

Positive Bradford Curves through Sharpening

Graham D. Finlayson
School of Information Systems,
The University of East Anglia,
Norwich NR4 7TJ
United Kingdom

Mark S. Drew
School of Computing Science
Simon Fraser University
Vancouver, British Columbia V5A 1S6
Canada

Abstract

The Bradford curves are used in the CIECAM97 colour appearance model, and are ‘sharpened’ in the sense that they have narrower support than the cone fundamentals. Spectral sharpening is a method which finds the linear combination of a set of sensors that is most sensitive to a given interval of the visible spectrum. Here we investigate the relationship between the apparently sharp Bradford curves and the spectral sharpening method (since spectral sharpening was not used to derive the Bradford curves). We find that Bradford curves can be derived using spectral sharpening but the sharpening intervals are not the ones we would have expected or wished for. The Bradford intervals are far from the ‘prime wavelength’ intervals: those parts of the visible spectrum where there is maximal visual sensitivity. Also, independent of any sharpening argument, the Bradford curves are unexpected in the sense that they have some negative sensitivity. Here we address both these concerns and produce sharpened versions of the Bradford curves that are both all-positive and also sharpened within wavelength intervals around the prime wavelengths. In a sense, we are continuing the work of MacAdam, and Pearson and Yule, in forming positive combinations of the colour-matching functions. However, the advantage of the spectral sharpening approach is that not only can we produce positive curves, but the process is ‘steerable’ in that we can produce positive curves with as good or better properties for sharpening within a given set of sharpening intervals. At base, however, it is positive *colours* in the transformed space that are the prime objective. Therefore we also carry out sharpening of sensor curves governed not by positivity of the curves themselves but of colours resulting from them. Curves that result have negative lobes, but generate positive colours. We find that this type of constrained sharpening generates the best results, almost as good as for unconstrained sharpening but without the penalty of negative colours.

1. Introduction

Spectral sharpening is a method of transforming colour camera, scanner, or other optical device multispectral image pixel values into new values that would have resulted from sensors with more narrow-band spectral sensitivities [1]. The utility of such a transform is that for many computer vision and colour image processing algorithms, sharper sensors result in better performance. Consider, for example, the simplest form of colour correction, the von Kries diagonal transform for correcting from RGB values under one illuminant to those under a second illuminant. Theoretical sensors that act as delta functions would exactly obey a diagonal transform, and it was shown in [1]

that spectral sharpening could greatly benefit such a colour constancy strategy.

Spectrally sharpened curves are derived as a matrix transform of the human colour-matching functions. The set of three Bradford curves, used in colour appearance modelling, are themselves also written as a matrix transform of the colour-matching functions, and appear to be ‘sharpened’: they are sensitive to smaller bands of the visible spectrum. Yet, Bradford sensors were not derived using spectral sharpening. Rather, they were the result of fitting psychophysical corresponding colour data [2]. In this paper, we use spectral sharpening to investigate Bradford sensors.

Our first result is to show that there exist three intervals of the visible spectrum with respect to which spectral sharpening delivers sensitivity curves which are close to the Bradford sensors. We observe, however, that the sharpening intervals we discover do not seem to make much sense. From a mathematical standpoint the intervals are poor because the resulting sharpened curves really are not very sharp with respect to these intervals. That is, the derived sensors cannot behave like narrow-band sensors positioned in the sharpening intervals because they are sharp elsewhere! From a practical viewpoint the sharpening intervals are far from the ‘prime wavelength’ regions: those parts of the visible spectrum for which the visual system is maximally sensitive.

Sharpening apart, another potential problem for the Bradford curves is that they have negative lobes (i.e., negative sensitivities). This means that it is possible that certain highly saturated colours could under some lights induce a zero or negative sensor response. If this were to happen then colour constancy algorithms predicated on positive RGB values, e.g. those that utilise a maximum value in an image (cf. [3]) could not be used directly. However, perhaps the most compelling reason for sharpening with positivity is colour ratio stability: if an algorithm uses colour ratios, as in [4], then if values fall close to zero or change sign then colour ratios can vary substantially. As well, if we consider log-based homomorphic filtering for image enhancement, then the advantage of positive values is evident.

