Analytic Solution for Separating Spectra into Illumination and Surface Reflectance Components

Mark S. Drew and Graham D. Finlayson

Mark S. Drew
School of Computing Science
Simon Fraser University
Vancouver, B.C.
Canada V5A 1S6
mark@cs.sfu.ca

Graham D. Finlayson
School of Computing Sciences
University of East Anglia
Norwich
UK NR4 7TJ
graham@cmp.uea.ac.uk

Mailing address:
Dr. Mark S. Drew
School of Computing Science,
Simon Fraser University,
Vancouver, B.C.,
Canada V5A 1S6

Telephone: 1–604–291–4682
Fax: 1–604–291-3045

Running Head: Analytic Solution for Separating Spectra

http://www.cs.sfu.ca/cs/∼mark
http://www.cmp.uea.ac.uk/people/graham

The measured light spectrum is the result of an illuminant interacting with a surface. The illuminant spectral power distribution multiplies the surface spectral reflectance function to form a *colour signal* — the light spectrum that gives rise to our perception. Disambiguation of the two factors, illuminant and surface, is difficult, without prior knowledge. Previously,[1,2] one approach to this problem applied a finite-dimensional basis function model in order to recover the separate illuminant and surface reflectance components that make up the color signal, using principal component bases for lights and for reflectances. In this paper, we introduce the idea of making use of finite-dimensional models of logarithms of spectra for this problem. Recognizing that multiplications turn into additions in such a formulation, we can replace the original iterative method with a direct, analytic algorithm, with no iteration.

Moreover, in the new, logarithm-based approach, it is straightforward to further design new basis functions, for both illuminant and reflectance simultaneously, such that the initial basis function coefficients derived from the input colour signal are optimally mapped onto separate coefficients that produce spectra that more closely approximate the illuminant and the surface reflectance, for any given dimensionality. This is accomplished using an extra bias correction step that maps the analytically determined basis function coefficients onto the optimal coefficient set, separately for lights and surfaces, for the training set. The analytic equation plus the bias correction is then used for unknown input colour signals.

In tests, we found that in terms of RMS error in fact the original, non-logarithm, method does comparably to a logarithm-based additive method. However, once the extra correction step is added, the log-based method can outperform the original method. As well, a log-based approach is several orders of magnitude faster.

Also, we investigate the bias correction idea per se, considering the simplest possible basis. We determine that the bias correction idea produces fairly good results no matter what initial basis set is used.

The main contribution of the paper is to devise an additive-based approach, which speeds up calculation impressively because it is analytic.

Keywords (OCIS): (150.0150) Machine vision: Machine vision; (150.2950) Machine vision : Illumination

OCIS codes: 150.0150, 150.2950.

1. Introduction

Spectral data has become more available recently, and is of primary importance for tasks such as spectral imaging, digital archiving, and reproduction.[3,4] On a fundamental level, light is more than simply the three numbers encapsulated in RGB pixel values. If we assume spectral information is captured at least every 10nm, then over the visible spectrum from 400 to 700nm, data consists of 31 samples. The acquisition and use of full-spectral information is effective in a number of practical settings, e.g. in visualization of data,[5] in medical imaging (e.g., Refs. [6, 7]), or in accurate modelling using spectrum-based rendering for product development, internet imaging, accurate colour appearance modelling, and precise reflectance modelling for applications such as skin colour and texture.[8] Multispectral imaging is crucial for such spectrum-critical tasks as “digital rejuvenation” in the conservation of archival artworks.[9]

A basic issue in computer vision is the recovery of intrinsic images,[10] consisting of the illuminant contribution and reflectance contribution *separated* out of the measured light spectral data. An initial approach to this problem in the spectral realm was Ref. [1], which recovered both illuminant and reflectance spectral data by making use of statistical distributions for each.

This original separation algorithm[1] was based on finite-dimensional models of illuminants and reflectances.[11] The investigation of the power of linear models was continued in Ref. [12], which focussed on the accuracy of spectral models when filtered through human or device sensors. Here, we concentrate on whole-spectrum error measures since entire spectra are needed for such tasks as illuminant-change widgets for visualization,[5] and thus this work is independent of any particular sensor set; as well, we go over to logarithmic spectra, with a very large increase in speed. Logarithms of spectra in themselves are not new. In Ref. [13], a principal component analysis on the logarithm of the illuminant was used in order to cast the colour-signal formation process into a simple form for very low dimensionality, for use in studying illumination invariants. A low-dimension finite model was used in Ref. [14] for establishing the best spectrum that could be reliably associated with a particular observed RGB triple. Logarithmic basis functions were used in Ref.[15] to model reflectance data; a logarithmic image is also used in Ref.[16]. And in Ref. [17], a finite model of logarithms of spectra is used as a tool to help characterize the distribution of colours in an image.

Most recently, Chandra and Healey[2] re-examined the original separation method,[1] applying the method to more materials and a much wider set of wavelengths (from 350nm to 1740nm). Surprisingly, they found the method could be used effectively with only a 2-dimensional model for lights, but a quite high dimension for surfaces (dimension 11). Their main concern was to examine how the method performed as dimensionality was raised, and to develop a reasonable algorithm for deciding when adding more dimensions would actually decrease, rather than increase the accuracy of the method. Usually, the number of basis functions suggested for daylights is typically 3,[18] while that for reflectances might be as high as about 8.[19]

In this paper, we examine how using bases for logarithms of spectra affects the light–surface separation process. To begin, in §2 we briefly outline the original, non-logarithm based separation method. Then, in §3, we set out a new, analytic formulation based on logarithms, for separating light from surface spectrum. This is based on optimizing the description of the colour signal, as a *sum* of light plus surface logarithm, but is not simply a projection: instead, the light basis and surface basis collaborate to produce the best approximation of the incoming signal. Notwithstanding this coupling, the resulting equation is *linear* and is therefore very fast. In §4 a method is given for improving matters by optimizing for both coefficients as well as the basis vectors themselves: we start with the log basis function coefficients obtained using the above linear set of equations, but then carry out a further bias correction from this set separately onto the correct (training set, and thus known) log basis function coefficients for the lights and the reflectances separately. This is used in a training phase for determining a final set of basis vectors, and then applied (as in §3) to new, unknown signals. In order to account for noise, we use Tikhonov regularization in forming the regression. The role of the extra regression step per se. is investigated in §5, where the simplest basis is used for both lights and surfaces: the orthonormal set of basis functions corresponding to the projector onto the unit vector, which we call the Uniform basis functions in log-spectral space. It is found that the bias correction itself is responsible for much of the accuracy of the method. In §6, accuracy of all methods is explored by testing against a wide set of input colour signals Results show that the minimum error is achieved by optimizing the basis sets by the further regression step, and starting from a principal components log basis.

2. Non-Logarithm-Based Separation

To briefly recapitulate the method of Ref. [1], we model a measured colour signal $\hat{C}\,(\lambda)$ (using hats to denote absolute, non-logarithm quantities) as comprised of the product of illuminant spectral power distribution (SPD) $\hat{E}\,(\lambda)$ and surface spectral reflectance function $\hat{S}\,(\lambda)$:

$$\hat{C}\,(\lambda) \;=\; \hat{E}\,(\lambda)\,\hat{S}\,(\lambda) \tag{1}$$

Suppose now that we have determined a subspace approximation for illuminants, $\hat{E}_{\,i}(\lambda)$, $i = 1..m$ with $m \leq s$ for spectra sampled at s wavelengths over the visible spectrum. E.g., s used in this paper is 31. It has been shown[18] that about $m = 3$ basis functions are sufficient to account for most of the variance in daylight illuminant spectra.

