Recovering Spectral Sensitivities with Uncertainty

Ali Alsam and Graham Finlayson
School of Information Systems, University of East Anglia
Norwich, United Kingdom
graham@sys.uea.ac.uk

Abstract

It is well established that in order to obtain the best colour performance of a colour input device such as a scanner or a camera, that one needs to know the device spectral sensitivities. Unfortunately measuring sensitivities outside the laboratory is hard and moreover, manufacturers are reluctant to give the user specifications. Thus, there has been considerable interest in developing numerical techniques for estimating the spectral sensitivities.

These methods are based on taking images of known spectral targets and then, using knowledge of the image formation process, solving for the sensitivities using numerical methods. It is important to state that while these methods perform reasonably well, the problem is inherently ill-posed. There is simply not enough degrees of freedom in the spectral profile of a reflectance target to recover device sensitivities.

In this paper we tackle this uncertainty head on and develop a method to recover device sensitivities with uncertainty error bars. Experiments with a Megavision camera return a sensor estimate together with error bars. The error bars are sufficient to explain the discrepancy in the recoveries delivered by single-answer estimation algorithms and the actual sensitivities.

1. Introduction

It is accepted that in order to attain the best performance in a colour reproduction pipeline that input colour devices, such as scanners and cameras, should be spectrally calibrated. Unfortunately, determining the spectral response of scanners (or cameras) is not an easy task. In the laboratory camera sensitivities can be accurately measured using a monochromator. The monochromator is used to generate narrow spectral bands of light with known power. The device spectral response is estimated by noting the RGB response for different spectral bands across the visible spectrum. Unfortunately, a monochromator is a very expensive piece of test equipment (it is at least an order of magnitude more expensive than a camera) and so cannot be used outside the lab. An alternate approach would be to use a sequence of filters which transmit only in narrow wavelength bands.[1] Unfortunately, these filters are themselves expensive.

As a consequence researchers have sought other simpler methods for arriving at a spectral calibrations. The basic approach followed is to cast the sensor estimation problem as an equation solving exercise. Indeed, assuming linear device response, it is well known that the sensor response to spectral stimuli can be modelled as a linear equation of the form:

$$Ax = b \tag{1}$$

Where the matrix encapsulates the spectral stimuli, $\underline{b}$ the sensor responses (red, green or blue responses) and $\underline{x}$ the spectral sensitivity of the sensor. Standard numerical methods can be used to solve equations of this form to arrive at an estimate of the sensor spectral response. Unfortunately the degrees of freedom in the spectral sensitivity function are usually greater than those available in the spectral stimuli. As such A is 'ill-conditioned' and so solving for using naive methods returns a sensor which is highly tuned to the spectral stimuli in but which is far from the true spectral shape of the sensor.[2]

To get around this problem we might usefully incorporate domain knowledge in to the solution strategy thereby *regularizing* the solution making it more stable. For example, we know that spectral sensitivities are all positive functions[3] of wavelength. Moreover, they will tend to be somewhat smooth and also have a single dominant peak sensitivity.[2] Various methods have been proposed for solving (1) with these additional constraints[2,1,3,4] and good results are often attained.

Unfortunately, regularised solutions are only as good as the regularizing assumptions are appropriate. If for example, smoothness is assumed and the sensor has significant high frequency components then these must be missed in the sensor recovery. Moreover, the competing techniques for sensor estimation do not return consistent answers and, as a consequence, the user is left in an unsatisfactory position. They can, in good faith, implement the suite of calibration methods and 'solve' for the 'best' sensor according to each method. But, the user is given no guidance on which sensor is best overall?

A similar thought process has been followed in the colour constancy literature. Forsyth[5] formally demonstrated that the colour constancy problem is ill-posed: given a scene captured under unknown illumination there is a set of possible light colours that can account for the data. Forsyth presented an elegant algorithm for recovering the whole set of possible lights. It is only at the second stage that a single light is chosen from the set of possibilities. The best single illuminant can be chosen according to well founded statistical principles. For example, the mean illuminant colour[6] or the median[7] have been proposed as reasonable estimates. Additionally with prior information about the likelihood of different reflectances we might make a maximum likelihood selection.[8,9] Moreover, it is a simple matter to return to the user the uncertainty in the illuminant estimate.[10]

In this paper we apply the same sort of reasoning to the problem of sensor recovery. We begin by adopting only the weakest constraints: sensors should be positive, *somewhat* smooth and that the recovered sensors should predict RGB data within certain very liberal error bounds. Subject to these constraints we show how the *set of all plausible sensor estimates* can be created. We then set forth methods to calculate the mean and covariance of this set. In this framework our calibration returns a mean estimate together with error bars.