The negative lobe problem led us to develop a constrained spectral sharpening [5] that returns sharp sensors that are all positive. This naturally forces RGB colour values in the ‘sharpened’ space to be positive. This work is in a sense the natural completion of that begun by MacAdam, and Pearson and Yule [6]. These authors formed linear combinations of the colour-matching functions, adding

various proportions of the curves until negatives resulted. Here we use a straightforward optimisation technique instead, but in addition make the Pearson-Yule procedure *steerable*, as it were, by also insisting that the optimisation concentrate each curve's 'energy' within a given sharpening interval.

Since at base it is only positive *colours* that are needed, we also carry out an optimisation that sharpens sensor curves subject to the constraint that colours in the transformed space are non-negative. I.e., we do not insist on positive curves, but only positive results. This final approach turns out to produce the best results for constrained sharpening, almost as effective as the best possible, unconstrained sharpening.

2. Bradford Curves and Colour Appearance

Light entering the eye is a function of the surfaces in a scene and the prevailing illumination. If $E(\lambda)$ denotes illumination and $S(\lambda)$ surface reflectance then the reflected light is proportional to $E(\lambda)S(\lambda)$. In modelling the visual response to this light one first calculates its XYZ tristimulus coordinates: $X = \int_{\omega} \bar{x}(\lambda)E(\lambda)S(\lambda)d\lambda$

$$\begin{aligned} Y &= \int_{\omega} \bar{y}(\lambda)E(\lambda)S(\lambda)d\lambda \\ Z &= \int_{\omega} \bar{z}(\lambda)E(\lambda)S(\lambda)d\lambda \end{aligned} \quad (1)$$

where $\bar{x}(\lambda)$, $\bar{y}(\lambda)$ and $\bar{z}(\lambda)$ are the CIE colour matching functions and ω denotes the visible spectrum. It is clear from eq. (1) that the spectral characteristics of $E(\lambda)$ strongly affects the (X, Y, Z) tristimulus values.

Colour appearance models attempt to quantify how different surface colours appear when viewed in different viewing conditions. For example, it is apparent that if we reverse the roles of $E(\lambda)$ and $S(\lambda)$ in (1) we arrive at the same tristimulus values (and perceive the same colour percept): we cannot distinguish between a white wall viewed under red light and vice versa. Of course, surface colours are rarely seen in isolation and the colour for a surface seen in context tends to be perceived as more or less the same colour across contexts (the visual system has colour constancy).

Colour appearance models use a chromatic adaptation transform to model illumination. Suppose $x^{E,S}$ denotes the XYZ tristimulus for a surface $S(\lambda)$ viewed under $E(\lambda)$. We would like to find the tristimulus vector x that induces the same colour appearance under a second illuminant E' . This is calculated using the Bradford adaptation transform:

$$x = T^{-1} D^{E,E'} T x^{E,S} \quad (2)$$

In (2), T is a fixed 3×3 matrix and $D^{E,E'}$ is an illuminant-dependent diagonal matrix. It is important to note that x need not equal $x^{E',S}$, though it will generally be quite similar. For our purposes, (2) can be usefully simplified by premultiplying both sides of eq. (2) by T , yielding

$$T x = D^{E,E'} T x^{E,S} \quad (3)$$

In (3), the linear transform can be thought of as defining new sensor functions $r(\lambda)$, $g(\lambda)$ and $b(\lambda)$ (plotted in Fig. 1):

$$\begin{bmatrix} r(\lambda) \\ g(\lambda) \\ b(\lambda) \end{bmatrix} = \begin{bmatrix} t_{11} & t_{12} & t_{13} \\ t_{21} & t_{22} & t_{23} \\ t_{31} & t_{32} & t_{33} \end{bmatrix} \begin{bmatrix} \bar{x}(\lambda) \\ \bar{y}(\lambda) \\ \bar{z}(\lambda) \end{bmatrix} \quad (4)$$

Denoting the RGB response for $r(\lambda)$, $g(\lambda)$ and $b(\lambda)$ as $\rho = [R \ G \ B]^t$ it follows that:

$$\rho = D^{E,E'} \rho^{E,S} \quad (4)$$




Figure 1: Human colour-matching functions (dashed lines) and Bradford-transformed curves (solid lines).