As well, if we use $n \leq s$ basis functions for reflectances, with reflectance basis $\hat{S}_{\,j}(\lambda)$, then

$$\hat{E}\,(\lambda) \;\simeq\; \sum_{i=1}^{m} \hat{\epsilon}_{\,i} \hat{E}_{\,i}(\lambda)\;,\; \hat{S}\,(\lambda) \simeq \sum_{j=1}^{n} \hat{\sigma}_{\,j} \hat{S}_{\,j}(\lambda)\;,\; \hat{C}\,(\lambda) \;\simeq\; \sum_{i=1}^{m}\sum_{j=1}^{n} \hat{E}_{\,i}(\lambda)\hat{S}_{\,i}(\lambda)\hat{\epsilon}_{\,i}\hat{\sigma}_{\,j} \tag{2}$$

with coefficients $\hat{\epsilon}_{\,i}$ and $\hat{\sigma}_{\,j}$.

In order to solve for $\hat{\epsilon}_{\,i}$ and $\hat{\sigma}_{\,j}$ given driving data $\hat{C}\,(\lambda)$, and hence disambiguating the illuminant and reflectance from the colour signal, we are faced with a nonlinear problem. The problem can be expressed in terms of a least-squares minimization. Suppose we wish to find coefficients $\hat{\epsilon}_i$, $\hat{\sigma}_j$ that minimize the objective function

$$min\,\hat{\mathcal{I}} \;=\; \int \left[\sum_{i=1}^{m} \sum_{j=1}^{n} \hat{\epsilon}_i \hat{\sigma}_j \hat{E}_i(\lambda) \hat{S}_i(\lambda) \;-\; \hat{C}(\lambda) \right]^2 d\lambda \qquad (3)$$

where we integrate over the visible. Setting partial derivatives with respect to $\hat{\epsilon}_i$ and $\hat{\sigma}_j$ to zero in turn, we have a coupled, nonlinear set of equations

$$\hat{\boldsymbol{M}}\,\hat{\boldsymbol{\sigma}} \;=\; \hat{\boldsymbol{F}}\,\hat{\boldsymbol{\epsilon}} \;, \quad \hat{\boldsymbol{N}}\,\hat{\boldsymbol{\epsilon}} \;=\; \hat{\boldsymbol{F}}\,\hat{\boldsymbol{\sigma}} \;, \qquad (4)$$

where we define

$$\hat{M}_{qj} \;\equiv\; \sum_p \sum_i \hat{\epsilon}_i \hat{\epsilon}_p \hat{H}_{ipjq} d\lambda \;, \quad \hat{N}_{pi} \;\equiv\; \sum_q \sum_j \hat{\sigma}_j \hat{\sigma}_q \hat{H}_{ipjq} d\lambda \;,$$

$$\hat{H}_{ipjq} \;\equiv\; \int \hat{E}_p \hat{E}_i \hat{S}_q \hat{S}_j d\lambda \;, \quad \hat{F}_{ij} \;\equiv\; \int \hat{E}_i \hat{S}_j \hat{C}\, d\lambda$$

To solve, we first initialize $\hat{\boldsymbol{\epsilon}}$, and then solve for $\hat{\boldsymbol{\sigma}}$. Substituting into the first of the pair (4), we re-determine $\hat{\boldsymbol{\epsilon}}$. Iterating, we find both $\hat{\boldsymbol{\epsilon}}$ and $\hat{\boldsymbol{\sigma}}$. Essentially, we have created a nonlinear fixed-point problem, of the form $\hat{\boldsymbol{\epsilon}} = f(\hat{\boldsymbol{\epsilon}})$. Therefore we can evaluate our initial guess for uniqueness and convergence rate using the Contraction Mapping Theorem.[20] This involves numeric calculation of partial derivatives of $f()$ over a reasonable parameter region. In this paper, for an m-vector $\hat{\boldsymbol{\epsilon}}$, we initialize with the simplest vector $(1, 0, 0, ...)$, and then iterate until convergence or an iteration limit (1000) is reached, and do not examine for uniqueness — different starting points therefore might indeed produce different results.

One complication that arises from using the energy objective function (3) is that a nonuniqueness is present corresponding to the absolute intensity of the illuminant: if we multiply $\hat{E}(\lambda)$ by constant $\hat{\kappa}$ and also divide $\hat{S}(\lambda)$ by $\hat{\kappa}$, the objective $\hat{\mathcal{I}}$ is unaffected. This corresponds to multiplying $\hat{\boldsymbol{\epsilon}}$ by $\hat{\kappa}$ and dividing $\hat{\boldsymbol{\sigma}}$ by $\hat{\kappa}$. Therefore to evaluate the resulting $\hat{\boldsymbol{\epsilon}}$ we firstly re-set the first component of $\hat{\boldsymbol{\epsilon}}$ to 1 at every iteration, and then map the intensity to that of the actual illuminant when an error is calculated, in order not to bias an absolute error level.

To make clear just what algorithm is being discussed at any time, we shall name the algorithms developed in this paper. We call the original method, outlined above, **Method SEP**.

3. Logarithm-Based Separation

Since colour signal formation is a multiplicative process, separation should properly be placed in a logarithmic setting, and that is what is addressed in this paper. However, we should note that, once we move away from absolute errors, as in (3), we effectively introduce the effect of Weber's Law (see, e.g., Ref. [21]), in that errors in the logarithm result correspond to relative errors, in the original spectral vector space. Nevertheless, since logarithms are monotonic, we may assume that reducing errors in the log domain maps to a reduction in absolute errors. As well, all errors reported here correspond to exponentiated, non-logarithmic spectra.

Now we are mainly concerned with the *logarithm of spectra*, and will drop the hats when referring to these:

$$E(\lambda) \;=\; \log\left(\hat{E}\,(\lambda)\right)\;, \quad S(\lambda) \;=\; \log\left(\hat{S}\,(\lambda)\right)\;, \tag{5}$$

Introducing a basis decomposition for light and surface logarithmic spectrum, with different bases than in (2), we have

$$E(\lambda) \;\simeq\; \sum_i \epsilon_i E_i(\lambda)\,,\; S(\lambda) \simeq \sum_j \sigma_j S_j(\lambda)\,,\; C(\lambda) \;\simeq\; \sum_i \epsilon_i E_i(\lambda) \,+\, \sum_j S_j(\lambda)\sigma_j \tag{6}$$

The natural modification of (3) to logarithmic spectra is

$$min\,\mathcal{I} \;=\; \int \left[\sum_{i=1}^{m} \epsilon_i E_i(\lambda) \,+\, \sum_{j=1}^{n} \sigma_j S_j(\lambda) \;-\; C(\lambda)\right]^2 d\lambda \tag{7}$$

In terms of the $s \times m$ set of illuminant basis vectors $\boldsymbol{E}$, and the $s \times n$ set of reflectance basis vectors $\boldsymbol{S}$, replacing integration by summation this is written

$$min\,\mathcal{I} \;=\; \int \left[\boldsymbol{E}\,\boldsymbol{\epsilon} \;+\; \boldsymbol{S}\,\boldsymbol{\sigma} \;-\; \boldsymbol{c}\,\right]^2 d\lambda \tag{8}$$

with $\boldsymbol{c}$ the $s \times 1$ samples of an input colour signal. Setting partial derivatives with respect to $\boldsymbol{\epsilon}$ and $\boldsymbol{\sigma}$ to zero, we obtain a set of coupled equations which is now *linear*:

$$\begin{cases} \boldsymbol{M}\,\boldsymbol{\epsilon} \;+\; \boldsymbol{N}\,\boldsymbol{\sigma} \;=\; \boldsymbol{f} \\ \\ \boldsymbol{O}\,\boldsymbol{\epsilon} \;+\; \boldsymbol{P}\,\boldsymbol{\sigma} \;=\; \boldsymbol{g} \end{cases} \tag{9}$$

where

$$\begin{array}{ll} \boldsymbol{M} \;=\; \boldsymbol{E}^{\,T}\boldsymbol{E}\;, & \boldsymbol{N} \;=\; \boldsymbol{E}^{\,T}\boldsymbol{S}\;, \\ \boldsymbol{O} \;=\; \boldsymbol{N}^{\,T}\,, & \boldsymbol{P} \;=\; \boldsymbol{S}^{\,T}\boldsymbol{S}\;, \\ \boldsymbol{f} \;=\; \boldsymbol{E}^{\,T}\boldsymbol{c}\;, & \boldsymbol{g} \;=\; \boldsymbol{S}^{\,T}\boldsymbol{c} \end{array} \tag{10}$$

with $\boldsymbol{c}$ the particular, driving, colour signal being examined. Here, $\boldsymbol{M}$ is an $m \times m$ matrix, $\boldsymbol{N}$ is $m \times n$, and $\boldsymbol{P}$ is $n \times n$.

