Illumination and Reflectance Spectra Separation of a Hyperspectral Image Meets Low-Rank Matrix Factorization

Yinqiang Zheng
National Institute of Informatics
yqzheng@nii.ac.jp

Imari Sato
National Institute of Informatics
imarik@nii.ac.jp

Yoichi Sato
The University of Tokyo
ysato@iis.u-tokyo.ac.jp

Abstract

This paper addresses the illumination and reflectance spectra separation (IRSS) problem of a hyperspectral image captured under general spectral illumination. The huge amount of pixels in a hypersepctral image poses tremendous challenges on computational efficiency, yet in turn offers greater color variety that might be utilized to improve separation accuracy and relax the restrictive subspace illumination assumption in existing works. We show that this IRSS problem can be modeled into a low-rank matrix factorization problem, and prove that the separation is unique up to an unknown scale under the standard low-dimensionality assumption of reflectance. We also develop a scalable algorithm for this separation task that works in the presence of model error and image noise. Experiments on both synthetic data and real images have demonstrated that our separation results are sufficiently accurate, and can benefit some important applications, such as spectra relighting and illumination swapping.

1. Introduction

A hyperspectral signal records the spectral radiance of a reflective surface, which is the compound of the illumination spectral power distribution and the surface reflectance spectra. To extract the illumination and the reflectance spectra from an observed hyperspectral signal has been a longstanding problem in photometric computer vision [8,10,15]. Obviously, this separation problem is underconstrained, since there are twice as many unknowns as observations. To resolve this issue, a typical way is to introduce the low-dimensional subspace model of reflectance and that of illumination. Although it is widely believed that any practical reflectance spectrum can be well approximated by a low-dimensional model [24, 25], this does not hold for the illumination spectra in general, except the daylight spectra family [19, 28] and a limited few spectra of artificial illuminants [27].

This paper addresses a more practical variant of the classical separation problem of a single spectral signal under restricted subspace illumination, namely, the illumination and reflectance spectra separation problem of a whole hyperspectral image captured under general spectral illumination, hereafter referred to as the IRSS problem. We believe that to explore this extension is necessary and worthwhile, especially when considering that an image usually contains greater color variation, which might be utilized to improve the separation accuracy and relax the restriction of subspace illumination. Unfortunately, the huge amount of pixels in a hyperspectral image does impose great challenges on the aspects of model expression and computational efficiency.

Under the ideal diffuse reflectance and spatially uniform illumination assumption, we have found that this IRSS problem can be modeled into a low-rank matrix factorization problem. By doing so, the solution properties, like uniqueness, can be analyzed elegantly by means of singular value decomposition (SVD). Actually, we have proved that this IRSS problem assumes a unique solution up to an unknown scale between the illumination and reflectance components, under the standard assumption that reflectance spectra lie in a low-dimensional linear subspace. Considering that this subspace model is not perfectly errorless and the image intensity values usually suffer from noise, we also develop a scalable algorithm that works in the presence of both model error and image noise. Rather than explicitly describing the physical imaging process of those complicated effects beyond diffuse reflectance, like shadows and highlights, we treat them as outliers to our low-rank model, which can be accounted for via low-rank matrix approximation operation of a nonnegative observation matrix under the robust L_1-norm error metric.

Quantitative experiments on both synthetic data and real images have demonstrated that, our separation results of scenes with sufficient color variation are reasonably accurate, and can benefit some important applications, such as spectra relighting of a single view and illumination swapping between two different views.

Our major contributions can be summarized as follows:

(i). To suggest a novel low-rank matrix factorization perspective for the IRSS problem under general spectra illumination; (ii) To analyze its solvability via singular value decomposition; (iii) To develop a scalable factorization algorithm that works in real scenarios with model error and image noise.

1.1. Related Works

1.1.1 Illumination and Reflectance Spectra Separation

Existing works on illumination and reflectance spectra separation originate either from the classical problem of separating a single hyperspectral color signal or from the broad class of literatures on color constancy of a trichromatic RGB image.

As for the spectra separation task of a single signal, Ho et al. [15] introduced the subspace models of illumination and reflectance, and formulated it into a noncovex bilinear program. To better constrain this problem, Chang and Hsieh [8] enforced proper constraints on reflectance and illumination, and used simulated annealing to solve the bilinear program so as to reduce the risk of being trapped into local minima. Drew and Finlayson [10] successfully avoided the bilinearity by manipulating it in the logarithmic space. Although some heuristic countermeasures have been suggested in [10], the resulting linear system might suffer from numerical instability, since taking logarithm of an infinitesimal number would result in a huge one.

