

This paper was published in JOSA A and is made available as an electronic reprint with the permission of OSA. The paper can be found at the following URL on the OSA website:
<http://dx.doi.org/10.1364/JOSAA.33.000589>

© 2016 Optical Society of America. One print or electronic copy may be made for personal use only. Systematic reproduction and distribution, duplication of any material in this paper for a fee or for commercial purposes, or modifications of the content of this paper are prohibited.

Rank-based camera spectral sensitivity estimation

Graham Finlayson, Maryam Mohammadzadeh Darrodi and Michal Mackiewicz

Vol. 33, No. 4 / March 2016 / J. Opt. Soc. Am. A

Rank-Based Camera Spectral Sensitivity Estimation

GRAHAM FINLAYSON¹, MARYAM MOHAMMADZADEH DARRODI^{1*}, AND MICHAL MACKIEWICZ¹

¹School of Computer science, University of East Anglia, Norwich, United Kingdom

*Corresponding author: m.darrodi@uea.ac.uk

Compiled January 29, 2016

In order to accurately predict a digital camera response to spectral stimuli, the spectral sensitivity functions of its sensor need to be known. These functions can be determined by direct measurement in the lab - a difficult and lengthy procedure - or through simple statistical inference. Statistical inference methods are based on the observation that when a camera responds linearly to spectral stimuli, the device spectral sensitivities are linearly related to the camera rgb response values, and so can be found through regression. However, for rendered images, such as the jpeg images taken by a mobile phone, this assumption of linearity is violated. Even small departures from linearity can negatively impact the accuracy of the recovered spectral sensitivities, when a regression method is used.

In our work, we develop a novel camera spectral sensitivity estimation technique that can recover the linear device spectral sensitivities from linear images and the *effective* linear sensitivities from rendered images. According to our method, the rank order of a pair of responses imposes a constraint on the shape of the underlying spectral sensitivity curve (of the sensor). Technically, each rank-pair splits the space where the underlying sensor might lie in two parts (a feasible and infeasible region). By intersecting the feasible regions from all the ranked-pairs, we can find a feasible region of sensor space. Experiments demonstrate that using rank orders delivers equal estimation to the prior-art. However, the Rank-based method delivers a step-change in estimation performance when the data is not linear and, for the first time, allows for the estimation of the effective sensitivities of devices that may not even have 'raw mode'. Experiments validate our method. © 2016 Optical Society of America

OCIS codes: (150.1488), Calibration; (330.1690), Color; (330.1710), Color, measurement; (330.1715), Color, rendering and metamerism; (330.1730), Colorimetry; (040.1490), Cameras; (130.6010), Sensors; (000.4430), Numerical approximation and analysis; (110.2960), Image analysis; (150.1135), Algorithms; (110.5200), Photography

1. INTRODUCTION

Knowledge of a camera spectral sensitivity functions is important for many applications including multi-spectral imaging [1, 2], spectral reflectance and illuminant estimation [3–5], colour correction [6, 7] and colour constancy [8, 9]. Since these functions are not available from the camera manufacturers, they must be measured or estimated.

In the lab, a camera sensor spectral sensitivities are measured by recording the device's rgb response to light for narrow-band stimuli at a range of wavelengths across the visible spectrum. The narrow-band stimuli might be produced by a monochromator [10], integrating sphere [11] or sets of narrow band interference filters [12, 13]. However, this lab-based method is both time-consuming and expensive to carry out, and moreover is not available to most users.

There has been a sustained interest in the digital imaging community in developing estimation methods for recovering device sensitivities. The idea is to infer the spectral shape of the r, g and b spectral sensitivity curves given known spectral measurements of the light entering the camera and the corresponding recorded camera rgb values. Assuming the image is linear - i.e. when light is doubled the response is doubled - the estimation problem can be posed as a linear regression. Thus, we can solve for the spectral function of the sensor mechanisms over the visible spectrum [14–16].

Of course, nowadays most images - for example, from mobile phones - are not linear but rather they are *rendered*. The difference between raw and rendered is shown in Figure 1. To a tolerable first approximation, raw images are mapped to rendered counterparts in two stages. First the raw rgb is mapped to




Fig. 1. Image of the Macbeth SG colour chart. Left: Raw (colour corrected), Right: Rendered camera output.

a corresponding display RGB via a 3×3 matrix transform. This, so called, colour correction transform accounts for the display and colour preference [17]. The colour corrected RGB values are then mapped by an increasing function $f()$ which accounts for gamma, tonality and preference [18–21]. In this paper, for rendered images, we are also interested in recovering the *effective* device spectral sensitivities which we define to be the device sensors multiplied by the (unknown) colour correction transform.

Our method is based on rank orders of the R, G and B pixel values in an image. Our observation is that if the response for one patch in a colour target is larger than that measured for a second patch, then this places a constraint on the shape of the spectral sensitivity curve (of the sensor).

Let us illustrate this concept in a simple example. Suppose that we have two surfaces, one of which is a red patch, which is highly reflective in the red part of the visible spectrum, and the other a purple patch which is slightly less reflective in the long wavelengths but has significant reflectivity in the shorter wavelengths. Now, we measure the red reflectance (illuminated by white light) with a red sensitive sensor. The red patch response is larger than that recorded for the purple patch. This makes sense, as the red sensor is most sensitive to longer wavelengths. In contrast, if the sensor were sensitive only to short wavelengths, then we would record an opposite rank order i.e. the response to the purple reflectance would be greater than for the red one. Thus, by looking at a single rank ordering of two reflectances we can immediately distinguish - albeit in coarse terms - whether a sensor is more likely to be long-wave or short-wave sensitive. Of course, we actually know whether we are trying to calibrate a red or blue sensor but this rank ordering idea turns out to be very powerful in determining the shape of the underlying spectral sensitivity curve (of the sensor). We show that each pairwise rank ordering splits the space of all spectral sensitivities into two parts: a feasible and an infeasible region. The intersection of feasible regions (computed for pairs of responses) delimits a feasible region of sensor space from which we pick our overall sensor estimate.

