor absorbs, again on a per-wavelength basis. The response of a sensor, which we denote ρ_i can thus be expressed as:

$$\rho_i = \int_\omega E(\lambda)S(\lambda)Q_i(\lambda)d\lambda \tag{2}$$

where ω represents the range of wavelengths for which the device has non-zero sensitivity. Our own, and most other visual systems are trichromatic and thus light from any point in a scene is represented as a triplet of sensor responses which we denote $\underline{\rho}$.

For the purposes of the derivations which follow, we approximate light, surface, and sensor by their values at a set of n discrete sample points across the range of the visible spectrum so that Eq. (2) becomes:

$$\rho_i = \sum_{k=1}^{n} E(\lambda_k)S(\lambda_k)Q_i(\lambda_k)\Delta\lambda \tag{3}$$

where the scalar $\Delta\lambda$ accounts for the size of the sampling interval. Or dropping the dependence on wavelength (λ) and representing our discrete representations of light, surface, and sensor as vectors we obtain:

$$\rho_i = S^T diag\,(E)\,Q_i \tag{4}$$

where $\Delta\lambda$ has been incorporated into Q_i and $diag$ is the diagonal operator which transforms its $n \times 1$ vector argument into an $n \times n$ diagonal matrix, whose non-zero entries are the entries of the vector.

A further simplification which is often made is to represent lights, surfaces, or both by low-dimensional linear models, a simplification justified by statistical analyses of large sets of typical lights [5] and surfaces [6]. In this paper we will make use of such a representation for surfaces so that any surface S is represented:

$$S(\lambda) = \sum_{k=1}^{M} \sigma_k S_k(\lambda) \tag{5}$$

where M, the number of basis functions, is typically much less than n, the number of sample points used to represent lights and surfaces. The image formation equation now becomes:

$$\rho_i = ([S_1 \dots S_M]\underline{\sigma})^T \, diag\,(E)\, Q_i \tag{6}$$

where $[S_1 \dots S_M]$ is a matrix whose columns are the basis functions of the linear model. Such a representation is helpful since it can lead to a simplified model of illumination change. For example, it can be shown [7] that if surface reflectance is 3-dimensional, then sensor responses to the same surface under two different lights are a linear transform apart:

$$\underline{\rho}^{E_1} = T\underline{\rho}^{E_2} \tag{7}$$

where T is a 3×3 matrix. The *von* Kries model of adaptation simplifies this relationship even further, proposing that corresponding sensor responses under two different illuminants are related by a diagonal matrix D:

$$\underline{\rho}^{E_1} = D\underline{\rho}^{E_2} \tag{8}$$

A slightly more general model of illumination change, is the so called generalised coefficient proposed by Finlayson *et al* [1]. In this model illumination change is again modelled by a diagonal transform, but sensor responses are first transformed by a fixed 3×3 linear transform prior to this. In this case responses under different illuminants are related:

$$T_s\underline{\rho}^{E_1} = DT_s\underline{\rho}^{E_2} \implies \underline{\rho}^{E_1} = T_s^{-1}DT_s\underline{\rho}^{E_2} \tag{9}$$

The matrix T_s can be seen as a transformation of the original sensors to some new set of sensors. Finlayson *et al* set out to derive transforms such that the resulting new set of sensors better supported *von* Kries adaptation. That is, they derived sensors for which Equation (8) is a good model of illumination change. They showed that improved *von* Kries adaptation could be achieved by "sharpening" sensors, that is, by finding a transform T_s resulting in a set of sensors which are responsive only in a restricted region of the visible spectrum, an approach motivated by the fact that in the limiting case of sensors responsive to light at only a single wavelength, *von* Kries adaptation affords perfect colour constancy.

The remainder of the paper aims to provide a theoretical justification of the von Kries model which explains its good practical performance. We begin with a re-formulation of the work of West and Brill who have previously considered the necessary and sufficient conditions for the model to be exact.