Efforts have been made to handle more hyperspectral signals. For example, Ikari et al. [17] utilized similar subspace assumptions and constraints of [8], and were able to separate dozens of hyperspectral signals. One might seek to extend their method further to handle a high-resolution hyperspectral image with at least tens of thousands of pixels. Unfortunately, the classical bilinear formulation and the typical simulated annealing algorithm [8, 15, 17] could not handle such a large-scale spectra separation problem. This reveals the necessity of developing a compact model to describe, and a scalable algorithm to solve, this large separation problem.

Computational color constancy of a trichromatic image [2, 11, 13] is to simultaneously estimate the RGB color of the illuminant and the true scene color. A number of priors on the scene, ranging from the plain grey-world [5] or white-world [20] assumption to more complicated statistical knowledge [3, 12], as well as special optical effects, like highlights [29], have been utilized for this task.

Researchers have generalized the assumptions and methods in trichromatic computational color constancy for the illumination and reflectance separation problem of a hyperspectral image [26]. For example, Huynh and Robles-Kelly [16] assumed that the scene could be segmented into several homogeneous surface patches, and were able to estimate the illumination and reflectance spectra under the dichromatic reflectance model. In principle, some other methods for trichromatic color constancy, like [5, 20, 29], could be extended to handle a hyperspectral image as well, when their underlying assumptions on the scene or the illuminant indeed apply.

Rather than relying on any restrictive spatial distribution assumption of the scene, we resort to the more general and widely accepted assumption that any reflectance spectrum lies in a low-dimensional subspace [24, 25]. In this sense, our work is more closely related to those on separating one or multiple hyperspectral signals [8, 10, 15, 17] than to those arising from the scenario of computational color constancy. Interestingly, this subspace assumption on reflectance spectra leads naturally to a compact mathematical description of the IRSS problem via low-rank matrix factorization. It also enables us to easily analyze its solvability, and to readily develop scalable algorithms for solving by virtue of the latest theoretical and algorithmic advancements in low-rank matrix factorization.

1.1.2 Low-Rank Matrix Factorization

A great variety of problems in science and engineering assume low-dimensional subspace structures, and have been solved via low-rank matrix factorization. Particularly, in photometric and geometric computer vision, the underlying low-rank structure of various entities has been exploited for photometric stereo [32], camera calibration [21, 33], rigid [30, 35] and nonrigid [4, 9, 35] structure-from-motion and so on.

The fundamental tool for low-rank matrix factorization is singular value decomposition (SVD), which offers the optimal solution when the measurement matrix is corrupted by Gaussian noises [18]. In recent years, much research attention has been paid to the more practical scenario that a measurement matrix is corrupted by gross outliers. A plenty of theoretical fruits on the quality of various convex surrogates [6, 7] have been reported, followed by tremendous algorithmic advancements in developing different kinds of fast first-order optimization algorithms (see [22, 23, 34] and many others).

This work adds another member into the broad family of low-rank problems in photometric computer vision. Unlike a typical low-rank problem, IRSS is particular in the sense that its measurement matrix is nonnegative due to the physical restrictions on illumination and reflectance spectra.

2. Low-Rank Factorization Model of IRSS

For a diffuse surface point, its radiance d_i at the i-th spectral band recorded by a hyperspectral camera is proportional to the product of the illumination l_i and the surface reflectance r_i, that is,

$$d_i = l_i r_i, 1 \le i \le m, \quad (1)$$

where the proportional scalar, accounting for such factors as gain and exposure time, has been omitted, and m denotes the number of spectra bands. It is also assumed that the spectral sensitivity function of the hyperspectral camera has been precorrected to be one at all spectral bands. To separate the observed radiance spectra, one has to assume that both the illumination and the reflectance spectra are low-dimensional [8,10,15], since otherwise there would be more variables than constraints. In the following, we try to utilize a huge amount of spectral signals in a hyperspectral image to assist the separation, without imposing any restriction on the illumination spectra. It is worthy to note that the logarithmic transformation [10] does not apply to our separation problem, since certain bands (unknown beforehand) of a general illumination spectrum might be close to zero (e.g. a banded illumination).

For a hyperspectral image with n pixels under spatially uniform illumination, the intensity value d_{ij} of the j-th pixel at the i-th spectral band reads

$$d_{ij} = l_i r_{ij}, 1 \leq i \leq m, 1 \leq j \leq n, \tag{2}$$

which can be stacked into a matrix system

$$\underbrace{\begin{bmatrix} d_{11} & \cdots & d_{1n} \\ \cdots & \cdots & \cdots \\ d_{m1} & \cdots & d_{mn} \end{bmatrix}}_{D_{m\times n}} = \underbrace{\begin{bmatrix} l_1 & & \\ & \cdots & \\ & & l_m \end{bmatrix}}_{L_{m\times m}} \underbrace{\begin{bmatrix} r_{11} & \cdots & r_{1n} \\ \cdots & \cdots & \cdots \\ r_{m1} & \cdots & r_{mn} \end{bmatrix}}_{R_{m\times n}}. \tag{3}$$

The above system has mn constraints in the observation matrix D and $m(n+1)$ variables in the diagonal illumination matrix L and the reflectance matrix R, which means that the IRSS problem without any assumption on illumination nor reflectance is underconstrained, no matter how many signals we have.