Equation (4) is very significant. It informs us that the effect of the illumination can be modelled by simple scalars operating individually on each of the R , G and B (the diagonal matrix has only three non-zero terms). In comparison, the relationship between corresponding XYZs is much more complex.

Yet, why should a diagonal matrix model illumination change for the Bradford curves but not for the XYZ functions? One explanation is that the sensitivity of the Bradford curves is concentrated in a small interval of the visible spectrum and it is well known that, the narrower sensors are, the more accurately will a diagonal matrix model illumination change. One might reasonably hypothesise, therefore, that the Bradford sensors are the human visual system's attempt at synthesising narrow-band sensors.

However, this hypothesis can be criticised. First, other work has shown that much more narrow-band sensors might be constructed. In terms of narrow-bandedness the Bradford curves are far from optimal. Second, on examining the Bradford curves more closely one observes that while the curves are more narrow-band they do have significant sensitivity outside these intervals. Moreover, this residual sensitivity may be 'negative' and so the Bradford responses can be driven to 0 or driven towards negative numbers. The issues of optimising narrow-bandedness and maintaining positivity are addressed in the next section, where we mean to derive alternate versions of the Bradford curves that are all-positive (non-negative), while maintaining or improving any benefits that derive from spectral sharpening.

3. Spectral Sharpening: Relation to the Bradford Curves

3.1. Spectral Sharpening
Spectral sharpening means choosing three specific 'sharpening intervals' ψ within the visible spectrum, in which we would like energy in sensor curves to be concentrated. If the visible spectrum consists of wavelengths ω , then our objective is to decrease the amount of energy for wavelengths $\phi = \omega - \psi$ outside a desired sharpening interval ψ .

Suppose in general that there are p sensors (e.g., $p = 3$), and an $s \times p$ matrix Q of sampled sensor values. E.g.,

s might be 31 if we sample between 400nm and 700nm at a 10nm interval. We may choose a different sharpening interval for each of the p sensors and hence carry out a separate minimisation for each of the p colour channels.

Thus spectral sharpening consists of finding a p -component vector c that minimises the least squares summation

$$\min_{\lambda \in \phi_k} \sum [Q(\lambda)c]^2 + \mu \left\{ \sum_{\lambda \in \omega} [Q(\lambda)c]^2 - 1 \right\} \quad (5)$$

for $k = 1..p$ where μ denotes a Lagrange multiplier that ensures the resulting derived sensor has unit length in the L_2 norm. Let us define an operator Δ_α that picks out wavelength indices in the sharpening interval α within any sum. E.g., the operator Δ_{ψ_k} picks out wavelength indices in the sharpening interval ψ_k . Using this operator, it is further useful to define a $p \times p$ matrix involving the summation

$$\Lambda(\alpha) = \sum_{\lambda \in \alpha} Q^t(\lambda)Q(\lambda) = Q^t \Delta_\alpha Q \quad (6)$$

Then taking partial derivatives with respect to unknown vector c and equating to the zero vector produces the Euler equation, which can be written

$$\Lambda(\phi_k)c - \mu[\Lambda(\omega)c] = 0 \quad (7)$$

Note that here $\Lambda(\omega)$ is just $Q^t Q$.