By construction, our *Rank-based* method is invariant to any increasing function that is applied to the raw rgb values of an image (since the application of an increasing function will not change the pair-wise rank orders). It follows then that the Rank-based method applied to rendered images can recover the effective device linear spectral sensitivities (the actual raw spectral

sensitivities multiplied by the colour correction matrix).

We benchmark the spectral sensitivities recovered by different algorithms against the spectral sensitivity functions of a Nikon D5100 and a Sigma SD1Merill cameras that we measured at the National Physical Laboratory [10]. When the camera response is linear, we demonstrate that the proposed Rank-based method provides estimated spectral sensitivity functions that are equally accurate as those returned by cited antecedent spectral sensitivity estimation methods. However, when camera response is non-linear, Rank-based spectral estimation is shown to provide a step change in our ability to estimate the effective camera spectral sensitivities.

In section 2, we review colour image formation along with prior-art camera spectral sensitivity estimation algorithms. Rank-based spectral sensitivity estimation is described in section 3. We evaluate our method on linear and rendered images in section 4. The paper concludes in section 5.

2. BACKGROUND

The camera response at j th pixel from i th sensor can be modelled as:

$$p_{ij} = \int_{\omega} E(\lambda) S^j(\lambda) Q^i(\lambda) d\lambda, \quad i = \{R, G, B\} \quad (1)$$

where p_{ij} denotes the camera response, $E(\lambda)$ is the spectral power distribution of the scene illuminant, $S^j(\lambda)$ is the surface reflectance imaged at pixel j and $Q^i(\lambda)$ is i th spectral sensitivity function. In sensor estimation it is more useful to operate on spectral functions as discrete vectors. For example, let us sample spectral quantities at every 10 Nanometre intervals [22], starting from 400 and ending at 700 Nanometres, and replacing integral by a summation. That is, each spectral function is represented as a 31-vector. In general, we can do this for any sampling interval $\Delta\lambda$ resulting in an N-vector spectral function, where $N = \frac{700 - 400}{\Delta\lambda} + 1$. Thus, we can rewrite the integral in Eq. 1 as:

$$p_{ij} = \sum_{l=1}^N E(\lambda_l) S^j(\lambda_l) Q^i(\lambda_l) \Delta\lambda, \quad i = \{R, G, B\} \quad (2)$$

where $\lambda_l = 400 + (l - 1) * \Delta\lambda$ Nanometres. Here and throughout this paper underlining denotes vector quantities. We will also, henceforth, assume that $\Delta\lambda$ is incorporated into the vector of spectral sensitivity measurements. Denoting light, reflectance and sensor by the vectors \underline{E} , \underline{S}^j and \underline{Q}^i we rewrite (2) as:

$$p_{ij} = \sum_{l=1}^N E_l S_l^j Q_l^i, \quad i = \{R, G, B\} \quad (3)$$

We combine the light multiplied by reflectance into a single spectral quantity, the colour signal, denoted by $C(\lambda) = E(\lambda)S(\lambda)$. In the discrete representation where spectral quantities are represented by vectors, the colour signal for the j th reflectance, \underline{c}^j , is computed as:

$$\underline{c}_j = \text{diag}(\underline{E}) \underline{S}^j \quad (4)$$

where the operator *diag* transforms any $m \times 1$ vector into an $m \times m$ diagonal matrix (the *diag* operator draws attention to the fact that component-wise vector multiplication is not in simple linear algebra). We are now in a position to rewrite the integral

in Equation 1 using the language of linear algebra. For a single pixel, the sensor response is the dot-product of the colour signal spectrum and the sensor; or the transpose of the colour signal spectrum multiplying the sensor.

$$p_{ij} = \underline{c}^j \cdot \underline{q}^i = [\underline{c}^j]^t \underline{q}^i \quad (5)$$

Let C denote the $n \times N$ matrix of colour signal spectra (one per row). The corresponding n responses are written as:

$$\underline{P} = C\underline{q} \quad (6)$$

In spectral sensitivity estimation, it is assumed that the colour signal matrix C and the camera responses are known. This restriction is not overly onerous since we might probably know the spectral reflectances in a colour target and we can calibrate our camera under a light source. We can solve for the sensor spectral sensitivity by minimising:

$$\min_{\underline{q}} \| C\underline{q} - \underline{P} \|^2 \quad (7)$$

The least-square closed-form solution to Equation 7 is written as:

$$\underline{q} = C^+ \underline{P}. \quad (8)$$

where the matrix $C^+ = (C^t C)^{-1} C^t$ is the Moore-Penrose pseudo inverse [23]. Equation 7 is a simple, elegant but ultimately, naive formulation of spectral sensitivity estimation. In the top of Figure 2 we show our ground truth measured Nikon camera sensitivities. The results, by simple linear estimation using the raw image in the left side of Figure 1 together with corresponding rendered colour signal spectra, are the ‘jaggy’ sensitivities shown at the bottom of Figure 2.

At the heart of the Moore Penrose pseudo inverse is the raw cross product $C^t C$. This matrix is only *stably* invertible if it has full rank (the jaggy recovery shown in Figure 2, in essence informs us that the matrix doesn’t meet this condition). The raw cross product matrix is full rank if and only if, the set of colour signal spectra spans the set of all possible spectra. In practice, this is unlikely to be the case, since surface reflectance are inherently smooth [24]. Parkkinen [25], amongst others [13, 26], estimates that the dimensionality of spectral reflectance is between 6 and 9. When the raw-cross product matrix has low rank this means that very small perturbations to the measurements (in this case \underline{P}) result in very different regression results. Both the Ground-truth and jaggy sensors in Figure 2 describe the data very well. There exists a small perturbation to \underline{P} , that would result in the correct ground-truth sensors being recovered.