It has been widely agreed that reflectance spectra lie in a low-dimensional linear subspace. Therefore, we introduce the subspace model of reflectance R, and rewrite eq.(3) into

$$D_{m\times n} = L_{m\times m} R_{m\times n} = L_{m\times m} B_{m\times s} C_{s\times n}, \tag{4}$$

in which B and C denote the spectral bases and coefficients, respectively, and s the subspace dimensionality. According to [24,25], s is often chosen to be around 8 so as to reach the best tradeoff between expression power and noise resistance in the process of fitting reflectance spectra.

Interestingly, we have found that the low-rank formulation of IRSS in eq.(4) is very similar to that of the nonrigid structure-from-motion problem (NRSfM) [1,4,9,14] in geometric vision, thus can be regarded as a spectra-domain counterpart to the NRSfM problem in time-domain. Due to page limitation, the detailed comparison of these two problems will be presented in the supplementary materials. This analogy might interest researchers in photometric and geometric computer vision in drawing on experience of solving one problem to better handle the other.

3. Ambiguity Analysis

In addition to the physically meaningful factorization in eq.(4), we can also factorize D mathematically via the partial singular value decomposition as follows

$$D_{m\times n} = U_{m\times s} S_{s\times s} V^T_{n\times s} = UQQ^{-1} S V^T = (UQ)\left(Q^{-1} S V^T\right), \tag{5}$$

in which Q is an arbitrary $s \times s$ invertible matrix. By comparing eq.(4) with eq.(5), we can recognize that our IRSS problem is to find a proper matrix Q such that

$$UQ = LB. \tag{6}$$

When the spectral bases matrix B is unknown, eq.(6) is trivial[1]. In other words, to simultaneously separate the hyperspectral image and learn the bases is infeasible.

In the following, we assume that B has been learned via a principle component analysis (PCA) of a spectra dataset [24, 25]. Considering that $U^T U = I_s$, in which I_s is the $s \times s$ identity matrix, we can multiply U^T at both sides of eq.(6), and obtain $Q = U^T LB$, which can be plugged back into eq.(6), such that

$$UQ = UU^T LB = \left(UU^T\right) LB = LB. \tag{7}$$

From eq.(4), the bounds on s should be $1 \leq s \leq \min\{m, n\}$. Considering that n is much greater than m in general, we can assume that $1 \leq s \leq m$. When $s = m$, U becomes a square matrix, and we have $U^T U = UU^T = I_s$. Under this condition, Eq.(7) reduces to a trivial equation $LB = LB$, which can be trivially satisfied by any nonnegative diagonal L. When $s \leq m - 1$, eq.(7) implies that LB is the s eigenvectors of UU^T corresponding to the eigenvalue 1, which results in a linear system with ms constraints and m variables in L. This means that, when $1 \leq s \leq m - 1$, eq.(7) has a unique solution in general, up to an arbitrary scale. Given L, the reflectance R can be calculated as

$$R = BC = BQ^{-1} S V^T = B\left(U^T LB\right)^{-1} S V^T. \tag{8}$$

Therefore, the scale ambiguity between L and R could not be resolved.

Now we reach the following proposition on the solvability of our IRSS problem.

Proposition 1. *Under the standard assumption that scene reflectance lies in a s-D low-dimensional subspace spanned by known bases, the IRSS problem assumes a unique solution, except that the absolute magnitudes of illumination and reflectance are ambiguous. The illumination spectra can be very general, as long as there are at least s nonzero bands.*

[1]For example, when letting $L = I_s$, B can be chosen to be UQ for an arbitrary invertible Q.

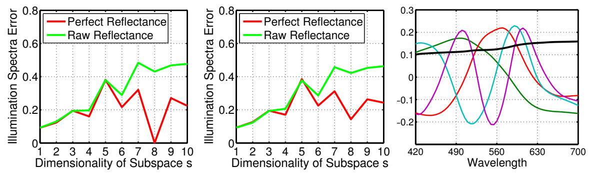

(a) 0% Noise (b) 5% Noise (c) PCA Bases

Figure 1. Illumination spectra estimation error of the linear solution using perfect and raw reflectance, with respect to varying dimensionality of reflectance subspace. A perfect reflectance spectrum is obtained by projecting a raw reflectance spectrum onto the 8-D subspace.