3. West and Brill's Analysis

West and Brill [2] set out the mathematical conditions for the von Kries model (Equation 8) to be exactly invariant to a change in illumination. Their analysis consists of two main cases: first, illumination is represented as a finite-dimensional linear model and the authors derive the conditions for a reflectance to be illuminant-invariant with respect to the von Kries model. In a second case the authors adopt a linear model of reflectance and derive the illuminants which are von Kries invariant with respect to this model of reflectance and some reference illuminant. Here we present our own solution to the second case.

We begin by adopting a linear basis of surface reflectance functions as defined by Equation (5). Now, let $\underline{\rho}_k^c$ represent the triplet of responses of a device to the kth basis function $S_k(\lambda)$ when viewed under a reference illuminant c and let $\underline{\rho}_k^o$ be the corresponding response under an arbitrary illuminant o. It can be shown that if responses $\underline{\rho}_k^o$, $(k = 1 \dots M)$ are related to responses $\underline{\rho}_k^c$, $(k = 1 \dots M)$ by a *von* Kries transform then so too are the responses to any other surface representable in the basis. This follows from two facts: first that any such surface is a linear combination of basis functions and so, by the linear nature of image formation, any resulting response is that same linear combination of the responses to the basis functions, and second the fact that if responses are related by a diagonal transform then so too are linear combinations of these responses.

So, considering just the basis functions, let $\underline{q}_i^c$ be a $M \times 1$ vector representing the response of the ith sensor to the basis functions when viewed under illuminant c. That is, for a trichromatic device:

$$\begin{aligned} \underline{q}_1^c &= [S_1 \dots S_M]^T \, diag(Q_1)E^c \\ \underline{q}_2^c &= [S_1 \dots S_M]^T \, diag(Q_2)E^c \\ \underline{q}_3^c &= [S_1 \dots S_M]^T \, diag(Q_3)E^c \end{aligned} \tag{10}$$

or

$$\underline{r}^c = B^T E^c \tag{11}$$

where $\underline{r}$ is the three vectors $\underline{q}_i^c$, $i = 1 \dots 3$, stacked one on top of the other and B^T, the three matrices $[S_1 \dots S_M]^T diag(Q_i)$, $i = 1 \dots 3$, stacked in the same way.

Now, for an arbitrary illuminant to be a *von Kries illuminant* with respect to the reference light the following must be true:

$$\underline{\rho}^o = D\underline{\rho}^c$$

for some 3×3 diagonal matrix D. That is, responses under the two lights must be related by a *von* Kries transform. Now, any diagonal matrix has three degrees of freedom and so can be expressed as a linear combination of any three linearly independent basis matrices:

$$D = a_1D_1 + a_2D_2 + a_3D_3 \tag{12}$$

and thus any *von* Kries illuminant can in turn be expressed as the same linear combination of illuminant spectra which are solutions to the equations:

$$\underline{r}^i = B^T E^i, \quad i = 1 \ldots 3 \tag{13}$$

where the elements of $\underline{r}^i$ are related to $\underline{r}^c$ by one of the matrices D_1, D_2, or D_3. For each value of i, Equation (13) represents a linear system of equations of $3M$ equations in n unknowns. In general $3M < n$ (the number of basis functions, M is less than the number of sample points n divided by 3) which implies the system is underconstrained and thus has either zero or infinitely many solutions [8].

In the case that the system has solutions, any solution can be expressed as a particular solution to Equation (13) plus a solution to the corresponding homogeneous system [8]:

$$\underline{0} = B^T E^o \tag{14}$$

This homogeneous system is independent of $\underline{r}^i$ and thus independent of the basis matrices D_i. So to find all *von Kries* illuminants we must find three particular solutions corresponding to the three sets of equations in (13) together with the solutions to Equation (14). That is, any *von* Kries illuminant can be expressed as:

$$E^o = a_1 E^1 + a_2 E^2 + a_3 E^3 + E^b \tag{15}$$

A particular solution to Equation (13) can be found: $E^o = \left(BB^T\right)^{-1} B\underline{r}^i$. The illuminant E^b, a solution to the homogeneous system is any illuminant in the space orthogonal to the columns of B, a space which West and Brill referred to as the forbidden subspace. In physical terms E^b represents an illuminant which elicits zero response from the sensors.