A specific and interesting condition which, immediately, forces $C^t C$ to be rank deficient, is when the number of measurements are less than the dimensionality of the matrix. When we have fewer known variable than unknowns, there are (infinitely) many sensors that *exactly* account for the data. A classic approach is to find the ‘least-norm’ sensor from amongst these sensors. Suppose we assume that the sensor \underline{q} is in the row-space of C i.e. $\underline{q} = C^t \underline{c}$. Clearly, one solution for \underline{c} would be:

$$\underline{c} = (C C^t)^{-1} \underline{P} \quad (9)$$




Fig. 2. Top: Spectral response function of Nikon camera measured at NPL [10]. Bottom: Jaggy (high norm) spectral response function estimated using Eq. 7 on 140 reflectance patches of Xrite SG chart.

In other words, the sensor is defined as $\underline{q} = C^t (C C^t)^{-1} \underline{P}$. It is straightforward to show that this solution, compared against all others, has least-norm [27]. This idea of solving for a sensor while implementing a secondary constraint, in the case of least norm, is often useful in spectral sensitivity estimation. Referring back to Figure 2, the norm (magnitude) of the actual sensor is much less than the simple least-squares estimation since the more jaggy the sensor the higher its norm. This idea is exploited by the prior-art estimation algorithms.

A. Prior-art estimation algorithms

I Regularisation: Equation 7 is *regularised* by adding a penalty term on the norm:

$$\min_{\underline{q}} \| C\underline{q} - \underline{P} \|^2 + \gamma \| T\underline{q} \|^2 \quad (10)$$

where $\gamma > 0$ is a user defined Regularisation parameter and T is a linear transform. Where T is equal to the $N \times N$ identity matrix, Eq. 10 attempts to minimise both the least-squares fitting error and the norm of \underline{q} . This procedure is known as Tikhonov Regularisation [28]. Alternatively, a 2nd order derivative matrix operator [29] can be used to define T as:

$$T = \begin{bmatrix} 1 & -1 & 0 & 0 & \cdots & \cdots & \cdots & \cdots & 0 \\ -1 & 2 & -1 & 0 & \cdots & \cdots & \cdots & \cdots & 0 \\ 0 & -1 & 2 & -1 & \cdots & \cdots & \cdots & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \ddots & \ddots & \ddots & \cdots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & 0 & -1 & 2 & -1 \\ 0 & 0 & 0 & 0 & \cdots & 0 & 0 & -1 & 1 \end{bmatrix}. \quad (11)$$

The regularised least-squares solution to spectral sensitivity recovery is written in closed form as:

$$\underline{q} = (C^t C + \gamma T^t T)^{-1} C^t \underline{p}. \quad (12)$$

A natural question arising upon the application of either Regularisation methods, is how to set γ . A good solution is one which makes both the least-squares error and the norm, simultaneously small. The best compromise solution is usually determined by trying a variety of different γ values and then applying a statistical criterion to choose the best gamma overall [30].

II Subspace Methods: Suppose, the spectral sensitivity of sensors is described by a low-dimensional linear model:

$$Q(\lambda) = \sum_{k=1}^m \alpha_k B_k(\lambda) \quad (13)$$

where $B_k(\lambda)$ denotes the k th basis function and α_k is a scalar coefficient. Equivalently, $q = B\alpha$ where B is a $N \times m$ matrix where the k th column is the discrete sampled version of a basis function. Here α is an m -vector of basis weights. Substituting $q = B\alpha$ into Equation 7, sensor estimation involves finding the best basis weights α for:

$$\min_{\alpha} \|CB\alpha - \underline{p}\|^2, \quad (14)$$

The least-squares solution of Equation 14 is

$$\alpha = [CB]^+ \underline{p} \quad (15)$$

In [11, 31] linear models were derived by the principal component analysis (PCA) of measured spectral sensitivities of 28 different Nikon and Canon camera models. Remarkably, it was found that even using just the first 2 principal components of this data-set, it was sufficient to account for their data. Zhao [32] proposed using a simple *Sine* basis expansion, where between 7 and 10 basis functions were used.

III Convex Optimisation: Equation 10 is constrained by incorporating *soft* penalties into the problem formulation. However, there are also *strong* constraints which we might usefully apply. For example, we know that the spectral sensitivities of the devices are zero at both ends of the visible spectrum and also they tend to have one, or at most two peaks. These "harder" constraints are often described using linear equalities or inequalities. Optimising a linear objective such as Equation 6 subject to linear equality/inequality constraints is a convex optimisation.

In [33] sensor estimation is posed as minimising a linear objective function subject to linear constraints and the sensors

are found using linear programming. Finlayson and Hordley [14] proposed that assuming a known peak sensor response at the z th sample point, the spectral sensitivity is estimated as:

$$\min_{\underline{q}} \|C\underline{q} - \underline{p}\|^2$$

Subject to the constraints:

$$\begin{aligned} \underline{q} &= B\alpha \\ q_l &\geq 0 \\ q_l &\leq q_{l+1}, l = 1, \dots, z-1 \\ q_l &> q_{l+1}, l = z, \dots, m \end{aligned} \quad (16)$$

The standard Fourier basis is used for B which up to a given band limited frequency composed of *Sine* and *Cosine* basis functions. The formulation in Equation A is a quadratic programming problem, (quadratic objective with linear constraints). The global best solution for QP and linear programming problems can be solved by efficient simplex-type search algorithms.

B. Measured spectral sensitivity functions

In this article, we use the measured spectral sensitivities of Nikon D5100 and Sigma SD1Merrill cameras as the benchmark to evaluate spectral estimation algorithms. The camera were rigorously measured at the the National Physical Laboratory [10]. The resulting measured spectral response functions for both Nikon and Sigma, are illustrated in Figure 3 and are also accessible online [34].




Fig. 3. Measured Nikon (Top) and Sigma (Bottom) camera sensitivities from NPL calibration facility [10, 34].

3. THE RANK-BASED METHOD

In this section, we show how we can use a single image of a colour chart to estimate the camera spectral sensor sensitivities irrespective of whether the image is linear or rendered (non linear). That we can solve the latter case, is an improvement over all the prior-art.