3.1. Verification Using Synthetic Data

To verify our theoretical analysis, we use the raw reflectance spectra of ten color patches in the Macbeth color checker as the test target. The reflectance bases are trained via PCA on the spectra data of the Gretag Macbeth Color Checker Digital Camera (CCDC) target. We have also tried to use other spectra data to learn the bases, and verified that similar results can be obtained. Considering that the raw reflectance spectra do not lie perfectly in the subspace, we project the raw spectra onto the low-dimensional subspace to simulate 'perfect reflectance spectra' for the purpose of numerical verification. According to [24, 25], the dimensionality of reflectance subspace is chosen to be 8 in the projection process. We generate random illumination spectra and normalize the scale such that the maximum value is 1. We measure the root-mean-square-error (RMSE) of the estimated illumination spectra from the linear solution in eq.(7), with respect to the corresponding ground truth. To synthesize noisy signals, we add zero-mean Gaussian noise with derivation of 5% relative magnitude. We vary the dimensionality s from 1 to 10.

As shown in Fig.1(a), the estimation error is zero, when $s = 8$ for perfect reflectance. This verifies our assertion in *proposition 1*. However, the solution from eq.(7) is quite inaccurate, when raw reflectance spectra that slightly deviate from the subspace are used. This indicates that the linear solution in eq.(7) is very sensitive to the subspace model error, even in no presence of image noise.

Surely, using the first principle component to express reflectance should not be accurate. However, an interesting observation from Fig.1(a) and (b) is that, when $s = 1$, the estimation accuracy of illumination is the highest in the case of raw reflectance. This phenomenon links to the well-known grey-world assumption [5]. Specifically, when $s = 1$, eq.(6) reduces to $\mathbf{u}_1 q_{11} = L\mathbf{b}_1$. As shown in Fig.1(c), the first principle component (in thick dark) of reflectance $\mathbf{b}_1$ is almost flat, the illumination spectra vector $\mathbf{l}, \mathbf{l} = \begin{bmatrix} l_1, & \cdots, & l_m \end{bmatrix}^T$, is thus almost proportional to $\mathbf{u}_1$, the primary left singular vector of D. In addition, for a nonnegative matrix, like our observation matrix D, its primary left singular vector is close to its row-wise average. Therefore, the illumination spectrum l_i at the i-th band almost equals the mean pixel value of the i-th band image, which is exactly the grey-world assumption. Since there are ten different colors in this simulation, the grey-world assumption holds true to a large extent.

In the following, we propose a scalable algorithm for the separation problem that works in the presence of model error and image noise.

4. Scalable Algorithms

4.1. Handling Model Error and Image Noise

As recognized in [8, 17], the scene reflectance is bounded between 0 and 1, that is $0 \leq r_{ij} \leq 1, i = 1, 2, \cdots, m, j = 1, 2, \cdots, n$. According to eq.(3), the constraints on the illumination spectra are

$$l_i \geq \max_{1 \leq j \leq n} \{d_{ij}\}, 1 \leq i \leq m. \tag{9}$$

Let us observe that the constraints in eq.(9) link to the well-known white-world assumption [20], which assumes that the scene contains a white surface and the illumination spectrum is equivalent to the radiance of that surface. This is exactly the special case of eq.(9) when equality holds. It is widely known that this assumption is sensitive to image outliers, caused for example by sensor defects or highlights. We are going to relieve this problem by using the robust estimation technique described later in this section.

Based on eq.(6), the illumination spectra can be estimated by solving the following minimization problem

$$\min_{\mathbf{l}, Q} \|UQ - LB\|_F^2, s.t., \; l_i \geq \max_j \{d_{ij}\}, l_z \leq \max_{i,j} \{d_{ij}\}, \tag{10}$$

in which we introduce another constraint $l_z \leq \max_{i,j}\{d_{ij}\}, 1 \leq i \leq m, 1 \leq j \leq n$, to restrict the absolute scale of $\mathbf{l}$, and z denotes the row index of the maximum value $\max_{i,j}\{d_{ij}\}$. The optimization problem in eq.(10) is a small convex quadratic program (QP), thus can be easily solved.

After obtaining $\mathbf{l}$, we can further improve the separation accuracy by minimizing the following cost function on the basis of the original low-rank model in eq.(4)

$$\begin{aligned} &\min_{\mathbf{l}, C} \|D - LBC\|_F^2, \\ &s.t., 0 \leq (BC)_{ij} \leq 1, l_i \geq \max_j \{d_{ij}\}, l_z \leq \max_{i,j} \{d_{ij}\}. \end{aligned} \tag{11}$$

Considering that C is extremely large in size, a typical descending method would be impractical. We thus develop a simple but scalable alternating projection algorithm to minimize eq.(11). Specifically, given $\mathbf{l}$, C can be updated as $C = (LB)^+ D$, in which $(LB)^+$ denotes the Moore-Penrose pseudoinverse of LB. Then, the bound constraints on reflectance BC are enforced by clamping the negative elements of BC