To solve at once for both $\boldsymbol{\epsilon}$ and $\boldsymbol{\sigma}$, we combine the above set into a single equation

$$\boldsymbol{A}\,\boldsymbol{\alpha} \;=\; \boldsymbol{h} \tag{11}$$

where $\boldsymbol{A}$ combines $\boldsymbol{M}$ and $\boldsymbol{N}$, in a block row, and $\boldsymbol{O}$ and $\boldsymbol{P}$ in a second block row, $\boldsymbol{\alpha}$ concatenates $\boldsymbol{\epsilon}$ and $\boldsymbol{\sigma}$ into an $(m+n)\times 1$ vector, and $\boldsymbol{h}$ concatenates $\boldsymbol{f}$ and $\boldsymbol{g}$ into an $(m+n)\times 1$ vector:

$$\boldsymbol{A} \;=\; \begin{pmatrix} \boldsymbol{M}\,, & \boldsymbol{N} \\ \\ \boldsymbol{N}^{\,T}\,, & \boldsymbol{P} \end{pmatrix} \;, \quad \boldsymbol{\alpha} \;=\; \begin{pmatrix} \boldsymbol{\epsilon} \\ \boldsymbol{\sigma} \end{pmatrix}\,, \quad \boldsymbol{h} \;=\; \begin{pmatrix} \boldsymbol{f} \\ \boldsymbol{g} \end{pmatrix} \tag{12}$$

The nonuniqueness of absolute scale of the illuminant still persists in this reformulation, but now it consists of an additive constant κ. Therefore we take the absolute intensity for the recovered E and S to be not recoverable by this method. Notice that the solution of (12) is indeed unique — the method analytically finds weights to make the log colour signal as well modelled as possible, so the minimization is doing it's job. However, here we have a different sub-goal: that of generating approximations of lights and surfaces. We guide this separation by utilizing eigen-subspaces, but intensity is not recovered. This issue is discussed further in §5.

Notably, the formulation (12) is not simply a projection of the colour signal vector onto each subspace formed by the illuminant and surface reflectance basis sets: instead, the two basis sets collaborate to produce the best overall match to the colour signal, coupling the two sets. Therefore, eq. (12) will do better matching the input signal than if the bases are used separately.

Since matrix $\boldsymbol{A}$ is the outer product of the $s \times (m+n)$ array $(\boldsymbol{E}\ ,\ \boldsymbol{S}\ )$ with itself, its maximum rank is $min(s,(m+n))$. If $\boldsymbol{E}$ =$\boldsymbol{S}$, then the rank is only m. Generally, it is possible to find basis sets $\boldsymbol{E}$ and $\boldsymbol{S}$ to make matrix $\boldsymbol{A}$ full-rank, provided $(m+n) \leq s$; but a rank check should be included, to signal using the Moore-Penrose pseudoinverse to determine $(\boldsymbol{\epsilon}\ ,\ \boldsymbol{\sigma}\ )^T$ via a truncated Singular Value Decomposition (SVD). For the analytic equation eq. (11), we will usually focus on low-dimension representations $(m+n) \leq s$ so the inversion is full-rank, with no significant noise magnification that regularization would be aimed at reducing. However, below, in § 4, we add Tikhonov regularization to guard against the effects of noise.

Note that, in the approach developed using logarithms, disambiguation of illuminant and reflectance is a non-iterative, analytic, and hence very fast algorithm. We start from knowledge of illuminant and reflectance bases, and derive weights $\boldsymbol{\epsilon}$ and $\boldsymbol{\sigma}$ in one step.

Of course, once logarithms are introduced we must guard against negative or even small values: in most image processing applications there are special cases to be considered and that is the case here. In our program, the practical and simple way of dealing with low values is to flag such and replace those samples in the log-signal with zero. Using the flag, exponentiated spectra are reconstituted with zero values at problem locations. Any generated negative values are clamped to zero.

The issue of precisely which bases to use is left unresolved at this stage. One naturally thinks of a principal component basis as capturing most of the information for each space, illuminants and reflectances, in an optimal fashion aimed at subspace reduction to a smaller dimension m and n. However, as pointed out by Healey,[2] if the illuminant and reflectance basis subspaces substantially overlap, then any separation algorithm will not succeed as well as possible, since again each component is confounded. This will hold true for the log version as well as for the original version of separation algorithm given above. Therefore in experiments we consider two kinds of fixed basis: principal component bases for illuminants and for reflectances, derived via an SVD; and also a nonorthogonal basis for each space derived via an Independent Component Analysis (ICA).[22]

For creation of eigenvectors, flagged small spectral values, that would create a problem in the log domain, are set to zero and then simply carried along.

We found that an ICA basis did not improve the log-separate algorithm given above in (11), and so do not include it below.

Note that ICA is aimed simply at blind source separation to explain additive contributions to a signal. Here we also have extra knowledge of the problem, namely that we wish to separate light from surface. Therefore we may ask further that bases be designed specifically for this task. The next section addresses this aim by re-mapping the log basis function coefficients derived from eq. (7) onto the known coefficients for lights and surfaces, separately. This step amounts to a bias correction aimed at re-fashioning the light and surface basis functions themselves, with a view to using these in the specific application at issue here. Results are found to improve substantially using this extra step.

4. Optimizing Bases via Bias Correction

One problem with using fixed, predefined bases for lights and for surfaces, defined independently of the task at hand here, is that the subspace overlap is fixed, for given subspace dimensions m and n. Simply increasing dimensions will not necessarily help matters, since this may have the effect of increasing this overlap: after all, if we allow both m and n to be the largest possible dimension, s, where s is the number of wavelength samples, then overlap is complete and there is no different information in the basis for lights than in the basis for surfaces. We would expect a separation algorithm to fail completely in this situation.

However, here we have extra knowledge in our training set: we actually know the correct illuminants and reflectances. If we insert an additional regression step, for each of lights and surfaces, taking us from the recovered (training set) vectors $\boldsymbol{\alpha}$ to the correct vectors $\boldsymbol{\epsilon}$ and $\boldsymbol{\sigma}$, we will be fully utilizing all the information we have.

Thus we wish to *guide* the light and surface coefficients determined from eq. (11) by *forming a regression*, for the training set, from the $(m+n)$-vector solutions $\boldsymbol{\alpha}$ onto the correct illumination-coefficient vectors $\boldsymbol{\epsilon}$, and separately onto the correct surface-coefficient vectors $\boldsymbol{\sigma}$. This amounts to prepending an extra $m \times (m+n)$ regression matrix $\boldsymbol{Q}$ onto the inverse of matrix $\boldsymbol{A}$, to guide solution vectors towards the optimum illuminant-coefficients. As well, we also prepend another regression matrix, $\boldsymbol{R}$, to guide solutions towards the correct reflectance-coefficients.