A. Rank-based spectral estimation

The rank order of two single-channel linear responses p_a and p_b is encapsulated by writing an inequality:

$$p_a > p_b. \quad (17)$$

We note that, if we apply a monotonically increasing function $f()$ to both sides of the above inequality, the rank order will still be preserved:

$$f(p_a) > f(p_b), \quad (18)$$

Let us revisit our vector-based model of image formation, Eq. 5, and write the pair-wise rank order of two responses:

$$\underline{c}_a \cdot \underline{q} > \underline{c}_b \cdot \underline{q} \Rightarrow (\underline{c}_a - \underline{c}_b) \cdot \underline{q} > 0, \quad (19)$$

where \underline{c}_a , \underline{c}_b and \underline{q} denote respectively the colour signals for reflectances a and b and the spectral sensitivity of the sensor. Given n colour samples we can choose $\binom{n}{2}$ pairs of responses (and the corresponding pairs of colour signal spectra) to form a set of rank orders, each with the form shown in Equation 19. Notice, that each pair-wise constraint is the dot-product of the difference in a pair of colour signal spectra against the spectral sensitivity of a sensor we would like to estimate. Let us now denote the j th colour signal difference vector as \underline{d}^j ,

$$\underline{d}^j \cdot \underline{q} > 0. \quad (20)$$

Equation 20 has a useful geometric interpretation. It teaches us that \underline{q} must belong in one half of the space of all sensors. In other words, if \underline{q} satisfies the inequality, then $-\underline{q}$ does not. Geometrically, Eq. 20 defines a half-space constraint. We denote this half-space as $H(\underline{d}^j)$. The intersection of all the half-spaces delimits an unbounded *feasible* convex region of sensor space. According to rank orderings, the actual true spectral sensitivity must lie in this feasible set. Mathematically, we write:

$$\underline{q} \in \bigcap_j H(\underline{d}^j). \quad (21)$$

Let us visualise Eq. 20, by assuming the (grossly over) simplified problem, where we seek only to estimate spectral sensitivities at two wavelengths. With respect to this simplification, each half-plane constraint of Eqn. 20, divides the Cartesian plane into two halves. In Figure 4 we show 4 half-plane constraints where solid arrows show which side of the line is the feasible half-space. The overall bounded feasible convex cone is the hatched area shown. In this paper, we are assuming that our spectral sensitivities are 31-vectors (the value of the sensor at 10Nm sampling intervals for the visible spectrum of 400 to 700Nm). Thus, the feasible set is a 31-dimensional unbounded convex cone. The unbounded cone, encapsulates the idea that with rank ordering constraints alone, we can constrain the *shape* (or, vector orientation) but not its *scale* (not the vector magnitude).




Fig. 4. Illustration of the intersection of 4 half planes with blue diagonal stripes.

B. The Rank-based method for rendered images

To a tolerable approximation [18] raw rgb values are mapped to rendered counterparts in two steps. First, the raw rgb values are mapped by a 3x3 colour correction matrix M (which accounts for the display primaries [35] and preference [17]), resulting in RGB values. The effective sensitivities for a rendered image, \underline{Q} are defined as:

$$\underline{Q} = \underline{Q}M \quad (22)$$

In a second step, an overall, or per channel, transfer function $f()$ is applied to the colour corrected RGB values. Every transfer function is increasing and implements display gamma and tone adjustment [17]. The mathematical model for rendered image is written as:

$$p_i = f_i(\underline{c}^t \underline{Q}^i) \quad (23)$$

where \underline{Q}^i denotes the i th column of \underline{Q} . Returning to our rank ordering idea, for rendered images:

$$p_a^i > p_b^i \Rightarrow f_i(\underline{c}_a^t \underline{Q}^i) > f_i(\underline{c}_b^t \underline{Q}^i) \Rightarrow (\underline{c}_a^t - \underline{c}_b^t) \underline{Q}^i > 0, \quad (24)$$

Note, we drop $f()$ because this function is increasing (does not change the ranks of the responses). For rendered images we can only recover the effective sensitivity \underline{Q}^i .

C. Implementation

In spectral sensitivity estimation, in general, we attempt to recover the spectral sensitivities of a camera by taking a single picture of a colour target (e.g. see Figure 1) where we assume that the reflectances of the colour target are known and that the prevailing illuminant spectral power distribution is measured. In Rank-based estimation, each pair of responses (per channel) can be related to a colour difference signal and a linear inequality of the form in Equation 19. Taken together the set of inequalities delimits an overall feasible region of sensor space where the spectral sensitivity must be. To choose a single member from this set, we seek a sensor that has zero response at the both ends of the visible spectrum and integrates to 1. Also inspired by the work of Zhao [32], we seek for a sensor that spans a

7-dimensional Sine basis. By these, we assure that the chosen sensor has the commonly addressed usual characteristics of spectral sensor sensitivity functions. The Rank-based spectral sensitivity is recovered by minimising the following quadratic program

$$\min_{\underline{q}} \| C\underline{q} - \underline{p} \|^2$$

Subject to the constraints:

$$\begin{aligned} \underline{d}^j &> 0 \quad j = 1, 2, \dots, N \\ q_1 &= 0 \\ q_N &= 0 \\ \underline{q} &= B\underline{\alpha} \\ \sum_i q_i &= 1 \end{aligned} \quad (25)$$

where B is an $N \times 7$ matrix of *Sine* basis functions. Note that for the case of rendered image, where we are effectively solving for $\underline{Q} = QM$, we remove the positivity constraint from the formulation.