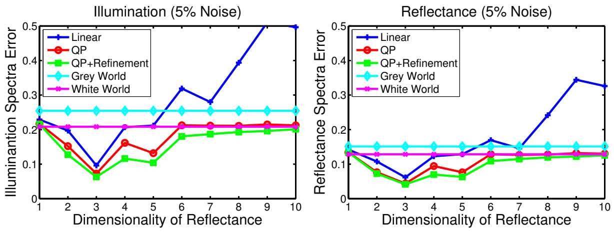

(a) Random Illumination Spectra

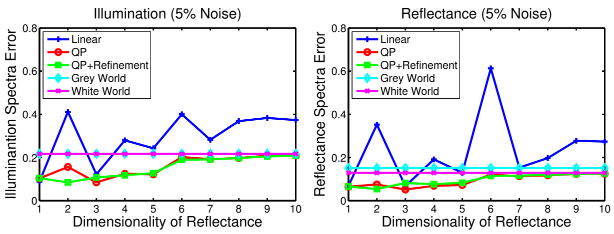

(b) Daylight Illumination Spectra

Figure 2. Separation accuracy w.r.t. the dimensionality of reflectance under random illumination (a) or daylight illumination (b), in the presence of 5% relative image noise.

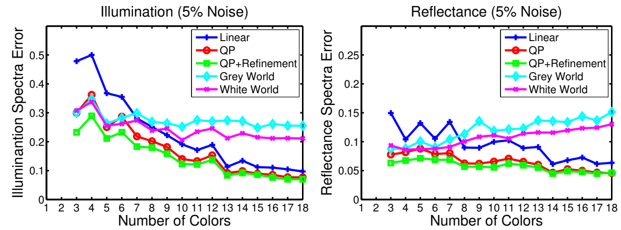

(a) Using Color Patches Only

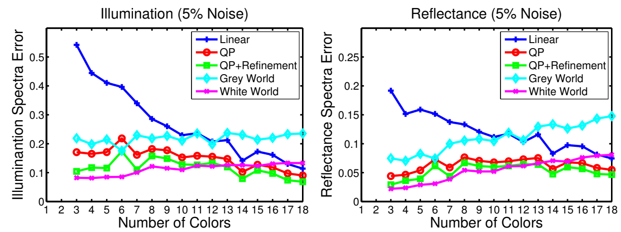

(b) Adding One Grey Patch

Figure 3. Separation accuracy w.r.t. the number of colors under general illumination in the presence of 5% relative image noise. (a) shows the results by using color patches only, while (b) the improved results by randomly adding a grey patch to assist the separation.

to be 0 and those elements greater than one to be 1. Given BC, to update $\mathbf{l}$ is simpler, since $l_i, i = 1, 2, \cdots, m$, are independent. By letting $R = BC$, $l_i = (\hat{\mathbf{d}}_i^T \hat{\mathbf{r}}_i)/(\hat{\mathbf{r}}_i^T \hat{\mathbf{r}}_i)$, in which $\hat{\mathbf{d}}_i$ and $\hat{\mathbf{r}}_i$ denote the transpose of the i-th row of D and R, respectively. The bound constraints on l_i in eq.(11) can be enforced posteriorly by clamping those elements beyond the bounds to be the bounded value. Starting from the initial $\mathbf{l}$ by solving eq.(10), we update C and $\mathbf{l}$ alternatively, until the maximum iteration number (20 in all our experiments) is reached.

4.2. Handling Outliers

Let us recall that our low-rank factorization model in eq.(4) relies on the assumption of diffuse surface reflectance and spatially uniform illumination. Here, we slightly relax this assumption by treating those effects beyond diffuse reflectance as outliers to our low-rank model, and accounting for them in a robust factorization framework.

We introduce a mask matrix W with the same size of the observation D. Its element w_{ij} is 0, when the j-th pixel at the i-th band is apparently saturated, i.e., the maximum allowed value (255 for 8-bit images) is reached. Otherwise, w_{ij} is 1. Given an observation matrix D with outliers, we try to find a nonnegative matrix $\tilde{D}$, whose rank is at most $\tilde{s}$, by minimizing the following robust L_1-norm error

$$\min_{\tilde{D}} \left\|W \odot (D - \tilde{D})\right\|_1^1, s.t., \tilde{D} \geq 0, rank(\tilde{D}) \leq \tilde{s}, \tag{12}$$

in which $\odot$ denotes element-wise multiplication of matrices.

To eliminate the rank constraint, we adopt the bilinear expression such that $\tilde{D} = \tilde{U}\tilde{V}^T$, in which $\tilde{U}^T\tilde{U} = I_{\tilde{s}}$. Eq.(12) can be reformulated as

$$\min_{\tilde{D},\tilde{U},\tilde{V}} \left\|W \odot (D - \tilde{D})\right\|_1^1, s.t., \tilde{D} \geq 0, \tilde{D} = \tilde{U}\tilde{V}^T, \tilde{U}^T\tilde{U} = I_{\tilde{s}}. \tag{13}$$

Challenging as it is, the large-scale nonconvex optimization in eq.(13) can be directly solved by using the first-order augmented Lagrange multiplier method of [34]. The rank upper bound $\tilde{s}$ is chosen to be 8 in all the experiments.