But this means that, effectively, we have optimally changed the basis sets, since the resulting combined solution will combine the effect of using *both* sets $\boldsymbol{E}$ and $\boldsymbol{S}$, to form new coefficients.

Note again that the strategy to be adopted is to take a large training set of known illuminants and surfaces, derive optimized basis sets $\boldsymbol{E}$ and $\boldsymbol{S}$, and then apply the simple log-separation algorithm (11) plus regressions to new, measured, testing colour signals that derive from unknown lights and surfaces. In this new approach, we include regression matrices that optimally take solutions onto light and surface coefficients for the training set.

4.A. Regression onto Light and Surface Coefficients

Since we wish to start from a large set of lights and surfaces, let us define new symbols that encompass these sets. Let $\widetilde{\mathscr{C}}$ be the set of $s \times N$ training set colour signals, with N the number of cases: i.e., N is the number of lights in our training data set times the number of surfaces used. We also define a set of illuminants $\widetilde{\mathscr{E}}$ of the same size, $s \times N$, by replicating illuminants, and similarly an $s \times N$ set of reflectances $\widetilde{\mathscr{S}}$. Recall that the $s \times m$ basis for lights is $\boldsymbol{E}$ and the $s \times n$ set of reflectance basis vectors is $\boldsymbol{S}$. As well, let the set of $m \times N$ light-basis coefficients for all the training cases be $\widetilde{\boldsymbol{\epsilon}}$ and the $n \times N$ surface-basis coefficients be $\widetilde{\boldsymbol{\sigma}}$. And finally let the set of all $(m+n)$-vectors $\boldsymbol{\alpha}$ be $\widetilde{\boldsymbol{\alpha}}$.

Once we have determined the set $\widetilde{\boldsymbol{\alpha}}$ for all training set colour signals, we wish to optimally map $\widetilde{\boldsymbol{\alpha}}$ onto the known values $\widetilde{\boldsymbol{\epsilon}}$ and, additionally, onto the known values $\widetilde{\boldsymbol{\sigma}}$ using another regression. Thus we wish to solve for two regression matrices, $\boldsymbol{Q}$ and $\boldsymbol{R}$, via two objective functions

$$
\begin{aligned}
min\,\mathcal{I}_1 \;&=\; \sum \|\widetilde{\boldsymbol{\epsilon}} \;-\; \boldsymbol{Q}\,\widetilde{\boldsymbol{\alpha}}\,\|^2 \\
min\,\mathcal{I}_2 \;&=\; \sum \|\widetilde{\boldsymbol{\sigma}} \;-\; \boldsymbol{R}\,\widetilde{\boldsymbol{\alpha}}\,\|^2
\end{aligned}
\tag{13}
$$

It makes sense to use two separate regressions, rather than mapping onto known vectors $\boldsymbol{\alpha}$ in a single regression, since in that case one would be trying to minimize errors in both light and surface simultaneously, not making use of the separate information we have for the training set. Clearly, the minimum L_2-norm solutions in this simple situation are the two pseudoinverses

$$
\begin{aligned}
\boldsymbol{Q} \;&=\; \widetilde{\boldsymbol{\epsilon}}\;\widetilde{\boldsymbol{\alpha}}^T \left(\widetilde{\boldsymbol{\alpha}}\,\widetilde{\boldsymbol{\alpha}}^T\right)^{-1} \;, \\
\boldsymbol{R} \;&=\; \widetilde{\boldsymbol{\sigma}}\;\widetilde{\boldsymbol{\alpha}}^T \left(\widetilde{\boldsymbol{\alpha}}\,\widetilde{\boldsymbol{\alpha}}^T\right)^{-1} \;,
\end{aligned}
\tag{14}
$$

Here, $\boldsymbol{Q}$ is an $m \times (m+n)$ matrix, and $\boldsymbol{R}$ is $n \times (m+n)$.

Thus combining these into a single $(m+n) \times (m+n)$ matrix $\boldsymbol{T}$, we have a simple form

$$
\widetilde{\boldsymbol{\alpha}} \;\rightarrow\; \widetilde{\boldsymbol{\alpha}}_{new} \;=\; \boldsymbol{T}\,\widetilde{\boldsymbol{\alpha}}\,, \quad \boldsymbol{T} \;=\; \begin{pmatrix} \boldsymbol{Q} & 0 \\ 0 & \boldsymbol{R} \end{pmatrix}
\tag{15}
$$

4.B. Tikhonov Regularization

So far we have not considered the effect of noise on our inversion problem. Instead, we have suggested simply using a Moore-Penrose pseudoinverse in any case where rank reduction is suspected, in order to carry out an inversion based on a truncated sum over eigenvectors. However, a straightforward extension of eq. (13) to take into account noise is to add a regularizing term,[23] penalizing large-norm solutions, into the solution eq. (14):

$$\begin{aligned}
\boldsymbol{W} &= \tfrac{1}{(m+n)} \sum diag\left(\widetilde{\boldsymbol{\alpha}}\,\widetilde{\boldsymbol{\alpha}}^T\right) \boldsymbol{I}_{\,(m+n)} \\
\boldsymbol{Q} &= \widetilde{\boldsymbol{\epsilon}}\ \widetilde{\boldsymbol{\alpha}}^T \left(\widetilde{\boldsymbol{\alpha}}\,\widetilde{\boldsymbol{\alpha}}^T + \lambda \boldsymbol{W}\right)^{-1}, \\
\boldsymbol{R} &= \widetilde{\boldsymbol{\sigma}}\ \widetilde{\boldsymbol{\alpha}}^T \left(\widetilde{\boldsymbol{\alpha}}\,\widetilde{\boldsymbol{\alpha}}^T + \lambda \boldsymbol{W}\right)^{-1},
\end{aligned} \tag{16}$$

where $\boldsymbol{I}_{\,(m+n)}$ is the $(m+n) \times (m+n)$ identity matrix, and parameter λ is typically in the range 10^{-2} – 10^{-6}. Here we use $\lambda = 10^{-3}$.

Below, in § 6, we shall see that the addition of a regression step substantially reduces errors in recovering light and surface spectra. The two steps of regression and regularization will greatly ameliorate the noise effects generated by logs of low values.

We call the above algorithm **Method REGLOGSEP**, to distinguish carrying out a regularized regression for the training set, for each of lights and surfaces, from the previous analytic method that omits a regression. For a new, testing colour signal, we first apply the simple, linear method LOGSEP, eq. (11), using a principal component basis for each of lights and surfaces, and then follow with the regression also determined from the training set.

5. Uniform-Vector Basis

5.A. Use of Simplest Basis in REGLOGSEP

The additive constant κ in eq. (11) is worth thinking about in more detail. First we observe that it is likely that bright and dim lights with the same spectral shape will occur and, similarly, we will have reflectances that have the same shape but different average reflectivity. That is, the constant vector should be a member of the span of both the light basis and the reflectance basis. In fact, each subspace, the span of $\boldsymbol{E}$ and the span of $\boldsymbol{S}$, contains the constant spectrum almost exactly, for dimension above 2. We generally expect the light and surface subspaces to overlap, and this is just a special case. As in the non-log space, we will set the magnitude of light (or surface) and so effectively fix the magnitude of the coordinate governing brightness for the light basis (or for the surface basis). Indeed, without loss of generality we could set the contribution of the uniform vector to reflectance to be zero, by projecting onto the subspace orthogonal to this direction. Then the surface basis would be completely orthogonal to the direction of intensity variation (cf. “black lights” not seen by sensors[24]).

Furthermore suppose that that the average light in the set were exactly in the uniform direction, in 31-space (this is not so difficult to believe as, if we allow intensity to vary, this intensity variation would be the direction of maximal variance). In this case we need only add the uniform vector to the illuminant basis and again find the eigenvector basis.