Differentiating (5) with respect to μ simply sets the scale of the resulting sensor to unity in the L_2 norm:

$$\sum_{\lambda \in \omega} [Q(\lambda)c]^2 = c^t \Lambda(\omega)c = 1 \quad (8)$$

Rearranging this equation, we see that solving for c (and consequently the sharpened sensor) is an eigenvector problem:

$$[\Lambda(\omega)]^{-1} \Lambda(\phi_k)c = \mu c \quad (9)$$

There are p solutions of the above equation, each solution corresponding to a stationary value, so we choose the eigenvector which minimises $\sum_{\lambda \in \phi_k} [Q(\lambda)c]^2$. It is straightforward to show [1] that c derived in this way is always a real-valued vector.

3.2. Sharpening Gives the Bradford Curves

Spectral sharpening applied to the human colour-matching functions creates curves with negative lobes such as appear in the Bradford curves. An interesting question arises, therefore: if the Bradford curves are purported to be ‘sharpened’ in some sense, is it possible to write them explicitly as the sharpened versions of the colour-matching functions? And, if so, then what are the sharpening intervals to which the Bradford curves correspond?

The all-non-negative 1931 2° colour matching functions $\bar{x}(\lambda)$, $\bar{y}(\lambda)$, $\bar{z}(\lambda)$ making up an $s \times 3$ matrix X , along with the Bradford-transformed curves making up a similar matrix Q_B , are shown in Fig. 1 (normalised to unity in the L_1 norm, so that each column sums to 1).

In terms of X , the Bradford curves Q_B are given by a matrix transform $Q_B = X B$ (10)

with 3×3 Bradford transform matrix $B = T^t$. If the Bradford curves are indeed close to being spectrally sharpened combinations of the colour matching functions X , then for each colour channel we should find that the corresponding column of the Bradford transform matrix B satisfies the eigenvalue problem (9).

Here, matrix $\Lambda(\omega)$, representing a summation over the entire visible spectrum, ω is given by

$$\Lambda(\omega) = X^t X$$

A priori, we do not know the *sharpening intervals* ψ_k , $k = 1..3$ and the main task is to determine these. Thus, for k in $1..3$, we mean to find the sharpening interval ψ_k , or its corresponding interval $\phi_k = \omega - \psi_k$, and hence *find the matrix* with components

$$\Lambda(\phi_k)_{mn} \equiv \sum_{i \in \phi_k} X_{im} X_{in}$$

that best satisfies (9):

$$\text{find } \phi_k \ni [\Lambda(\omega)]^{-1} \Lambda(\phi_k) b_k \simeq \mu b_k \quad (11)$$

where b_k is the k th column of the transform matrix B . This amounts to choosing a set of sharpening intervals so that the resulting matrix on the left hand side of eq. (11) is closest to being a matrix that has one of its eigenvectors equal to the column of the Bradford transform matrix that is in question.

Since we do not know just what eigenvalue μ we shall have, for any particular sharpening interval chosen, a simple algorithm is to **1.** pick a sharpening interval ψ_k ; **2.** form the vector $v_k = [\Lambda(\omega)]^{-1} \Lambda(\phi_k) b_k$ on the left hand side of (11); **3.** normalise the length of vector v_k ; **4.** also normalise the length of the vector b_k on the right hand side of eq. (11). **5.** Calculate the Euclidean distance between the normalised vectors v_k and b_k . We repeat **1** through **5** to find the minimum Euclidean distance overall and so the best solution to (11).

Once we have found sharpening intervals ψ_k we can go back to our original colour-matching curves X and sharpen them in ψ_k according to eq. (9), without these normalisations. The resulting curves can then be compared with the actual Bradford curves Q_B .

It turns out, using this analysis, that the Bradford curves correspond to sharpening in the following intervals: R: 580nm spike; G: 550–560nm interval; B: 460nm spike. Fig. 2 shows the Bradford curves (dashed curves) and the approximations (solid curves) to these formed by actually sharpening the colour matching functions. The striking agreement between the Bradford curves and those for sharpened colour-matching functions implies that we may claim: **the Bradford curves equal the human colour-matching functions, sharpened in narrow intervals around 460, 555, and 580 nm.** Although spectral sharpening can deliver Bradford curves the results are not