4.3. Computational Procedures

Now, we briefly summarize the computational procedures for our IRSS problem as follows: (*Step-1*). Construct the observation matrix D and remove potential outliers by solving eq.(13); (*Step-2*). Compute the partial SVD factorization of $\tilde{D}$; (*Step-3*). Solve the convex QP in eq.(10) to find the illumination spectra $\mathbf{l}$; (*Step-4*). Refine the illumination spectra via alternating minimization of eq.(11).

Actually, by solving eq.(11), we can obtain both the illumination spectra $\mathbf{l}$ and reflectance spectra R. However, as shall be shown in the following experiment section, we are going to compare our algorithm with the grey-world and white-world assumption, both of which directly estimate the illumination spectra from image observations. Given the estimated illumination spectra $\mathbf{l}$, we calculate the reflectance spectra for these two assumptions as

$$R = \tilde{B}(L\tilde{B})^{+}\tilde{D}, \tag{14}$$

in which $\tilde{B}$ is the reflectance bases with dimensionality $\tilde{s}$. To make the comparison fair, for all our experiment results in the following, we re-calculate the reflectance spectra by using eq.(14) as well, instead of using those from eq.(11).

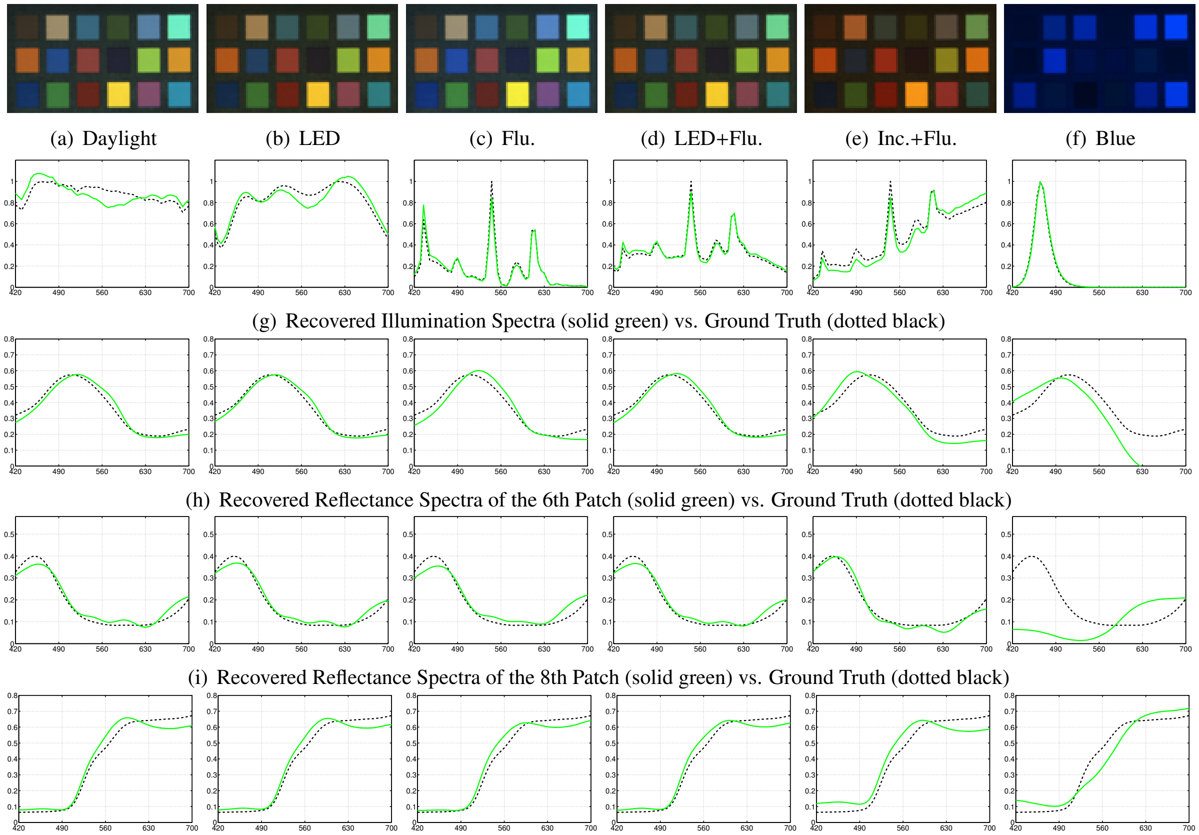

(a) Daylight (b) LED (c) Flu. (d) LED+Flu. (e) Inc.+Flu. (f) Blue

(g) Recovered Illumination Spectra (solid green) vs. Ground Truth (dotted black)

(h) Recovered Reflectance Spectra of the 6th Patch (solid green) vs. Ground Truth (dotted black)

(i) Recovered Reflectance Spectra of the 8th Patch (solid green) vs. Ground Truth (dotted black)

(j) Recovered Reflectance Spectra of the 12th Patch (solid green) vs. Ground Truth (dotted black)

Figure 4. Illumination and reflectance separation of the Macbeth color checker (18 color patches) under daylight (a), LED lamp (b), office fluorescent light (c), LED lamp with fluorescent light (d), incandescent lamp with fluorescent light (e) and blue light (f). The recovered illumination spectra (g) and representative reflectance spectra (h,i,j) are compared with their respective ground truth.