While we do not expect such a basis to provide the best results, simply used in eq. (11), use of such a generic basis should provide a useful test of just how far the regression step can take us, starting from a non-optimal basis. It turns out that the regression provides a powerful tool for determining spectra, even from basis vector subspaces that completely overlap.

We create a "uniform-vector" basis $\boldsymbol{U}$ by forming the projector onto the uniform vector $\boldsymbol{u} = (1, 1, 1, 1, ...)$. This projector is $\boldsymbol{u}\,\boldsymbol{u}^T$, up to normalization. This provides a 31-dimensional set of eigenvectors. We use this set for both lights and surfaces: note that such a generic basis includes the unit-vector in both light and reflectance bases. We denote this method, utilizing eq. (11) with this simplest basis, and then forming a regression with Tikhonov regularization onto light and surface coefficients, as **Method REGUNILOGSEP**.

5.B. Direct Mapping of Uniform-Basis Colour-Signal Coefficients onto Light and Surface Coefficients

Finally, since we are interested in how regression plus regularization performs in general, we can also simply form a regression from the Uniform basis coefficients of the log colour signal separately onto the separate log light and log surface coefficients, in this basis, for the training set. I.e., we form log colour signal coefficients $\widetilde{\boldsymbol{\alpha}}$, via

$$min\, \mathcal{I}_3 \;=\; \sum \|\widetilde{\boldsymbol{\alpha}} \;-\; \boldsymbol{U}\,\widetilde{\mathscr{C}}\,\|^2 \qquad (17)$$

and then carry on as in eqs. (13) and (16). We call this simplest method **Method REG**. I.e., instead of using eq. (11) we simply find the best least-squares basis coefficients for the log colour signal, and then regress onto the best log light and log surface coefficients, using separate regressions but using Tikhonov regularization to reduce the effect of noise in the samples. We note that the regularization parameter λ effectively culls the effective dimensionality; here again we use $\lambda = 10^{-3}$. Results below show that even in this non-optimal basis case, results using a regression step are quite good (but not as good as using eq. (11) with a principal component basis for lights and for surfaces, followed by regression). Results for regression using Method REG are found to be essentially the same as for using Method REGUNILOGSEP (using any basis, in fact, or separate bases for lights and surfaces).

6. Experimental Results

6.A. Method SEP Results

As a benchmark, we establish results for the original, nonlinear algorithm **SEP** first. As a training set of lights, we take 102 standard, measured illuminants.[25] these are shown in Fig. 1(a) — they represent a wide variety of natural and artificial lighting. The singular values for this set, plotted in

Fig. 1(b), indicate that a dimension of about $m = 4$ will adequately represent the illuminants. As a first reflectance set, we use the 170 reflectances of natural objects measured by Vrhel et al.[26] Colour signals are formed from products of these illuminants and reflectances (no transmission functions were included in the set of model colour signals). A plot of the surface reflectance functions is not shown as it is less enlightening. However, Fig. 1(c), plotting the singular values for this reflectance set, indicates that again a dimension $n = 4$ should be sufficient for modelling spectra *individually*, without reference to the problem at hand — separating light and surface from input colour signal.

Using the SVD vectors from these training sets of lights and surfaces, we then apply the algorithm SEP from §2 to a new set of testing colour signals. Here, we take as testing lights the set of five Judd daylights,[27] from correlated colour temperature 4,800° to 10,000°. As a first set of testing surface reflectances, we use the 24 colour patches of the MacBeth ColorChecker chart.[28]

Since we are equally interested in all wavelength components recovered, as error measure we use the simple RMS error for the spectrum approximation compared to the original: if $\tilde{E}(\lambda)$ is the approximation to the actual (non-logarithmic) $\hat{E}(\lambda)$, then

$$error \;=\; \|\tilde{E}(\lambda) \;-\; \hat{E}(\lambda)\|/\|\hat{E}(\lambda)\| \tag{18}$$

in the Euclidean norm. Before evaluating this error measure, any negative components are set to zero.

The relative timing of algorithms compared in this paper are shown in Table 1, with Method REGLOGSEP taken as unity. We see that the only substantially slower algorithm is the original algorithm, Method SEP.

For each set of light and surfaces dimensionalities, m and n, we determine the median of the above RMS error (in percent) for recovered illuminants, over all testing colour signals, for m and n ranging from 1 to 10. For Method SEP, the minimum of such median errors occurs at dimensionalities $m = 2$, $n = 9$, with error 15.01% (shown in Table 2). In contrast, the minimum of the median of the RMS surface reflectance errors, testing combinations of dimensionalities, occurs at $m = 1$, $n = 10$, at error value 16.22% (shown in Table 3). We know that for the colour signal itself, the best approximation is at maximum dimensionalities: Clearly, the more dimensionality available, the closer is the approximation to the colour signal. But when dimensionality increases, this does not make the separation algorithm perform better — in fact, it fares very poorly when dimension gets high. Instead, an optimum set of light and reflectance dimensions is observed.

Generally, we are most interested in using the illuminant error, rather than the reflectance error, as an indicator of success. Fig. 2 shows what the recovery of an illuminant with about the above median of calculated RMS errors of 15% looks like (the illuminant is daylight D65 and the surface used was Macbeth #6, “light bluish green”). We see that the approximation is reasonable, given that the only input information is the s samples of the colour signal, and the output is the $2s$ samples for separated illuminant and reflectance. Fig. 2 also shows the reflectance: the recovered curve has error 12.5% in this case. Table 2 compares the dimensionality for a minimum over dimensionalities of the median of illuminant RMS errors over all testing signals, for the methods presented here, and Table 3 shows minimum median of surface spectral reflectance function RMS errors. We note that for Method SEP, the dimensionalities that give the smallest RMS illuminant error in fact produce reflectance errors that are not well behaved in that their variability, as expressed in the standard deviation, is too high.

The bottom of each table reports results for a second, larger test: here we again use the 102 measured illuminants, but now use a standard set of 1995 reflectances[25] collected from various sources, but with the 170 Object reflectances removed, as our training set of reflectances. For testing we again use the five Judd daylights, combined with the 170 Object reflectances to make up 850 testing Colour Signals.

6.B. Method LOGSEP Results

Again using an SVD basis, but for logarithmic signals, we can apply the **LOGSEP** method of eq. (11) to the testing sets described above. Results for medians of RMS errors over all the testing signals are shown in Tables 2 and 3 (for exponentiated, not logarithmic, signals). Plots show that the median error for the LOGSEP method is much less dependent on the dimension than is the original, non-logarithmic method. We also notice, as per Table 2, that since the log-signals are more correlated than the non-logarithm ones, the separation of light and surface is not quite as good: for the first testing set the minimum median error is increased to 16.9%, at dimensions $m = 1$, $n = 4$. Nonetheless the algorithm proceeds far faster than the original method. The second, much larger, testing set behaves similarly.

Better minimum illuminant error results for the first testing set obtain if we go on to optimize the basis sets by following on with an extra regression step, as follows.

6.C. Method REGLOGSEP Results

Fig. 3 shows results for method **REGLOGSEP —** using an SVD basis in eq. (11), and then a regression onto light and surface coefficients. For the first testing set, the best illuminant separation occurs for dimensionalities $m = 4$, $n = 11$, as per Table 2, and the best reflectance recovery occurs for $m = 10$, $n = 20$, from Table 3. Except for low surface dimensionalities n, the results are fairly independent of dimensionality over a substantial range of dimensions, especially for the reflectance.

For the second, larger testing set (Objects), in fact Table 2 shows that Method SEP gives a smaller median RMS error. However, the minimum error is close to that of Method REGLOGSEP, and takes substantially longer to run. Table 3 shows that while Method REGLOGSEP is the clear best choice for the first testing set, it falls behind Method SEP for the second: however, the latter is inapplicable here because of the large variability in accuracy.