D. A cautionary remark

Potentially, the rendering Equation 23, is actually yet more complex [18]. In particular, it is often the case that cameras implement a gamut mapping step which we haven't included in our model. The need for gamut mapping is due to the fact that the colour correction matrix M will map raw image colours outside the gamut, which cannot be displayed. Assuming display RGB is the unit cube (all values between 0 and 1), some of the mapped colours might be negative or greater than 1. Gamut mapping [36] moves colours, in a proprietary non-linear way, towards the neutral axis so that all the colours in an image are displayable i.e. all colours lie in the unit cube. Denoting the gamut-mapping function as $\Gamma(\cdot)$, rendered image formation is written as

$$p_i = f_i(\Gamma(\underline{c}^t \underline{Q}^i)) \quad (26)$$

However, in rendering, the colour correction matrix and per channel functions $f_i(\cdot)$ provides most of the heavy lifting and in fact, if the gamut mapping step is left out, the resulting image is a tolerable approximation of the gamut mapped version. This is especially true for image colours which are not saturated (as the effect of any gamut mapping step can be ignored for colours that lie near the neutral axis).

Empirically, we found that the gamut mapping step did not change the rank-order of the responses. But, equally, we found that restricting ourselves to less saturated colours still provided a strong Rank-based constraint and our method continued to work well.

4. EXPERIMENTS AND RESULTS

A. Data Acquisition

We took an image of a 140 colour chip Xrite SG colour checker chart using the Nikon D5100 and Sigma SD1Merril cameras. This will give us 9730 colour pairs. The chart is imaged in VeriVide cabinet, facing upwards towards the D65 light source at

an angle of approximately 45 degrees. The $P_{140 \times 3}$ matrix of rgb values is the average camera response for the central 10x10 pixel region of each checker patch. We made sure that the images are non-clipped, by adjusting the exposure level of the camera. The dark current noise level is measured (by taking a picture with the lens cap on) and is subtracted from the raw image.

A PR670 spectroradiometer was, in analogy to the camera capture conditions, positioned approximately at 45 degrees relative to the colour checker facing the VeriVide cabinet. The spectroradiometer - which measures the spectrum of light at a single spot in the image - was held fixed. That is, the colour checker was moved so that each colour patch was measured at the same location (with respect to the same 45 degree viewing geometry). The per-chip spectral measurements comprise the matrix $C_{140 \times N}$ (there are 140 chips in the colour checker and we measure the visible spectrum, 400 to 700 Nm, at a 10Nm sample interval).

We benchmark spectral sensitivity estimation against measurements made at the National Physical laboratory presented in 2.B. The effective sensitivities - see section 3.B. for a discussion - of the Nikon and Sigma camera are found via radiometric calibration method [18].

We test three cited antecedent methods: Tikhonov based on derivatives [29], the Subspace method [11] (using just 2 basis functions per sensor class) and Quadratic programming [14]. In the data tables that follow these methods are respectively denoted 'Regularisation', 'Subspace' and 'QP'. Finally, we compute recovery statistics for our new method 'Rank-based'.

B. Error Metrics

A. Percentage Spectral Recovery Error. Let \underline{q} and $\hat{\underline{q}}$ denote the actual and estimated sensor. The percentage spectral recovery error is calculated as:

$$SE = 100 \times \frac{\|\underline{q} - \hat{\underline{q}}\|}{\|\underline{q}\|} \quad (27)$$

Although intuitive, the percentage spectral error tends to be sensitive to small spectral shifts in the position of the estimated sensor. The spectral error between a sensor and a duplicate that is shifted 10Nm towards the longer wavelengths can be very large despite the fact that both sensors have an integrated response that is similar. This discrepancy has led researchers to measure the goodness of a recovery in terms of how well they predict actual sensor values.

B. The Vora Value Metric. This metric was designed to compare two sensors in terms of how similar they integrate all theoretically possible spectra [12, 21]. Suppose that \underline{p}_i and $\hat{\underline{p}}_i$ denote the actual and predicted rgb camera response. The percentage Response Error is calculated:

$$RE = \frac{\|\underline{p}_i - \hat{\underline{p}}_i\|}{\|\underline{p}_i\|} \quad (28)$$

A key insight of Vora is that we might be able to reduce the RE errors by colour correcting the measured rgb values:

$$RE = \frac{\|\underline{p}_i - M\hat{\underline{p}}_i\|}{\|\underline{p}_i\|} \quad (29)$$

Where M denotes a 3×3 matrix. The Vora value calculates the average RE (for the optimal M) over all possible spectra. Below, we recapitulate the formula for the Vora value.

$$Vora = 100(1 - (\frac{trace(QQ^+ \hat{Q}\hat{Q}^+)}{3})) \quad (30)$$

Q and \hat{Q} respectively denote the actual and estimated sensitivities, $trace()$ returns the sum of the diagonal of a matrix and \hat{Q}^+ indicates the Moore-Penrose pseudo inverse. Vora values that are close to 0 indicate estimated sensors close to the measured ground truth data. The derivation of Equation 30 is given in [21].

C. The CIE Lab ΔE Error Metric. The CIE Lab formula provides a measure for how close two reflectances appear to a human observer under a known illuminant. The error metric is designed for the XYZ colour space (which to a good approximation is a linear transform from the cone sensitivities in the eye). To calculate the ΔE error, the 3-vectors of XYZ responses are measured from two surfaces and then converted to corresponding Lab triplets. The Euclidean distance between these triples roughly correlate with our perception of colour difference with a ΔE error of one signifying, on average, just a noticeable difference [37].

Before we can use the metric for camera sensitivities we need to - via a calibration procedure - best map the camera rgb responses to corresponding XYZ coordinates and then to Lab values. We adopt the methodology defined in [38]. Taking the 1995 reflectances measured in [39] and CIE D65 illumination we numerically calculate using Equation 2 the rgb values of actual versus estimated rgb values and transform these into Lab Errors. The mean, median and 95% quantile errors are then used to assess how well we have estimated the camera sensors.

C. Spectral Estimation for Raw Images

In Figure 5, the ground truth spectral sensor sensitivity functions [10] from Nikon (top) and Sigma (bottom) are plotted in solid lines. The recovered sensors using our Rank-based method are plotted with dashed lines. All sensors are scaled to have a maximum sensitivity of one. Overall curve (of the sensor) of the estimated spectral sensitivity curve (of the sensor) are similar to the NPL ground truth data throughout the spectrum.