5. Experiment Results

5.1. Synthetic Data

We use the reflectance spectra of all 18 color patches in the Macbeth color checker to simulate data. Random illumination spectra or random daylight spectra are generated for experiments. To resolve the scale ambiguity of illumination and reflectance, we normalize the illumination spectra such that the maximum value is 1. We measure the RMSE of estimated illumination and reflectance spectra with respect to their respective ground truth.

5.1.1 Accuracy w.r.t. Dimensionality of Reflectance

Here, we try to investigate the effect of reflectance dimensionality on separation accuracy. We vary the dimensionality s from 1 to 10, and measure the illumination and reflectance spectra error for each s. We add zero-mean Gaussian noise with 5% relative magnitude. We show the average error over 100 independent trials in Fig.2. To clearly reveal how estimation accuracy is improved hierarchically, we present the results of the linear solution (Linear) in eq.(7), the quadratic program solution (QP) in eq.(10) and the refined solution (QP+Refinement) in eq.(11). We also include the grey-world and the white-world assumption as baseline.

From Fig.2, our preliminary observation is that the most appropriate dimensionality s is 3. This is a little bit surprising to us, since we were expecting that the best s should be around 8, the dimensionality usually adopted to represent reflectance. One possible reason for this observation might be that the potential sway of error between illumination and reflectance will advocate a tighter subspace expression of reflectance spectra. Due to this observation, the subspace dimensionality s is chosen to be 3 in all the remaining experiments.

5.1.2 Accuracy w.r.t. Color Variation

In addition to the subspace dimensionality, color variation in a scene might also affect the separation accuracy. To investigate this factor, we randomly pick k color patches, varying from 3 to 18. The average illumination and reflectance spectra error over 100 runs are shown in Fig.3(a), from which we can observe that color variation indeed has significant impact on the estimation accuracy. The accuracy

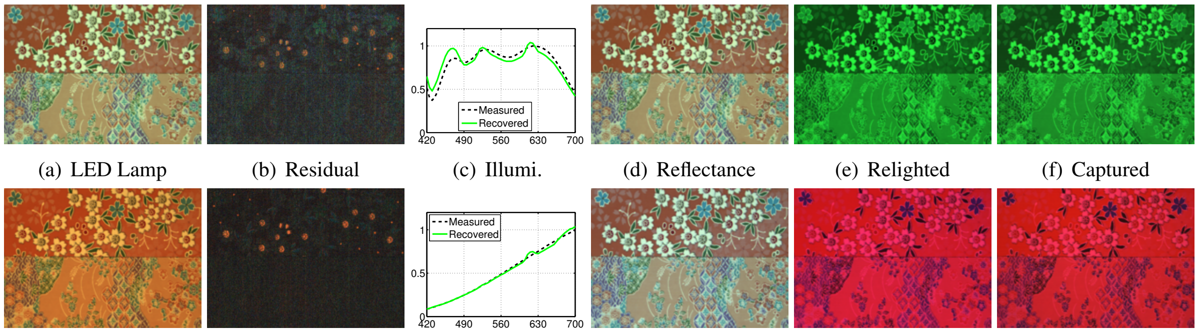

(a) LED Lamp (b) Residual (c) Illumi. (d) Reflectance (e) Relighted (f) Captured

(g) Inc. Lamp (h) Residual (i) Illumi. (j) Reflectance (k) Relighted (l) Captured

Figure 5. Spectra separation and relighting of the textured paper scene. For each row, the subfigures from left to right are the original scene, residual image, recovered illumination spectra, reflectance image, relighted scene under another illumination and the captured scene under that illumination, respectively.