We also tried applying a bias correction to the original, nonlinear method, Method SEP, but results were not improved.

6.D. Methods REGUNILOGSEP and REG Results

Table 4 shows results of using Method **REGUNILOGSEP** and Method **REG** of § 5. Recall that these two methods utilize a Uniform basis, as defined there. The first method, REGUNILOGSEP, solves eq. (11) and then regresses onto the best training set light and surface weights. The second method, REG, simply regresses from colour signal weights onto light and surface weights. Table 4 shows that both methods perform about equally. Plots of the error surfaces for Method REGUNILOGSEP and for Method REG are very similar. Fig. 4 shows the error surfaces for Method REGUNILOGSEP.

Both of these methods use log spectra; for illuminant error they do not do as well as regressing from an SVD basis, but do quite well for reflectance error. This shows the power of the regression method itself.

Overall, results for all methods tested indicate that while the LOGSEP approach, involving an analytic solution for light and surface coefficients, is fast, it is substantially improved by an extra regression step. **Method REGLOGSEP** produces the lowest-error results in our first set of test results, and the second-lowest of usable results for our second testing set (SEP is slightly better, for illuminant error). Given the great speed advantage, we therefore suggest using Method REGLOGSEP as the method of choice.

7. Conclusions

We have presented a method for separating the incoming colour signal spectrum into an illuminant spectrum and a surface reflectance spectrum using an analytic, linear equation. The method is based on utilizing light and surface basis sets from statistical distributions of logarithms of illuminant and surface spectra. Moreover, the best results for illuminant error, using a fast, log-based closed form and unique method, derive from further carrying out an additional, regression step from the calculated basis function coefficients onto separate illumination and surface coefficients, making use of Tikhonov regularization to mitigate the effects of noise. We find that the extra regression step is a powerful method for guiding the solution, even using a generic basis. For the datasets examined here, the best results derive from using a principal component basis shaped by the extra regression. The original non-log method, SEP, actually can demonstrate good performance, but sometimes behaves poorly and with unacceptably wide variability. The main advantage of the new method is the speedup that results from replacing an iterative procedure with an analytic one.

The bias correction step itself is seen to be powerful in itself (Table 4), when applied to log signals, and is again much faster than the original method, SEP.

Further work will involve amalgamating spectral evidence from multiple pixels to estimate the illuminant, and adding additional constraints beyond the norm constraint implicit in Tikhonov regularization, such as controlling for positivity of the output spectra, for bounding the error, and for spectral smoothness. However, even the results obtained using the simple approach presented here are already satisfactory, and form a fast algorithm.

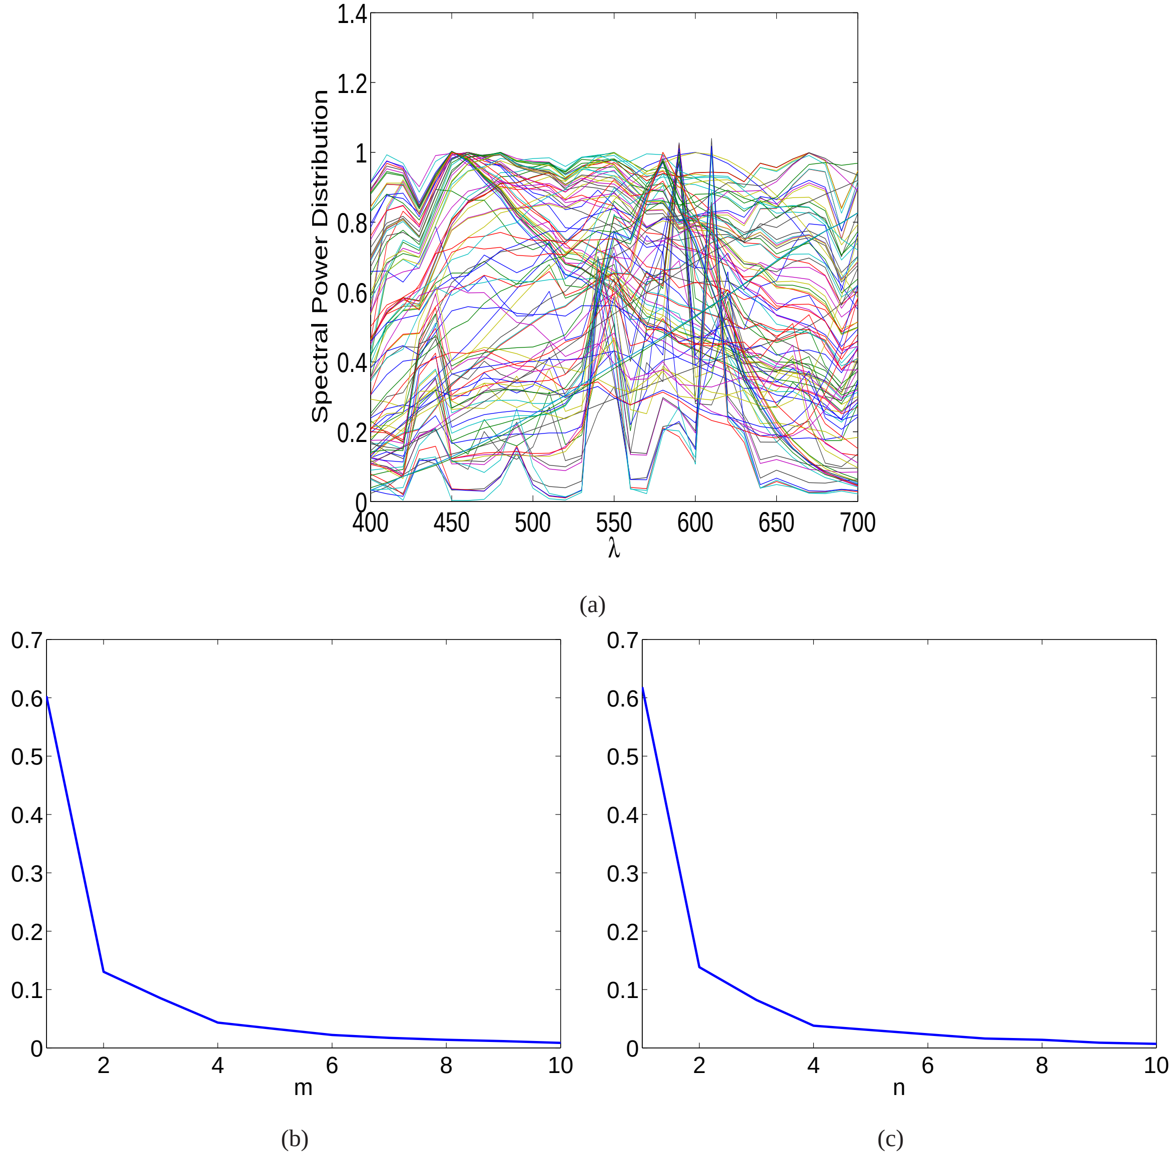

Fig. 1. Training illuminants and reflectances. (a): Set of 102 standard illuminant spectra. (b): Singular values for illuminant set. (c): Singular values for reflectance set.