Figure 2: Human colour-matching functions sharpened by unconstrained optimisation (solid lines), optimised for the same sharpening intervals as the Bradford curves, and compared to the original Bradford curves (dashed lines).

as one might expect. Sharpening at 460, 555 and 580nm returns respectively sensors which peak at 450nm, 540nm and 600nm. That is, the maximum sensitivity is not in the sharpening interval. In some sense this indicates a sharpening failure; we have been unable to concentrate sensitivity in the sharpening interval. Interestingly, 450nm, 540nm and 600nm are very close to the ‘prime wavelengths’, i.e. those parts of the spectrum to which we are maximally sensitive. Maybe intervals anchored in these regions would be more appropriate?

4. Sharpening with Positivity, and Positive Bradford Curves

We can ensure positivity of spectrally-sharpened sensors in two different ways, and each of these ways gives rise to a different approach to sharpening. Firstly, since we start with positive sensor curves, the simplest approach to developing a transform with positivity is to constrain the optimisation to a solution with positive, or non-negative, weights.

A second approach is to relax the above condition by allowing positive or negative weights, but directly constraining the optimisation so that the resulting *sensors* themselves are non-negative.

We refer to the first approach as an optimisation method with *constrained coefficients* and the second approach as an optimisation method with *constrained sensors*.

Using numerical optimisation schemes, here we investigate the effect of using an L_1 objective with an L_1 constraint, and an L_2 objective with an L_2 constraint.

4.1. Constrained Coefficients Sharpening

In this case our objective is to carry out a numerical optimisation with objective function

$$\begin{aligned} \min \sum \phi_k |Q(\lambda)c|^\nu \\ \text{with constraints} \\ \left\{ \begin{array}{l} \sum \omega |Q(\lambda)c|^\nu = 1, \quad L_\nu \text{ normalisation} \\ c \geq 0, \quad \text{non-neg. coeff's} \end{array} \right. \end{aligned} \quad (12)$$

where the exponent ν is 1 for an optimisation based on an L_1 norm, or is 2 for a least squares, L_2 norm based approach.

Firstly, we note two important theorems for these cases that prove that in fact the above minimisation need not be carried out throughout the c -space [5]:

Theorem 1 *Convexity implies that the solution for L_1 - L_1 sharpening, with coefficients constrained to be non-negative, lies on the boundary of the set of possible vectors c .*

The L_2 - L_2 case is the same as original spectral sharpening, but makes use of constrained optimisation. In this case, as in Theorem 1 for the L_1 - L_1 case, we have a convexity result that allows us to examine only the boundary of possible values of vector c .

Theorem 2 *Convexity implies that the solution for L_2 - L_2 sharpening lies on the boundary of the set of possible vectors c , if those coefficients are constrained to the non-negative range.*

4.2. Constrained Coefficients Bradford Curves

We may in fact select any sharpening intervals that are of interest, and here we wish to use sharpening intervals that encompass the prime wavelengths 450, 540, and 610nm. These particular wavelengths were shown to be the closest set of spikes to the human colour-matching functions, for

a collection of uniformly-distributed spectra [7] (they are the best positions for unconstrained sharpening), and are areas of maximal visual sensitivity[8].

The particular intervals we choose are [440–460] for blue, [530–550] for green, and [600–620] for red. Then *with these sharpening intervals unconstrained* spectral sharpening of the X colour-matching curves results in the sharpened sensors of Fig. 3 (solid curves).




Figure 3: Human colour-matching functions (dashed lines), compared to best unconstrained L_2 - L_2 optimisation (solid lines).

Using sharpening **with positivity**, Theorems 1 and 2 guarantee that weighting vectors c lie on the boundary of the search space, and that conclusion is indeed borne out here. However, we find that in fact the ‘sharpest’ positive sensors resulting from X , using the prime wavelength intervals, are just the original curves themselves, except for the sharpened version of the $\bar{z}(\lambda)$ curve, which is a combination of $\bar{z}(\lambda)$ and a small amount of $\bar{x}(\lambda)$. That is, the dashed lines in Figure 3 can be interpreted as the best sharpening of the XYZ curves. It is clear that, in this case, little sharpening has been achieved.