Experiments show that the recovered solution set contains, as we would expect, solutions delivered by regularization methods. Moreover, we have found that often the error bars are rather small indicating to the user that they can be confident of the appropriateness of their calibration.

In section 2 we give an overview of the estimation. Section 3 discusses how to recover the feasible set of solutions and thereafter how to calculate the mean estimate and the error bars. Experiments on a Megavision camera are reported in section 4. The paper finishes with a short conclusion in section 5.

2. Background

The response of a linear sensor to a spectral stimulus can be modelled as:

$$p = \int_{\omega} E(\lambda)S(\lambda)R(\lambda)d\lambda \;\; (k = R,G,B) \tag{2}$$

where λ is wavelength, p is the sensor response, E is the illumination and is the surface reflectance and R is a camera sensitivity function. Integration is performed over the visible spectrum ω.

Assuming that spectral functions can be represented by sampling at 10 Nanometer intervals across the visible spectrum (they can[11]), Equation (2) is rewritten as:

$$p_k = \sum_{i=1}^{31} E_i S_i R_i \Delta\lambda \tag{3}$$

where $\underline{E}$, $\underline{S}$, and $\underline{R}$ are 31 component vectors for light, reflectance and spectral sensitivity. The subscript indexes wavelength and the scalar $\Delta\lambda$ accounts for the 10 Nanometre sampling distance. Henceforth we will assume $\Delta\lambda$ is accounted for in $\underline{R}$.

Let us now consider the sensor response to set of spectral stimuli. Let *i*th row of an N x 31 matrix denote the *i* reflectance multiplied by the scene illumination (either the scanner light or the prevailing light for cameras). The *N* device repsonses by the vector *p* can be written as:

$$\underline{p} = CR \tag{4}$$

where the sum of multiplicands in (3) has been rewritten, in the usual way, as a matrix multiplication. We see that (4) is exactly in the same 'regression' form as Eq. (1).

Let us assume that the number of spectral stimuli *N* is much more than 31. Assuming that each column of *C* is linearly independent from each other column, the sensor $\underline{R}$ can be calculated by applying the pseudo inverse[12] to both sides of (4):

$$\left[C^t C\right]^{-1} C^t \underline{p} \approx \underline{R} \tag{5}$$

Notice that in solving for we have replaced the equality symbol in (4) with the approximation symbol in (5). We do this, because although the linear model formulation in (3) is reasonably accurate it is not perfect. Moreover, both the quantities and are measurements and are subject to measurement noise. Hubel et al[1] populated the rows of with narrow-band colour signals (constructed by placing a series of narrow-band interference filters in front of a camera). Thus, linear independence of was assured and the pseudo inverse solution returned adequate results.

Unfortunately, if the rows of are populated with colour signal spectra generated from a reflective target (e.g. the Macbeth colour checker or an IT8 chart) then the linear independence condition is not met and the pseudo inverse solution cannot be applied. Rather, it is well known that reflectances have at most 6 to 8 degrees of freedom[13,14] (since reflectances tend to vary smoothly with wavelength). As a consequence the inversion in (5) which assumes 31 degrees of freedom cannot be performed reliably. This is unfortunate since the reflectance target approach to calibration is by far the easiest and cheapest (interference filters are expensive) for a user to carry out.

The solution to this problem is to incorporate additional domain knowledge, or constraints, into the problem formulation. Constraints reported in the literature include:

(1) Sensor positivity:

$$R_i >= 0 \;\; (i = 1,2,\ldots,31)$$

(2) Sensor smoothness:

$$|R_i - R_{i-1}| \leq T \;(i = 2,\ldots,31)$$

(3) Unimodality

$$R_i > R_{i-1} \; (i = 1,2,\ldots,k) \text{ and } R_i < R_{i-1} \; (i = k+1, k+2,\ldots,31)$$

(4) Bounded prediction error:

$$\left|p_j - \underline{C}_j^t R_j\right| < \varepsilon \,(j = 1,2,\ldots,N)$$

where is a 1 x 31 row vector (for the th color stimuli).

In a general (and informal way) we might write the constrained, or regularised recovery problem as:

$$\overset{\min imize}{\underline{R}} \left|C\underline{R} - p\right| \tag{6}$$

subject to **(1)** AND/OR **(2)** AND/OR **(3)** AND/OR **(4)**

These constraints might be applied strongly in the sense that the solution has to meet them exactly. This is the approach followed in the methods.[2,3] Or they can be applied weakly. This is the 'penalty' approach used in Tikhonov regularization.[15,4] Moreover, there exists a range of optimization strategies that might be usefully applied including Linear programming,[16] Quadratic Programming[2] and Projection onto convex sets.[3]

Of course if different constraints are applied and different optimization strategies followed then the result is different sensor estimates. Since the recovery process is inherently ill posed the user cannot easily distinguish between different answers. Indeed, our own experiments (carried out where we know the spectral sensitivities) indicate that the 'best' methods depends on the device and data used in the analysis.