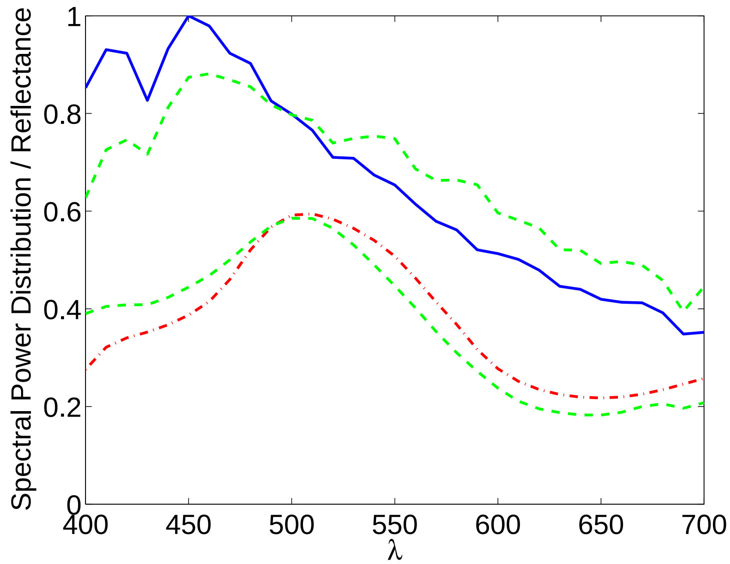

Fig. 2. Plot of an illuminant recovered with about 15% error (actual: solid line, approximation: dashed line); and reflectance recovered with about 13% error (actual: dot-dashed, approximation: dashed).

References

1. J. Ho, B. Funt, and M. Drew, "Separating a Color Signal into Illumination and Surface Reflectance Components: Theory and Applications," IEEE Trans. Patt. Anal. and Mach. Intell. **12**, 966–977 (1990). Reprinted in: Physics-Based Vision. Principles and Practice, Vol. 2, eds. G.E. Healey, S.A. Shafer, and L.B. Wolff, Jones and Bartlett, Boston, 1992, page 272.
2. G. Healey and K. Chandra, "Estimating visible through near-infrared spectral reflectance from a sensor radiance spectrum," J. Opt. Soc. Am. A **21**, 1825–1833 (2004).
3. S. Nevas, F. Manoocheri, and E. Ikonen, "Gonioreflectometer for measuring spectral diffuse reflectance," Applied Optics **9**, 6391–6399 (2004).
4. Http://www.multispectral.org/.
5. S. Bergner, T. Möller, M. K. Tory, and M. Drew, "A Practical Approach to Spectral Volume Rendering," IEEE Transactions on Visualization and Computer Graphics **11**(2), 207–216 (2005).
6. Q. Wu, L. Zeng, H. Ke, W. Xie, H. Zheng, and Y. Zhang, "Analysis of blood and bone marrow smears using multispectral imaging analysis techniques," in *Medical Imaging 2005: Image Processing*, pp. 1872–1882 (2005).
7. M. Roula, J. Diamond, A. Bouridane, P. Miller, and A. Amira, "A Multispectral Computer Vision System for Automatic Grading of Prostatic Neoplasia," in *Int. Symp. Biomedical Imaging* (2002).

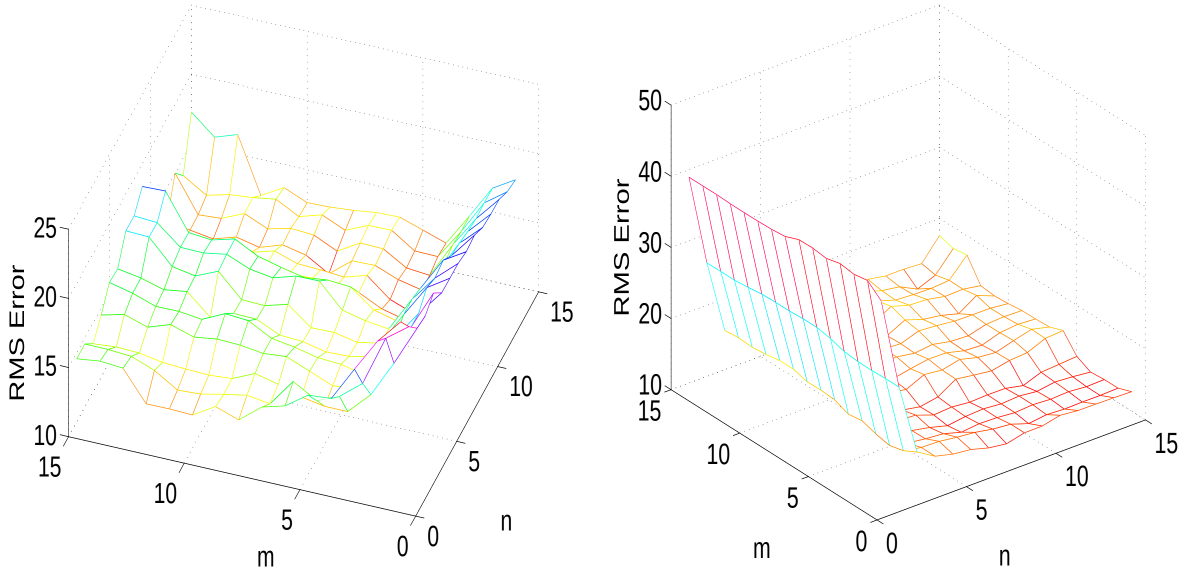

(a) (b)

Fig. 3. Median errors for **Method REGLOGSEP**. (a): Median errors for recovered illuminants. (b): Median surface reflectance errors.

8. N. Tsumura, "Appearance reproduction and multi-spectral imaging," in *AIC Colour 05: 10th Congress of the Int. Colour Assoc.*, pp. 119–123 (2005).
9. R. Berns, *Seurat and the Making of* La Grande Jatte*, ed. R.L. Herbert*, chap. "Rejuvenating Seurat's palette using color and imaging science: A simulation", pp. 214–227 (Univ. of California Press, 2004).
10. H. Barrow and J. Tenenbaum, "Recovering intrinsic scene characteristics from images," in *Computer Vision Systems*, A. Hanson and E. Riseman, eds., pp. 3–26 (Academic Press, 1978).
11. B. Wandell, "The synthesis and analysis of color images," IEEE Trans. Patt. Anal. and Mach. Intell. **PAMI-9**, 2–13 (1987).
12. D. H. Marimont and B. A. Wandell, "Linear models of surface and illuminant spectra," J. Opt. Soc. Am. A **9**, 1905–1913 (1992).
13. J. Marchant and C. Onyango, "Spectral invariance under daylight illumination changes," J. Opt. Soc. Am. A **19**, 840–848 (2002).
14. M. Drew and B.V. Funt, "Natural metamers," CVGIP:Image Understanding **56**, 139–151 (1992).
15. B. Funt, "Modelling Reflectance by Logarithmic Basis Functions," in *Color Imaging Conference: Color, Science, Systems and Applications*, pp. 68–71 (Society for Imaging Science & Technology (IS&T)/Society for Information Display (SID) joint conference, 1993).
16. E. Angelopoulou, "Objective Colour from Multispectral Imaging," in *ECCV 2000: European Conference on Computer Vision*, pp. 359–374 (2000).
17. R. Lenz, P. Meer, and M. Hauta-Kasari, "Spectral-based illumination estimation and color correction,"

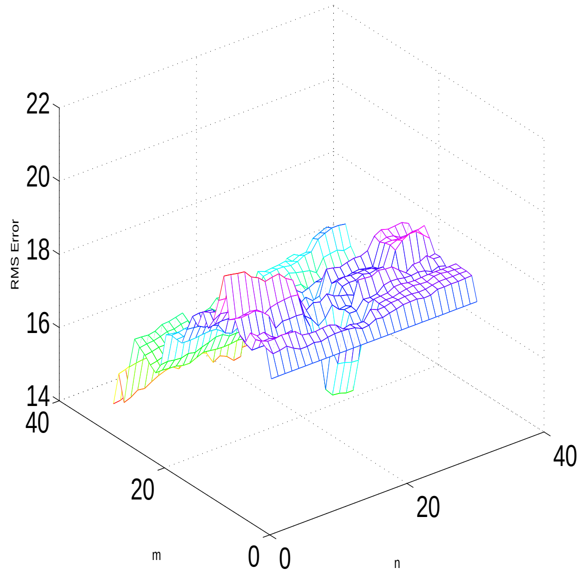

(a)