We evaluate the sensitivity estimation returned by the Rank-based method with the prior-art in Table 1. The first 5 columns of this table show, respectively, the spectral recovery error for the R, G and B colour channels; the average spectral recovery over all channels and the Vora Value. Bold-face numbers denote the best statistic for each column. The method with smallest statistic is the Subspace method of [11]. This method uses Equation 14 where a 2-dimensional sensor basis is used (derived by a PCA analysis of spectral sensitivity functions of 28 cameras). By chance the measured Nikon camera is very well modelled by this basis set.

However, we see that all methods have a low Vora Value indicating that the rgb values modelled by the recovered sensors are close to those of ground truth sensors. Encouragingly, the average spectral recovery method of the Rank-based method (excluding the 2-d subspace method) is similar to the prior-art. That is, a method that uses only the rank orders of the responses, works equally well as methods which fit the recorded data (we achieve the same performance using fewer of the degrees of freedom in the data).

In columns 6 to 10 of Table 1, we summarise the recovery performance for the Sigma Merrill camera. Again the Rank-based method works well and performs similar to Quadratic programming, in terms of the Vora Value. Notice that for this data set,

the subspace method works least-well (because the Sigma sensitivities are not in the assumed subspace).

Table 2 reports the CIE Lab error. Again the Rank-based method delivers good performance. Broadly, it supports performance similar to the prior-art.

D. Spectral Estimation for Rendered Images

As described in section 3.B. a rendered image (a typical jpeg from a smart phone) is non-linear. To a tolerable approximation the raw image colours are mapped by a 3x3 colour correction matrix and then tone-curves are applied to the colour corrected rgb values to make the final image (see Figure 1 for a raw and rendered pair). Because the Rank-based method is based on rank orders only, it should be able to recover the effective sensitivities of a rendered image. Effective sensitivities are the camera sensors multiplied by the colour correction matrix.

To evaluate how well we can recover the effective sensitivities we use the rendered image of the SG Macbeth - shown in the right of Figure 1 - colour chart as an input to the Rank-based method resulting in the estimated sensors illustrated with dashed lines in Figure 6. Here the ground truth spectral sensitivity functions (solid lines) are the NPL sensors [10], multiplied by a colour correction matrix M which was found using a calibration procedure described in [18]. Visually, the Rank-based method is able to recover the effective sensitivities of our two cameras.

Table 3 reports the error metrics (in the same format of Table 1). For both Nikon and Sigma cameras, we observe high spectral recovery errors and Vora values for all the three cited antecedent methods. This illustrates their inability to recover effective sensitivities in the face of the non-linearity present in a rendered image. However, the Rank-based method continues to work well. The CIE Lab errors are reported in Table 4 reaffirm this result. Rank-based delivers good estimation. In contradiction, all the prior-art are unable to accurately model colour outputs, when the calibration data is non linear.

5. RANK-BASED VERSUS LINEARISATION

The reader might wonder whether the poor performance of the prior-art methods working on rendered data might be mitigated by linearising the data first. We test this idea here. We consider two type of linearisation. First, we assume that a rendered image might be made approximately linear by undoing the gamma (e.g. undoing the sRGB gamma [35], roughly raising the image to the power 2.2). Second, we use the achromatic colours in a colour checker - here the 6 patches in the central part of the SG colour checker - to linearise the data. With respect to the second method, the idea is to plot the camera responses against the % spectral reflectance for the different grey-reflectances. Using linear interpolation we can define a function that, at least approximately, linearises the data (for more discussion see [40]).

In Table 5 we show the recovery errors for the prior-art methods for the simple gamma-only linearisation. It is clear this method is not sufficient for linearising the data (the sensor spectral estimation is poor). In Table 6, the 6 patch linearisation is shown. While the sensor spectral sensor estimation is improved, the recovery still has large errors. Neither linearisation recovers the effective sensitivities as well as the Rank-based method (last row of Table 3).

Initially, we were surprised that linearising the data with the achromatic scale led to such poor sensor estimation. After all, the prior-art does work tolerably well on actual linear data.




Fig. 5. Solid lines, the Nikon D5100 (top) and Sigma spectral sensor sensitivities and estimated Nikon spectral sensor sensitivities (bottom). Dashed lines are the sensors recovered by Rank-based spectral estimation.

Spectral Estimation	Nikon					Sigma				
	SE_R	SE_G	SE_B	$\mu(SE)$	Vora	SE_R	SE_G	SE_B	$\mu(SE)$	Vora
Regularisation	30	9	34	24	9	34	16	11	20	11
Subspace	11	6	7	8	1	46	42	54	47	56
QP	32	9	13	18	5	23	11	7	14	5
Rank-based	27	15	20	21	4	20	16	15	17	4

Table 1. The percentage spectral error of each channel, the average across all channels, and the Vora values are shown for raw images from Nikon (first 5 columns) and Sigma (last 5 columns) cameras.

Spectral Estimation	Nikon			Sigma		
	Mean	Median	95% percentile	Mean	Median	95% percentile
Regularisation	2	1.4	5	3.9	2.9	9.7
Subspace	1.8	1.1	5.4	18.3	12.7	49.5
QP	1.8	1.1	5.6	2.3	1.6	6.6
Rank-based	2	1.2	6	2.3	1.5	7

Table 2. CIE Lab ΔE errors between the predicted and estimated camera responses to the 1995 sample reflectances [39] under the D65 light [38] from Nikon (first 3 columns) and Sigma (last 3 columns) cameras.




Fig. 6. Solid lines, the Nikon D5100 (top) and Sigma spectral sensitivities and estimated Nikon spectral sensor sensitivities (bottom). Dashed lines are the sensors recovered by Rank-based spectral estimation using rendered image.