3. Solving for the Set of Sensors

Looking at (6) we notice that the sensor we are trying to recover must satisfy the sets of constraints. We might wonder therefore if the constraints themselves are sufficiently stringent that they themselves can lead to a sensor estimate. To see how we might proceed, remember that in the sensor estimation framework which we have developed, a sensor is a vector in 31-dimensional space. The positivity constraint simply restricts the sensor to lie in a particular part of the sensor space. Thus, just by applying positivity we have reduced the set of all possible sensor estimates. Similarly, the constraints on smoothness, unimodality and bounded prediction error, all demarcate regions of sensor space in which the recovered sensor must lie. The intersection of all these regions is exactly the **set of feasible solutions**. Any sensor that lies within the set is a perfectly reasonable answer to the sensor estimation problem.

Of course, we cannot usefully deal with a whole set of answers. Rather, to proceed we need to choose a single estimate from the feasible set. We propose that the centroid (or mean) of the set as a reasonable estimate. Moreover, we also calculate the standard deviation around the mean. In this way we arrive at an uncertainty estimate together with error bars. If the error bars are small then this indicates that we are quite certain that the mean estimate is correct. In contrast, large errors indicate higher uncertainty. In the limiting case if the errors are significant, this might necessitate the user to carrying out a second more stringent calibration (perhaps one involving a more expensive target).

3.1. Computational Methods

To compute the feasible set we first note that constraints (1) through (4) are defined by sets of linear inequalities. A single inequality in the 31-dimensional sensor space effectively splits the space into two parts: one half (which satisfies the inequality) in which the sensor might lie and the other which is infeasible. Given n inequalities, the sensor must simultaneously lie in each of the n half-spaces. That is it must lie in the intersection of halfspaces. This intersection region is a closed and convex subset of the 31-dimensional space. We might represent this convex set by a set of extreme vertices:

$$feasibleset = \{v_1, v_2, ..., v_M\} \tag{7}$$

All the interior points of the feasible set can be written as a convex combination of the vertices:

$$\underline{w} \in feasibleset \text{ iff } \underline{w} = \sum_{i=1}^{M} c_i \underline{v}_i$$
$$\left(\sum_{i=1}^{M} c_i = 1 \,\&\, c_i >= 0 \, (i = 1,2,...,M) \right) \tag{8}$$

The *feasible set* can be found using geometric algorithms. However, it should be pointed out that these algorithms have a large complexity. For the 31-dimensional case the worst case complexity is $O(n^{15})$ where n is the number of constraints. However, in practice the smoothness constraint diminishes the degrees of freedom in the sensor recovery. If for example smoothness is implemented by implementing the sensor band limit[2] to say 4 cycles across the visible spectrum then sensors actually belong to a 9-dimensional subspace of 31-dimensions and in this case the worst case complexity is $O(n^4)$ (which though large is tractable).

To calculate the centroid of a continuous convex set in higher dimensions is a computationally hard problem. Here we calculate the centroid through sampling. We randomly generate a sensor and then see if it lies in the feasible set. If it does we note it as feasible and then generate another random sensor. After many iterations we find the centroid by averaging all the feasible sensors we have generated. The feasible sensors can also be used for calculating the standard deviation around the mean for each wavelength across the visible spectrum. Thus, we can summarize the feasible set as:

$$\mu(\lambda) \pm \sigma(\lambda) \tag{9}$$

All computations were carried out in Matlab [www.mathworks.com] with external calls to Qhull [www.geom.umn.edu/software/qhull].

4. Experiments

To test our method we attempted to estimate the sensitivities of a Megavision camera. This camera was calibrated using a monochromator[17] and so we have a reasonable ground truth. Figure 1 shows the monochromator calibrated Megavision sensitivities. The RGB response for a Macbeth colour checker[18] taken under simulated D65 illumination and the corresponding colour signal spectra (measured with a PR650 spectrophotometer) are also known. Thus, we have all the information that is required to build the *feasible set*.