4. Improved von Kries Adaptation

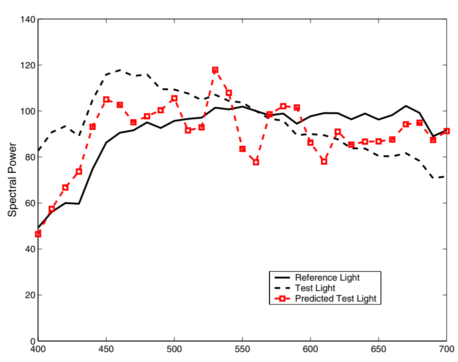

Figure 1: D50 Reference illuminant (solid line) and test illuminant (dashed line) together with the closest von Kries illuminant to the test light (dashed line, with squares).

Unfortunately (and in contradiction to empirical observations) West and Brill's analysis predicts that for real world lights, von Kries is a poor model of illumination change. Figure 1 illustrates this fact. For A D65 daylight illuminant (solid line in Figure 1) we determined the set of von Kries illuminants using the analysis above and a 3-d linear model of reflectance. Then we investigated whether a second (D50) daylight (the dashed line in Figure 1) was in this set. The answer is no: the dashed line with square markers in Figure 1 is the closest von Kries illuminant to the second light. That is, it represents the best approximation to the light using the basis expansion defined by Equation (15). Clearly the error in this approximation is large and yet were we to use a von Kries model to relate responses under the pair of illuminants we would find that the accuracy of the model is good. Thus there is a mismatch between theory and practice.

We reconcile this mismatch by reconsidering the analysis above. In particular we note that the von Kries model treats sensors independently of one another and so, rather than considering all three sensors together as West and Brill did, we instead consider each sensor in isolation and derive, a set of von Kries illuminants on a per-sensor basis. To begin, let us consider a single sensor Q_i from some arbitrary imaging device. In theory this sensor can have sensitivity across the whole range of the visible spectrum but in practice most sensors are sensitive across only a quite limited range of the spectrum. What is more, the work of Finlayson *et al* [1] has shown that improved *von* Kries adaptation is obtained by deriving a set of sharpened sensors which are sensitive in only a restricted region of the spectrum. Thus, without loss of generality we assume that an arbitrary sensor is sensitive only in a sub-region of the visible spectrum.

For this arbitrary sensor we would like to derive the set of von Kries invariant illuminants and, like West and Brill, we do so by first adopting a linear model of surface reflectance. However, we introduce here an important modification to the usual definition of a linear model which reflects the fact that we are treating sensors independently of one another. Suppose that our sensor is *active* (has non-zero responsivity) in a wavelength range $\lambda_s \leq \lambda \leq \lambda_e$. It follows that when modelling surface reflectances we are interested only in their behaviour within this range of wavelengths since variations outside this range do not affect the sensor.

Thus we define a *local* linear model of surface reflectance with respect to an arbitrary sensor by:

$$S(\lambda) = \sum_{k=1}^{N_i} \sigma_k S_k(\lambda), \quad \lambda_s \leq \lambda \leq \lambda_e \tag{16}$$

The basis functions themselves can be derived using the same techniques of statistical analysis set out by a number of previous authors [9, 6]. But importantly the derived basis functions capture the variation in reflectances across the region of the spectrum pertinent to the sensor of interest. The, number and nature of these basis functions will vary depending on the active range of the sensor Q_i and importantly we have found that for the linear models derived locally within the active region of a range of typical sensors 3-basis vectors are sufficient to capture 99% or more of this local variation whereas to capture similar variance using a global model needs 5 or more basis vectors.

Given this sensor specific model of surface reflectance let us now consider how we can derive the set of von Kries illuminants for the sensor Q_i. As in the West and Brill analysis, to ensure that an illuminant is a von Kries illuminant it is sufficient to ensure that the von Kries model holds for the surface basis functions since the result for an arbitrary surface within this basis follows from the linear nature of image formation. Thus, we begin by defining a reference illuminant E^c and a linear basis $[S_1 \dots S_{M_i}]$ (defined with respect to the active region of Q_i). Given these we can express the response of the sensor to the basis vectors by an $M \times 1$ vector $\underline{r}^c$ thus:

$$\underline{r}^c = [S_1 \dots S_{M_i}]^T \, diag(Q_i) E^c \tag{17}$$

For the case of a single sensor a von Kries illuminant is any illuminant o such that the responses of the sensor under it, $\underline{r}^o$ are related to the responses under the reference light by a simple scale factor:

$$\underline{r}^o = \alpha \underline{r}^c \tag{18}$$

Thus a von Kries illuminant is any illuminant E^o which satisfies the following equation:

$$[S_1 \dots S_{M_i}]^T \, diag(Q_i) E^o = \underline{r}^o = \alpha \underline{r}^c \tag{19}$$