4.3. Constrained Sensor Sharpening

Instead of constraining just the weighting coefficients, we can instead constrain the entire *sensor function* result. I.e., we may allow coefficients c to take negative values but constrain the resulting sensor function itself to non-negative values.

In this case (12) is modified. We no longer use a lower bound constraint on c , but instead we constrain Qc :

$$\begin{aligned} \min \sum \phi_k |Q(\lambda)c|^\nu \\ \text{with constraints} \\ \left\{ \begin{array}{l} \sum \omega |Q(\lambda)c|^\nu = 1, \quad L_\nu \text{ normalisation} \\ Q(\lambda)c \geq 0 \quad \text{non-neg. sensor result} \end{array} \right. \end{aligned} \quad (13)$$

This is a linear or a quadratic programming problem with one equality constraint and s inequality constraints (e.g., s may be 31) [5].

4.4. Positive Bradford Curves from Constrained Sensors

The results of this type of sharpening are shown in Fig. 4, which results from sharpening the colour-matching functions X . However, since the Bradford curves are themselves simply a matrix transform away from the colour-matching functions, we arrive at precisely the same curves if we instead try to sharpen the Bradford curves themselves.

To evaluate such sharpened curves, let us define a ‘degree-of-sharpness’ goodness measure ε showing how




Figure 4: Bradford curves sharpened by constrained L_2 - L_2 sharpening with the sensor result constrained to non-negativity (solid lines). Original colour-matching functions are also shown (dashed lines).

much ‘energy’ is concentrated in the sharpening interval ψ . Each sensor $k = 1..3$ will employ a different ψ_k . If ε measures the amount of energy contained in ψ relative to in the entire visible spectrum ω , we may define

$$\varepsilon = 100 * \frac{\sum \psi_k |q_k(\lambda)|^2}{\sum_{\omega} |q_k(\lambda)|^2} \quad (14)$$

for each of the $k = 1..p$ sensors. The second line of Table 1 shows how unconstrained spectral sharpening with an L_2 objective and L_2 norm, derived according to the minimisation (5), behaves with respect to the goodness measure. We note that generally spectral sharpening greatly improves the energy concentration.

Table 1 also shows results for the original Bradford curves and for the Bradford curves sharpened by constraining the sensor result to non-negativity, according to the minimisation (13).

We see that the sharpened functions that result from constrained optimisation have good energy concentrations, better than or about the same as those for the Bradford curves, even though the new curves are all-non-negative, while every Bradford curve actually has some negative values. The unconstrained minimisation can, of course, produce better energy concentration because we allow negative lobes. For comparison, we also show the energy concentration for the MacAdam curves quoted by Pearson and Yule [6].

Another useful feature of sharpened sensors is that any ‘crosstalk’ between sensors is usually diminished. Let us define crosstalk κ between channels i and j of sensors Q by the angle

$$\kappa = \cos^{-1} \left\{ \frac{|q_i^t q_j|}{\|q_i\| \|q_j\|} \right\} \quad (15)$$

where q_i is the i th column of Q . The ideal value for κ is 90° . Table 2 shows that the value for the crosstalk between channels is generally improved, using sharpening.

In the next section we see that in fact we need not have positive sensors in order to have positive transformed-RGB values, with even better results.

5. Data-Driven Optimisation

Above, we successfully found positive sensors by constraining the curves themselves. However, in applications all we actually require is that transformed RGB values be

non-negative. Could we not then use a type of ‘data-based sharpening’ to determine sensors with these desired properties?

In [1], it was shown that the least squares transform from a set of RGB’s under one illuminant to the set under another illuminant yields approximately the same sharpening transform as recapitulated in §3.1 if that least squares matrix is diagonalised. In [1] this phenomenon was referred to as ‘data-based sharpening’.