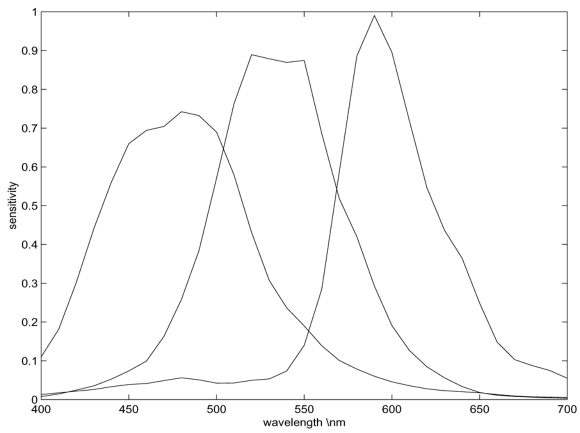

Figure 1. Megavision Spectral Sensitivities

The particular constraints used are however, important. First, we dropped the constraint of unimodality (it is a stringent constraint and one which we found to have little effect on the recovered estimate). Smoothness is implemented by forcing the recovered sensor to lie within the first 11 basis functions in the Sine basis expansion (11-dimensions were chosen as numerical tests indicted that this set was sufficient to account for a large class of sensors). We adopted a weak constraint on the bounded error. Estimates were forced to lie within an absolute error which constituted 5% of the signal range. Positivity of the estimate was enforced. Figure 2 is a visualization of the recovered *feasible set*. We show only the maximum and minimum sensitivity across the visible spectrum. The resulting envelope delimits a superset of the feasible set. All feasible sensors lie between the lower and upper bounds but not all sensors in this range are feasible.

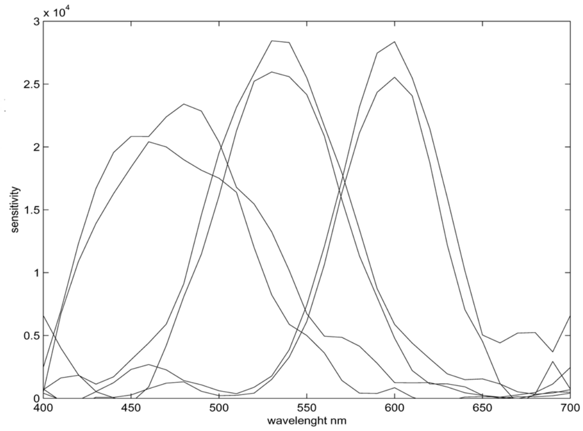

Figure 2. A visualisation of the recovered feasible set of sensors

We now calculated the mean of the feasible set together with the standard deviations. This is shown in Figure 3. Also in Figure 3 we have shown the sensor estimates recovered using the quadratic programming method of Finlayson and Hordley[2] and a constrained Tikhonov regularization based on a modified algorithm of that suggested by Dyas[15]. In this case the QP and Tikhonov recoveries are similar. But, notice they are quite far from the mean sensor. However, though far, they are in general within 1 standard deviation of the mean estimate. For the Megavision data we found that the mean sensor estimate delivered comparable prediction performance to that delivered by either the Tikhonov or Quadratic programming methods.

5. Conclusions

It is well established that in order to obtain the best colour performance of a colour input device such as a scanner or a camera, that one needs to know the device spectral sensitivities. Unfortunately measuring sensitivities outside the laboratory is hard and moreover, manufacturers are reluctant to give the user specifications thus, there has been considerable interest in developing numerical techniques for estimating the spectral sensitivities. These methods are based on taking images of known spectral targets and then, using knowledge of the image formation process, solving for the sensitivities using numerical methods. Though these methods often perform well, it is usually impossible for the user to understand how well they are working. Rather, the user is simply returned an estimate. This estimate may or may not be close to the true answer and so the calibration may or may not lead to the best colorimetric performance.

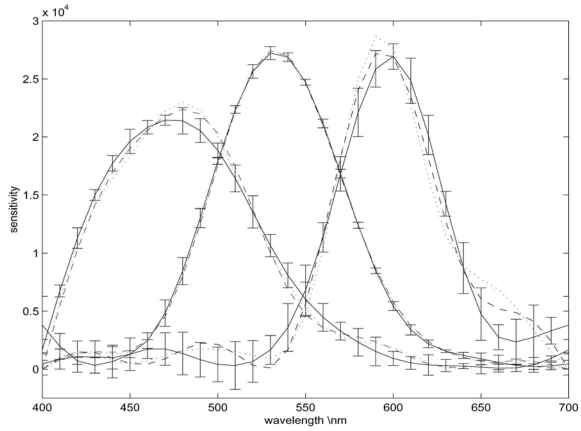

Figure 3. Solid line: mean sensor estimate. Dashed lines QP solution. Dotted line Tikhonov solution. QP and Tikhonov are almost within 1 standard deviation of the mean of the feasible set

In this paper we tackle this problem head on. We develop a method for recovering the *feasible set* of all reasonable sensor estimates. We then discuss how the mean and standard deviation for this set is calculated. Thus, our new calibration method returns a sensor estimate together with uncertainty. Experiments for a Megavision camera were carried out. For this camera we find the error bars are significant and graphically illustrates the inherent uncertainty in estimating sensor functions from reflective targets. The calibrations returned by the single-answer quadratic programming and Tikhonov methods are, as we might expect, within the feasible set.