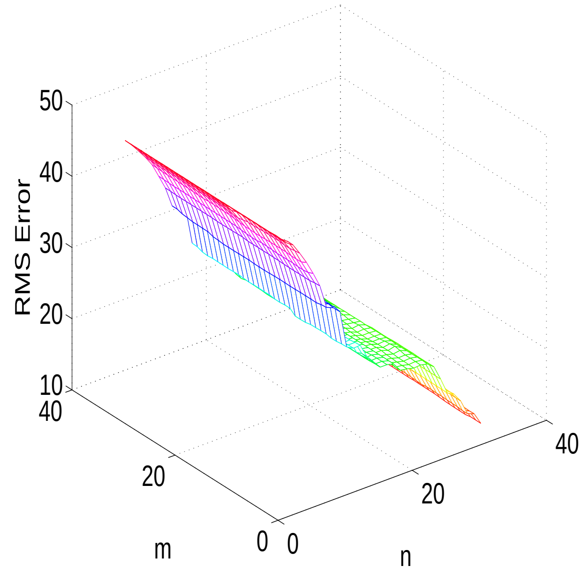

(b)

Fig. 4. Median errors for **Method REGUNILOGSEP**. (a): Median errors for recovered illuminants. (b): Median surface reflectance errors.

Color Research and Applications **24**, 98–111 (1999).

18. D. Judd, D. MacAdam, and G. Wyszecki, "Spectral Distribution of Typical Daylight as a Function of Correlated Color Temperature," J. Opt. Soc. Am. **54**, 1031–1040 (August 1964).
19. S. Nascimento, D. Foster, and K. Amano, "Psychophysical estimates of the number of spectral-reflectance basis functions needed to reproduce natural scenes," J. Opt. Soc. Am. A **22**, 1017–1022 (2005).
20. J. Faires and R. Burden, *Numerical Analysis*, 3rd ed. (Brooks/Cole, 2003).
21. Z.-N. Li and M. Drew, *Fundamentals of Multimedia* (Prentice-Hall, 2004).
22. A. Hyvärinen, J. Karhunen, and E. Oja, *Independent Component Analysis* (John Wiley and Sons, Inc., New York, 2001).
23. A. Tikhonov and V. Arsenin, *Solutions of ill-posed problems* (Wiley, 1977).
24. D. Brainard, B. Wandell, and W. Cowan, "Black light: How sensors filter spectral variation of illuminant," IEEE TRans. Biomed. Eng. **36**, 140–149 (1989).
25. K. Barnard, L. Martin, B. Funt, and A. Coath, "A Data Set for Colour Research," Color Research and Applications **27**, 147–151 (2002).
26. M. Vrhel, R. Gershon, and L. Iwan, "Measurement and analysis of object reflectance spectra," Color Research and Application **19**, 4–9 (1994).
27. G. Wyszecki and W. Stiles, *Color Science: Concepts and Methods, Quantitative Data and Formulas*, 2nd ed. (Wiley, New York, 1982).

28. C. McCamy, H. Marcus, and J. Davidson, “A color-rendition chart,” J. App. Photog. Eng. **2**, 95–99 (1976).

Method	m	n	Timing
SEP	10	10	185.9
LOGSEP	10	10	0.993
REGLOGSEP	10	10	1.000
REGUNILOGSEP	10	10	0.974
REG	31	31	1.467

Table 1. Average relative timings for algorithms considered (with unity = 0.027 seconds on a 2.8GHz PC running Matlab under Linux.)

Method	m	n	Minimum of Median **Illum.** Error	Std.Dev. Illum. Error	Median Reflec. Error	Std.Dev. Reflec. Error
Macbeths			RMS% Error			
SEP	2	9	15.01	13.87	*18.54*	*25.96*
LOGSEP	1	4	16.88	7.42	19.73	9.38
REGLOGSEP	4	11	**11.23**	7.90	12.22	8.66
Objects			RMS% Error			
SEP	2	8	**12.76**	**10.90**	15.49	12.30
LOGSEP	1	6	16.67	7.03	17.46	8.26
REGLOGSEP	4	22	13.62	8.56	14.90	9.00

Table 2. Minima, over methods examined, of median RMS errors. Two training and testing sets are examined: “Macbeths” indicates a training set consisting of 102 standard illuminants and 170 Object reflectances, and a testing set of 5 Judd daylights and 24 Macbeth colour patches; “Objects” indicates a training set with the same 102 illuminants plus 1825 measured reflectances, and a testing set of the 5 Judd daylights combined with the 170 Object reflectances. Dimensionalties m, n give minimum median Illumination error over all combinations of testing lights and surfaces; RMS Illuminant error for that m, n; standard deviation of RMS Illuminant error; RMS error and standard deviation for Reflectances at that dimensionalty. The values in boldface indicate the least-error method. The values in italics indicate an inapplicable method because of the wide variability in accuracy (high standard deviation).

Method	m	n	Median Illum. Error	Std.Dev. Illum. Error	Minimum of Median **Reflec.** Error	Std.Dev. Reflec. Error
Macbeths			RMS% Error			
SEP	1	10	16.50	7.71	16.22	9.48
LOGSEP	1	16	18.32	6.50	17.28	8.80
REGLOGSEP	10	20	12.17	7.68	**10.35**	**7.00**
Objects			RMS% Error			
SEP	2	10	13.77	12.23	*12.76*	*22.60*
LOGSEP	1	27	17.87	6.76	15.12	8.38
REGLOGSEP	4	22	13.62	8.56	**14.90**	**9.00**

Table 3. Dimensionalities for minima for Reflectance errors, and errors for Illumination and Reflectance at those dimensions. The values in boldface indicate the least-error method. The values in italics indicate an inapplicable method because of the wide variability in accuracy (high standard deviation).

Macbeths			RMS% Error			
Method	m	n	Median Illum. Error	Std.Dev. Illum. Error	Median Reflec. Error	Std.Dev. Reflec. Error
REGUNILOGSEP	31	31	14.70	9.45	12.24	9.56
REG	31	31	14.28	9.29	12.44	8.88
Objects			RMS% Error			
Method	m	n	Median Illum. Error	Std.Dev. Illum. Error	Median Reflec. Error	Std.Dev. Reflec. Error
REGUNILOGSEP	31	31	17.67	9.53	16.97	8.66
REG	31	31	16.84	9.28	16.93	8.31

Table 4. RMS errors for Uniform Basis Regression methods of §5, for testing set. Top line: using regression of solutions of eq. (11) in Uniform basis onto training set log light and log surface coefficients $\boldsymbol{\epsilon}$ and $\boldsymbol{\sigma}$. Bottom line: using regression of Uniform basis coefficients for log Colour Signal onto training set $\boldsymbol{\epsilon}$ and $\boldsymbol{\sigma}$.

Figure Captions

Fig. 1
Training illuminants and reflectances. (a): Set of 102 standard illuminant spectra. (b): Singular values for illuminant set. (c): Singular values for reflectance set.

Fig. 2
Plot of an illuminant recovered with about 15% error (actual: solid line, approximation: dashed line); and reflectance recovered with about 13% error (actual: dot-dashed, approximation: dashed).

Fig. 3
Median errors for **Method REGLOGSEP**. (a): Median errors for recovered illuminants. (b): Median surface reflectance errors.

Fig. 4
Median errors for **Method REGUNILOGSEP**. (a): Median errors for recovered illuminants. (b): Median surface reflectance errors.