Spectral Estimation	Nikon					Sigma				
	SE_R	SE_G	SE_B	$\mu(SE)$	Vora	SE_R	SE_G	SE_B	$\mu(SE)$	Vora
Regularisation	77	77	50	68	50	44	65	54	54	31
Subspace	68	42	18	43	20	53	20	23	32	15
QP	73	69	53	65	48	53	61	49	54	25
Rank-based	39	16	21	25	8	27	19	16	21	5

Table 3. The percentage spectral error of each channel, the average across all channels, and the Vora values are shown for rendered images from Nikon (first 5 columns) and Sigma (last 5 columns) cameras.

Spectral Estimation	Nikon			Sigma		
	Mean	Median	95% percentile	Mean	Median	95% percentile
Regularisation	8.3	5.9	20	4.4	3.1	10.2
Subspace	4.1	3	11.6	4.1	3	11.9
QP	5.8	3.3	17.1	5.3	3.2	15.2
Rank-based	2.3	1.6	6.9	1.8	1.2	5.6

Table 4. CIE Lab ΔE errors between the predicted and estimated colour corrected camera responses to the 1995 sample reflectances [39] under the D65 light [38] from Nikon (first 3 columns) and Sigma (last 3 columns) cameras.

Spectral Estimation	Nikon					Sigma				
	SE_R	SE_G	SE_B	$\mu(SE)$	Vora	SE_R	SE_G	SE_B	$\mu(SE)$	Vora
Regularisation	62	64	71	66	45	42	83	50	58	37
Subspace	72	30	34	45	30	51	56	50	52	30
Quadratic programming	54	60	55	57	37	38	67	58	54	34
Rank-based	38	16	21	25	8	27	19	16	21	5

Table 5. The percentage spectral error of each channel, the average across all channels, and the Vora values are shown for rendered images linearised by gamma correction from Nikon (first 5 columns) and Sigma (last 5 columns) cameras.

Spectral Estimation	Nikon					Sigma				
	SE_R	SE_G	SE_B	$\mu(SE)$	Vora	SE_R	SE_G	SE_B	$\mu(SE)$	Vora
Regularisation	58	47	41	48	27	43	51	54	49	25
Subspace	77	36	41	51	37	41	25	26	31	13
Quadratic programming	63	13	19	32	16	57	28	28	38	20
Rank-based	38	16	21	25	8	27	19	16	21	5

Table 6. The percentage spectral error of each channel, the average across all channels, and the Vora values are shown for rendered images linearised by interpolation from Nikon (first 5 columns) and Sigma (last 5 columns) cameras.

Our results would suggest that a small departure from linearity would hamper the prior-art from recovering the sensors. To test this idea we wished to measure how well a method recovers the spectral sensitivities of a camera as a function of increasing non-linearity.

Using our raw data as a starting point we made the data slightly non-linear by raising it to the power ρ . When $\rho = 1$ our data is linear and the recovery results shown in Table 1 hold. Now, we can increase or decrease the power, $\rho = 1.1$ or $\rho = 0.9$, and in either case the data remains quite linear. Even larger, or even smaller powers ($\rho = 2$ or $\rho = 0.5$) make the data progressively non-linear. In the left panel of Figure 7, for each of the 3 prior-art methods, we plot the Vora Value of the recovered sensitivities as a function of the ρ power. It is immediately clear how sensitive the prior-art is to the linearity of the data. Even small departures from linearity has a marked affect on the accuracy of the spectral sensitivity estimation. In the right of Figure 7 we plot the same result as a function of average Spectral Estimation error.

For each of the prior-art methods, even very small non-linearities can lead to a doubling in recovery error. Of course, the Rank-based method continues to work well in the face of small and large non-linearities. For the non-linearities there is no change in the Rank-based estimated sensors.

6. CONCLUSION

In this paper, we have shown how the rank order of camera responses is a powerful tool for estimating the camera spectral sensitivity functions. We have described how the rank orders of image rgb values impose half-space constraints as ranked pair splits the space where the spectral sensor sensitivity must lie in one of the two parts. A plausible sensor estimate, lies in this region where all half-space constraints intersect which can be found using constrained quadratic programming. The Rank-based method works well on linear raw (as well as prior




Fig. 7. Plot of Vora (left) and averaged spectral error (right) calculated for Nikon raw image when raised to power values indicated on the x axes.

art), but delivers a step change in performance given non-linear data. A key example of the latter is rendered images (e.g. normal Jpegs) where we show we can deliver the estimated sensitivities of a camera. In contradiction, the prior art does not work for rendered data even when it is linearised.

REFERENCES

1. N. Shimano, K. Terai, and M. Hironaga, "Recovery of spectral reflectances of objects being imaged by multi-spectral cameras," *J. Opt. Soc. Am. A*, **24**(10), 3211-3219 (2007).
2. F.H. Imai, and R. Berns, "Spectral Estimation Using Trichromatic Digital Cameras", in *Proceedings of International Symposium on Multispectral Imaging and colour Reproduction for Digital Archives*, (Society of Multispectral Imaging, 1999), pp. 42-49.
3. S. Han, I. Sato, T. Okabe, and Y. Sato, "Fast spectral reflectance recovery using dlp projector," in *Proc. Asian Conference on Computer Vision*, 2010.