Here we wish to investigate whether the requirement that the sensors be sharpened can be combined with the idea that under the ‘new’ sensors we wish to have only non-negative RGB values. We shall see that this scheme, which may yield sensors with negative lobes, leads to sharper sensors than those formed under the assumption of strict non-negativity of the curves themselves.

Suppose we consider an RGB triple ρ formed from a colour signal $C(\lambda)$ arriving at the camera sensors: if $E(\lambda)$ is the illuminant and $S(\lambda)$ is the surface spectral reflectance function corresponding to a particular pixel, then

$$C(\lambda) = E(\lambda)S(\lambda), \quad (16)$$

$$\rho = \sum_{\lambda \in \omega} C(\lambda)Q(\lambda)$$

Suppose we collect all such RGB triples ρ into an $n \times 3$ array \mathcal{R} , and also collect all the colour signals into an $n \times s$ array \mathcal{C} . Then we have

$$\mathcal{R} = \mathcal{C}Q \quad (17)$$

If the sensors themselves are changed to Q' via a matrix transform, then we obtain

$$Q' = QT \quad (18)$$

where column t_k pertains to sharpening the k th interval. The collection of RGB values changes to those seen under the new sensors:

$$\mathcal{R}' = \mathcal{C}QT = \mathcal{R}T \quad (19)$$

Now, we could set up a minimisation to achieve sharpening in a given interval, subject to non-negativity of *all* such sharpened RGB values \mathcal{R}' . However, this presents an unworkable set of constraints and in fact we can make use of convexity to work with convex hull points only.

Suppose we form the colour signal collection \mathcal{C} by using the reflectance spectra of the set of Munsell paint chips, various natural object reflectances, Dupont dye reflectances, and Macbeth ColorChecker reflectances. To form colour signals from these reflectances, let us use 11 illuminant spectra: A, D50, D65, D75, two fluorescent illuminants, and several measured SPDs.

Under each of the illuminants the set of calculated RGB’s forms a convex set [9].

Let us impose the reasonable constraint that the non-negativity of RGB points corresponding to the overall convex hull of the set \mathcal{R} be maintained under a transform (19).

Suppose the boundary set of RGB values is $\tilde{\mathcal{R}}$, with $\tilde{\mathcal{R}}$ an $\tilde{n} \times s$ matrix, where \tilde{n} is the number of samples in the boundary set. Our data-driven minimisation is thus

$$\min \sum_{\phi_k} [Q(\lambda)t_k]^\nu$$

with constraints

$$\begin{cases} \sum_{\omega} [Q(\lambda)t_k]^\nu = 1, & L_\nu \text{ normalisation} \\ \mathcal{R}t_k \geq 0, & \text{non-neg. RGB values} \end{cases} \quad (20)$$

Here we use an L_2 norm, with $\nu = 2$.

The sharpened sensors that result from this optimisation are displayed in Fig. 5.

The energy concentration is shown as the last line of Table 1 and is seen to be the best found using constrained




Figure 5: Sensors sharpened by a data-driven L_2 - L_2 sharpening with transformed convex hull points constrained to non-negativity (solid lines). Original curves are also shown, dashed.

optimisation. Interestingly, the sensors themselves are *not* all-positive. The data-driven sharpening results in sensors that are similar to the unconstrained ones of Fig. 3, but produce positive RGB's.

Thus we are left with the outcome that, for the sensors examined here, a data-driven sharpening that does not insist on non-negative sensors, but only on non-negative sensor response values, gives the best sharpening.

From the results of Tables 1 and 2, we can state the following: **The Bradford curves, that have been transformed via spectral sharpening with positivity, have superior energy concentration and crosstalk than the original curves.**

Energy Concentration (%)			
Algorithm	R	G	B
No sharpening	39.81	34.23	64.11
Unconstrained sharpening	50.53	49.95	64.62
Bradford	35.79	47.37	64.38
Pearson-Yule	24.77	44.87	64.11
Constrained-sensor sharpen'g	40.30	44.90	64.29
Data-driven sharpening	50.03	47.23	64.56

Table 1: Energy concentration in prime wavelength sharpening intervals, for original colour-matching functions X , sensors sharpened using unconstrained L_2 - L_2 optimisation, the Bradford curves, the Pearson-Yule curves, and sensors sharpened using constrained optimisation. The last line shows results for sharpening based on constraining RGBs, rather than the sensors themselves.