Equation (19) is the single sensor version of Equation (13) and represents a system of linear equations. There are M_i equations, corresponding to the dimension of the surface reflectance basis and the unknowns are the n elements of the illuminant E^o. In general M_i will be strictly less than n and so the system in Equation (19) is underconstrained. Once again we can use basic results of linear algebra [8] to characterise the solutions to this system as the sum of a particular solution to Equation (19) plus a solution to the corresponding homogeneous system:

$$[S_1 \dots S_{M_i}]^T \, diag(Q_i) E^o = \underline{0} \tag{20}$$

By inspection, αE^c is a solution to Equation (19) thus it remains to determine solutions to Equation (20). Algebraically Equation (20) can be interpreted as meaning that E^o is orthogonal to the space defined by the rows of the matrix $P = [S_1 \dots S_{M_i}]^T diag(Q_i)$. Or, in physical terms that E^o is an illuminant which elicits a zero response in the sensor. Solutions to Equation (20) can be determined by finding the null space of the matrix P, which we denote $P^\perp$ and which can be calculated $I - PP^+$ where $^+$ denotes the pseudo-inverse [8] of P. The dimensionality of this space is $n - M_i$ and thus any solution to Equation (20) can be written as a linear combination of the columns of $P^\perp$:

$$E^b = b_1 \underline{p}_1^\perp + b_2 \underline{p}_2^\perp + \dots + b_{n-M_i} \underline{p}_{n-M_i}^\perp \tag{21}$$

From which it follows that any von Kries illuminant can be expressed as:

$$E^o = \alpha E^c + E^b = \alpha E^c + b_2 \underline{p}_2^\perp + \dots + b_{n-M_i} \underline{p}_{n-M_i}^\perp \tag{22}$$

Both α and the elements $b_1, b_2, \dots b_{n-M_i}$ can be arbitrary scalars so that Equation (22) tells us that there are an infinite set of von Kries illuminants with respect to any given reference illuminant, surface reflectance basis, and sensor.

Such an analysis can be performed for any arbitrary sensor, thus for a trichromatic device we can derive three sets of von Kries illuminants, one set for each sensor. In the next section we investigate the practical implications of this theory.

5. Experimental Performance

We present preliminary results here for a set of sharpened *CIE* colour matching functions. These functions were obtained by first deriving a sharpening transform T_s, via the method of data based sharpening described in [1]. We then transformed the CIE colour matching functions by this transform which results in a new sensor set. The sharpening transform results in sensors which have regions of negative sensitivity so for the purposes of these experiments we ignored negative values by simply clipping them to zero (we note that the responses resulting from these clipped sensors are very close to those obtained using the set with negative response). For each of these sensors we derived a separate linear surface reflectance basis by considering variations in reflectance in the active region of each of the three sensors. These regions are approximately $400-550nm$, $460-610nm$, and $540-700nm$.

To derive the bases we performed a principal component analysis of a representative set of surface reflectance functions, specifically we used reflectances from a Macbeth colour checker and an Agfa IT8 72 chart. These local bases resulted in an improved fit in each region of interest and in each case a 3-d model was sufficient to capture approximately 99% of the variance as compared to 97% for a global model with 3 basis vectors. For each sensor, we used these local 3-d linear models of reflectance to derive the set of von Kries illuminants by the method set out above. In addition we used the method of West and Brill (with a global 3-d reflectance basis) detailed in Section 3 to derive a fourth set of von Kries illuminants.