4. G.D. Finlayson, P.M. Hubel, and S. Hordley, "Colour by correction," in Proceedings of the fifth colour Imaging Conference: colour Standards and colour Measurements (IST, 1997), pp. 6-11.
5. J. Park, M. Lee, M. D. Grossberg, and S. K. Nayar, "Multispectral imaging using multiplexed illumination," in *P of International Conf, on Computer Vision*, 2007.
6. G.D. Finlayson, and M.S. Drew, "White-point preserving colour correction," in *Proceedings of the Fifth colour Imaging Conference: colour Standards and colour Measurements* (IST, 1997), pp. 258-261.
7. G.D. Finlayson, and M.S. Drew, "Constrained least-squares regression in colour space," *Journal of Electronic Imaging*, pp. 484-493 (1997).
8. A. Forsyth, "A novel algorithm for colour constancy," *International Journal on Computer Vision*, **5**(1), 536 (IJCV, 1990).
9. G. Finlayson, S. Hordley, and P. Hubel, "Colour by correlation: A simple, unifying framework for colour constancy," in *IEEE Trans on Pattern Analysis and Machine Intelligence*, **23**(11), 12091221 (IEEE, 2001).
10. M. Mohammadzadeh Darrodi, G.D. Finlayson, T. Good-man, M. Mackiewicz, "A reference data set for camera spectral sensitivity estimation" *J. Opt. Soc. Am. A.* (2014).
11. J. Jiang, D. Liu, J. Gu, and S. Susstrunk, "What is the Space of Spectral Sensitivity Functions for Digital colour Cameras," in *IEEE Workshop on the Applications of Computer Vision* (IEEE, 2013), pp. 168-179.
12. P.L. Vora, J.E. Farell, J.D. Tietz, and D. Brainard, "Digital colour cameras - 2 Spectral response," in HP Technical Report (HP, 1997).
13. P.M. Hubel, D. Sherman, and J.E. Farell, "A Comparison of Methods of Sensor Spectral Sensitivity Estimation," in *proceedings of colour Imaging Conference: colour Science, Systems and Applications*, (IST, 1994), pp. 45-48.
14. G. D. Finlayson, S. Hordley, and P.M. Hubel, "Recovering device sensitivities with quadratic programming," in *Proceedings of The Sixth colour Imaging Conference: colour Science, Systems, and Applications*, 1998, pp. 90-95.
15. R. Martin, Z. Arno, and K. Reinhard, "Practical spectral characterization of trichromatic cameras," in *P of SIGGRAPH (ASIA, 2011)*.
16. P. Urban, M. Desch, K. Happel, and D. Spiehl, "Recovering camera sensitivities using target - based reflectances captured under multiple LED-illuminations," in 16th Workshop on colour Image Processing (2010), pp. 295-301.
17. R. Ramanath, W. E. Snyder, Y. Yoo, and M. S. Drew, "Color Image Processing Pipeline," in *IEEE Signal processing magazine* (2005).
18. S. J. Kim, H. T. Lin, Z. Lu, S. Susstrunk, S. Lin, and M. S. Brown, "A new in-camera imaging model for colour computer vision and its application," in *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **34**(12), 22892302 (IEEE, 2012).
19. R.W.G. Hunt, "The Reproduction of colour," 6th Ed. (Voyageur Press, 2004).
20. J. Singnoo, and G.D. Finlayson, "RGBE vs Modified TIFF for Encoding High Dynamic Range". 4Th Conference on colour in Graph., Imagi. and Vision, 431-436.36 (IV, 2010).
21. J. Trussell, and M. Vrhel, "*Fundamentals of Digital Imaging*," Cambridge University Press, (2009).
22. B. Smith, C. Spiekermann, and R. Sember, "Numerical methods for colorimetric calculations: Sampling density requirements," *Color Res. Appl.*, **1**, 394401 (1992).
23. E. Moore, "On the reciprocal of the general algebraic matrix", *Bull Am. Math. Soc.*, (1920), 26 pp. 394-395.
24. J. A. Worthey, "Spectrally Smooth Reflectances that Match," *Color Res. Appl.*, **19**, 395396 (1994).
25. J.P.S. Parkkinen, J. Hallikainen and T. Jaaskelainen, "Characteristic spectra of Munsell colors," *J. Opt. Soc. Am. A.* **6**, pp.318-322 (1989).
26. Y.H. Hardeberg, H. Brettel, and F.J. Schmitt, "Spectral characterization of electronic cameras, *Electronic Imaging: Processing, Printing, and Publishing*," (1998), pp. 100-109.
27. R. L. Burden and D. Faires, "*Numerical Analysis*," 8th ed. Brooks/Cole (2005).
28. A.N. Tikhonov, V.Y. Arsenin, "*Solution of Ill-posed Problems*," Washington (Winston Sons, 1977).
29. B. Dyas, "Robust colour sensor response characterization," in *P of Eighth colour Imaging Conference* (2000), pp. 144-148.
30. P.C. Hansen, and D.P. O'Leary, "The use of the L-curve in the Regularisation of discrete ill-posed problems," *SIAM J. Sci Comput.*, (SIAM, 1993), **14**, 1487-1503.
31. J. Gu ; [Availablefrom:http://www.cis.rit.edu/jwgu](http://www.cis.rit.edu/jwgu).
32. H. Zhao, K. Rei, T.T. Robby, and I. Katsushi, "Estimating basis functions for spectral sensitivity of digital cam-eras," in Meeting on Image Recognition and Understanding (2009), pp. 7-13.
33. J. J. Fuch, "Linear programming," in spectral estimation, Application to array processing, **6**, 3161 - 3164 (1996).
34. M. Mohammadzadeh Darrodi, [Availablefrom:http://www.spectralesimation.wordpress.com/data/](http://www.spectralesimation.wordpress.com/data/).
35. M. Stokes, M. Anderson, S. Chandrasekar, R. Motta. *A Standard Default Color Space for the Internet - sRGB*, Tech. Report, Hewlett-Packard, Microsoft, 1996, [Availablefrom:http://www.w3.org/Graphics/Color/sRGB](http://www.w3.org/Graphics/Color/sRGB).
36. J. Morovic, "*Color Gamut Mapping*," (Wiley,2008).
37. G. Wyszecki and W. S. Stiles, "Color Science: Concepts and Methods, Quantitative Data and Formulas," John Wiley Sons, Inc., London, 1967, **71**, 628 (1969).
38. J. Vazquez-Corral, D. Connah, and M. Bertalm, "Perceptual Color Characterization of Cameras," *Sensors*, **14**(12) pp. 23205-23229 (2014).
39. K. Barnard, L. Martin, B. Funt, and A. Coath, "A Data Set for Colour Research," *Color Research and Application*, **27**, 3, pp. 147-151 (2002).
40. P. Vora, L. Poorvi, J.E. Farell, J. Tietz, D.H. Brainard, *Digital Color Cameras - 1 - Response Models*, HP Labs Technical Reports, HPL-97-53.