Acknowledgements

Ali Alsam is grateful for studentship support from AGFA Gevaert. The authors gratefully acknowledge the help of Dr. Paul Hubel of Hewlett Packard Labs in carrying out the monochromator calibration of the Megavision camera.

References

1. P. M. Hubel, D. Sherman, and J. E. Farrel Drew. A comparison of methods for spectral sensitivity estimation. In *IS&T and SID's 2nd Color Imaging Conference*, pages 45–48. 1995.
2. G. D. Finlayson and S. D. Hordley. Recovering device sensitivities with quadratic programming. In *IS&T and SID's 6th Color Imaging Conference*. 1998.
3. G. Sharma and H. J. Trussel. Characterization of scanner sensitivity. In *IS&T and SID's Color Imaging Conference:*

Transforms&Transportability of Color, pages 103–107. 1993.

4. K. Barnard and B. V. Funt. Camera calibration for colour research. In *Human vision and electronic imaging IV*, volume 3644, pages 576–585. SPIE, 1999.
5. D. Forsyth. A novel algorithm for color constancy. *Int. J. Comput. Vision*, 5:5–36, 1990.
6. K. Barnard. Computational color constancy: taking theory into practice, 1995. MSc thesis, Simon Fraser University, School of Computing Science.
7. G. Finlayson and S. Hordley. Selection for gamut mapping colour constancy. *Image and Vision Computing*, 17(8):597–604, June 1999.
8. G. D. Finlayson, S. D. Hordley, and P. M. Hubel. Colour by correlation: a simple unifying theory of colour constancy. In *IEEE international conference on computer vision*, pages 835–842, 1999.
9. M. D'Zmura and G. Iverson. Probabalistic color constancy. In R.D. Luce, M. M. D'Zmura, D. Hoffman, G. Iverson, and K. Romney, editors, *Geometric Representations of Perceptual Phenomena: Papers in Honor of Tarow Indow's 70th Birthday*. Laurence Erlbaum Associates, 1994.
10. G. D. Finlayson and S. D. Hordley. Colour constancy with error bars. In *IEE conference on Image Processing and Applications*, 1999.
11. B. Smith, C. Spiekermann, and R. Sember. Numerical methods for colorimetric calculations: Sampling density requirements. *COLOR research and application*, 17(6):394–401, 1992.
12. G. H. Golub and C. F. van Loan. *Matrix Computations*. John Hopkins U. Press, 1983.
13. L. T. Maloney. Evaluation of linear models of surface spectral reflectance with small numbers of parameters. *J. Opt. Soc. Am. A*, 3:1673–1683, 1986.
14. J. P. S. Parkkinen, J. Hallikanen, and T. Jaaskelainen. Characteristic spectra of munsell colors. *Journal of the Optical Society of America A*, 6:318–322, 1989.
15. B. Dyas. Robust color sensor response estimation. In *IS&T and SID's 8th Color Imaging Conference*, pages 144–148. 2000.
16. F. Koenig and P. Herzog. Spectral characterization using linear programming. In *Color Imaging: Device Independent color, color hardcopy, and graphic arts V*, volume 3963, pages 36–56. SPIE, 2000.
17. The authors than Paul Hubel of Hewlett Packard corporation for making this data available.
18. C. McCamy. A nomograph for selecting light-balancing filters for camera exposure of color films. *Photog. Sc. Eng.*, 3:302–304, 1959.

6. Biography

Ali Alsam received a diploma in Mechanical Engineering from GTI Gothenburg's Technical Institute Sweden 1997 and a B.S in Mechanical engineering 2000 from the University of Surrey England. Since Sep 2000 he has been working towards a PhD in colour science at the University of East Anglia in England.

Professor Graham Finlayson leads the colour group at the University of East Anglia. In the last 5 years he has published almost 100 papers in the areas of colour in computer vision, colour science, content-based image retrieval and computational methods for colour calculation. In 2000 he was IS&T general chair for the 9th Color Imaging conference in Scottsdale Arizona. He is a member of the programme committee for CGIV.