In each case we took as our reference illuminant D65 daylight illumination To assess whether von Kries is a good model of illumination change we considered a second, test, illuminant and asked whether this illuminant was in each of the derived sets of von Kries illuminants. That is we determined the closest von Kries illuminant to the second illuminant using each of the four recovered sets of von Kries illuminants. An example of the results for the case when the test illuminant is D50 are shown in Figures 1 and 2. Note that the sets derived using only a single sensor recover an illuminant SPD (illustrated in Figure 2) only in the active region of the sensor since outside this region the sensor has zero response and so is unaffected by variations in illuminant or reflectance. In contrast the method of West and Brill results in an illuminant SPD (illustrated in Figure 1) across the full extent of the visible spectrum. The results in Figures 1 and 2 reveal two important and contradicting facts. First, the Brill and West analysis results in an illuminant which is far from the actual test illuminant, suggesting that von Kries will be a poor model of adaptation for this pair of illuminants. Second, and by contrast the illuminants derived using the single sensor analysis set forth in this paper are very close to the actual illuminant in each of the three regions. This suggests that von Kries adaptation should afford good colour constancy for this pair of illuminants. The latter theory is supported by an empirical study of how sensor responses change across these two illuminants, which reveals that von Kries adaptation is indeed a good model of the illuminant change. Of course the method we have set out results in one illuminant per sensor but in practice we require a single illuminant to provide good von Kries adaptation for all three sensors simultaneously. But in this case we can simply combine the three separate illuminants by an appropriate interpolation scheme to derive a single light. We note doing so will result in an illuminant which no longer affords perfect colour constancy but its accuracy will still be good, and importantly the recovered illuminant will be close to the test light.

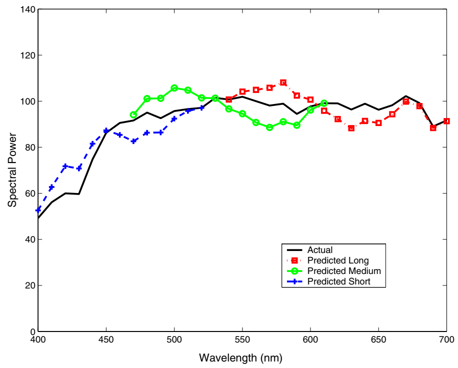

Figure 2: *Test Light (D50) together with the three local SPDs predicted using a sensor by sensor analysis.*

Finally we note that the results presented for this pair of illuminants are typical of the results we achieve for a range of common illuminants. Thus the theory set forth in this paper predicts the good practical performance of the von Kries model and so meets the initial aims of the paper.

References

[1] Graham Finlayson. *Coefficient Colour Constancy*. PhD thesis, Simon Fraser University, 1995.

[2] Gerhard West and Michael H. Brill. Necessary and sufficient conditions for von Kries chromatic adaptation to give color constancy. *Journal of Mathematical Biology*, 15(10):249–258, 1982.

[3] J. von Kries. Influence of adaptation on the effects produced by luminous stimuli. In D. L. MacAdam, editor, *Sources of Color Science*, pages 109–119. The MIT Press, Cambridge MA, 1970.

[4] Graham D. Finlayson and Brian V. Funt. Coefficient channels: Derivation and Relationship to Other Theoretical Studies. *Color Research and Application*, 21(2):87–95, 1996.

[5] D. B. Judd, D. L. MacAdam, and G. W. Wyszecki. Spectral distribution of typical daylight as a function of correlated colour temperature. *Journal of the Optical Society of America, A*, 54:1031–1040, 1964.

[6] D.H. Marimont and B.A. Wandell. Linear models of surface and illuminant spectra. *Journal of the Optical Society of America, A*, 9(11):1905–1913, 92.

[7] Laurence T. Maloney and Brian A. Wandell. Color constancy: a method for recovering surface spectral reflectance. *Journal of the Optical Society of America, A*, 3(1):29–33, 1986.

[8] G.H. Golub and C.F. van Loan. *Matrix Computations*. John Hopkins U. Press, 1983.

[9] Laurence T. Maloney. Evaluation of linear models of surface spectral reflectance with small numbers of parameters. *Journal of the Optical Society of America, A*, 3(10):1673–1683, 1986.