Crosstalk ($^{\circ}$)			
Algorithm	(R,G)	(R,B)	(G,B)
No sharpening	40.51	75.25	85.31
Unconstrained sharpening	82.42	87.53	86.14
Bradford	51.58	89.62	88.33
Pearson-Yule	34.25	88.51	83.50
Constrained-sensor sharpening	53.35	81.13	83.47
Data-driven sharpening	80.35	86.35	82.34

Table 2: Crosstalk κ , for original colour-matching functions, sensors sharpened using unconstrained optimisation, Bradford and Pearson-Yule curves, sensors sharpened using constrained optimisation, and sensors derived by constraining RGBs in the transformed space.

6. Conclusions

We applied techniques involving both L_2 and L_1 objectives and norms to the human colour-matching functions. Un-

constrained optimisation gives the best results for sharpening, but results in curves with negative lobes. An optimisation based on constraining only transformed RGBs, rather than the curves themselves, does best for a constrained sharpening, delivering almost as good results as for unconstrained sharpening but without the penalty of negative colours.

Note that the types of optimisation reported here can be applied to sensor systems of any dimensionality, although results are reported here only for the 3-band human system.

In a sense, the constrained-coefficients and constrained-sensors techniques presented here are a natural completion to the work of MacAdam, and Pearson and Yule [6]. The main advantage of using an optimisation, with positivity, that maximises energy concentration in desired sharpening intervals is that the process of making positive linear combinations of the colour-matching curves is guided not by simply decreasing crosstalk or making the most narrow curves, but by the practical necessity of sharpening within specific areas of the visible spectrum.

However, the data-driven approach, with an RGB constraint, is shown to be best overall, delivering almost as good sharpening as the best possible, but with none of the drawbacks. This being the case, we wonder whether these curves would lead to more accurate chromatic adaptation than the Bradford curves. We remind the reader that the latter were derived via a simple numerical fit for a single set of corresponding colour data. There is no good reason to suppose they are optimal in any sense.

Acknowledgement

Graham Finlayson gratefully acknowledges University of Derby (his former employer) for supporting this research.

References

- [1] G.D. Finlayson, M.S. Drew, and B.V. Funt. Spectral sharpening: sensor transformations for improved color constancy. *J. Opt. Soc. Am. A*, 11(5):1553–1563, May 1994.
- [2] K.M. Lam. *Metamerism and colour constancy*. PhD thesis, University of Bradford, 1995.
- [3] K. Barnard and B.V. Funt. Experiments in sensor sharpening for color constancy. In *6th Color Imaging Conference: Color, Science, Systems and Applications.*, pages 43–46. Society for Imaging Science & Technology (IS&T)/Society for Information Display (SID) joint conference, 1998.
- [4] B.V. Funt and G.D. Finlayson. Color constant color indexing. *IEEE PAMI*, 17:522–529, 1995.
- [5] M.S. Drew and G.D. Finlayson. Spectral sharpening with positivity, 1999. Submitted for Publication.
- [6] M.L. Pearson and J.A. Yule. Transformations of color mixture functions without negative portions. *J. Color and Appearance*, 2:30–35, 1973.
- [7] G.D. Finlayson and B.V. Funt. Coefficient channels: Derivation and relationship to other theoretical studies. *Color Research and Application*, 21:87–96, 1996.
- [8] M.H. Brill, G.D. Finlayson, P.M. Hubel, and W. Thornton. Prime wavelengths and color imaging. In *6th Color Imaging Conference: Color, Science, Systems and Applications*. Society for Imaging Science & Technology (IS&T)/Society for Information Display (SID) joint conference, 1998.
- [9] D.A. Forsyth. A novel approach to color constancy. In *ICCV88*, pages 9–18, 1988.