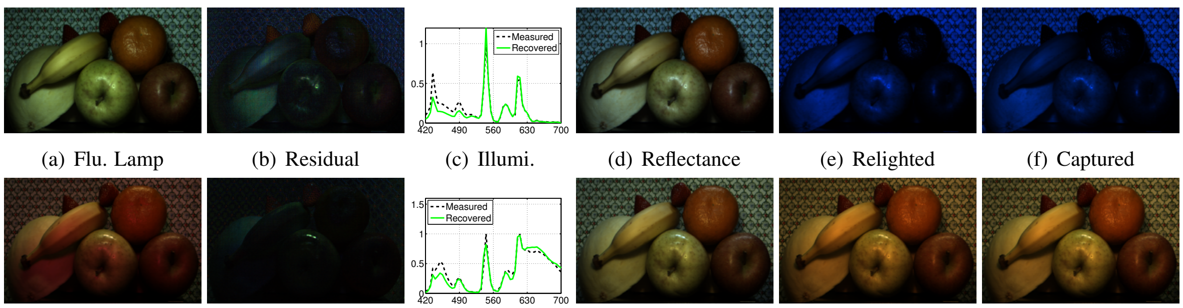

(a) Flu. Lamp (b) Residual (c) Illumi. (d) Reflectance (e) Relighted (f) Captured

(g) Pink+Flu. (h) Residual (i) Illumi. (j) Reflectance (k) Relighted (l) Captured

Figure 6. Spectra separation and relighting of the fruit scene. The subfigures are arranged in the same way as of Fig.5.

improves monotonically as the number of colors increases. Roughly, to ensure an accurate separation, the scene should contain at least 10 colors. In some real scenarios, the number of colors might be less than 10. As a remedy, it is not too demanding to assume that the scene usually covers some nearly white or grey surfaces. This is similar to the white-world assumption in a certain sense. We repeat the above experiment, and randomly pick a grey patch from the Macbeth color checker as well. The results are shown in Fig.3(b), from which we see that the separation accuracy is reasonably high, even in the presence of three or four colors.

5.2. Real Images

To capture hyperspectral images of real scenes, we use a NH-7 camera by EBA JAPAN Co. LTD, working in the visible range from 420nm to 700nm with an interval of 5nm. All hyperspectral images are shown in RGB for visualization.

5.2.1 Separation Accuracy

We first investigate the separation accuracy of real images under varying spectra illuminations. For convenience, the 18 color patches of the Macbeth color checker are used as test target, whose reflectance spectra are standard. As shown in Fig.4, we take images of the target under either natural or artificial lights, and compare the recovered illumination and reflectance spectra with their respective ground truth. From Fig.4, we can observe that our separation results are sufficiently accurate, no matter when the illumination spectra are relatively smooth (e.g. the household LED spectra in Fig.4(b)) or extremely spiky (e.g. the office fluorescent lamp spectra in Fig.4(c-e)).

An important observation from the last column of Fig.4 is that our method works very well to recover banded illumination spectra. However, for those intervals without irradiance, the recovered reflectance spectra rely primarily on the subspace model. Therefore, the accuracy of recovered reflectance might be poor.

5.2.2 Spectra Relighting

By recovering the reflectance component, it is possible to relight the scene under any novel illumination. We have obtained some preliminary results by using the textured paper scene in Fig.5 and the fruit scene in Fig.6. The robust non-negative low-rank matrix approximation scheme in Sec.4.2 has been used first to remove potential highlights. As shown in the residual images (2nd column) of Fig.5 and Fig.6, this robust scheme seems to work reasonably well and successfully identify some highlighted regions that conform to our

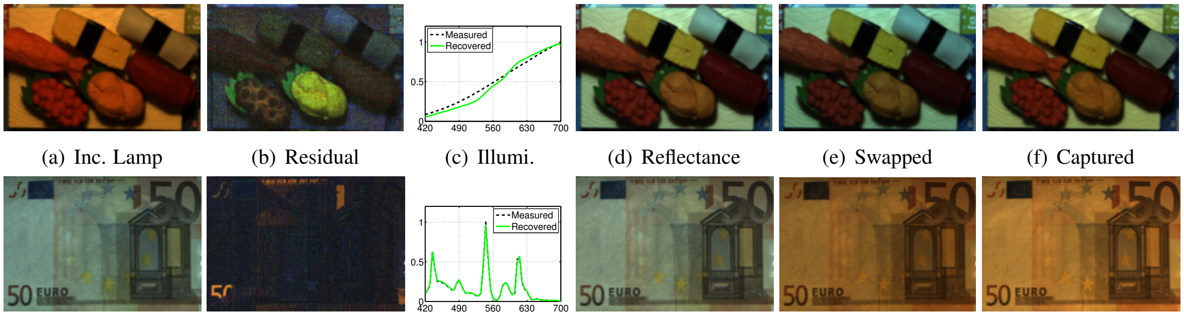

(a) Inc. Lamp (b) Residual (c) Illumi. (d) Reflectance (e) Swapped (f) Captured

(g) Flu. Lamp (h) Residual (i) Illumi. (j) Reflectance (k) Swapped (l) Captured

Figure 7. Spectra separation and illumination swapping of the sushi scene (1st row) and the currency note scene (2nd row). For each row, the subfigures from left to right are the original scene, residual image, recovered illumination spectra, reflectance image, scene under swapped illumination of the other row, and the captured scene under that illumination, respectively.

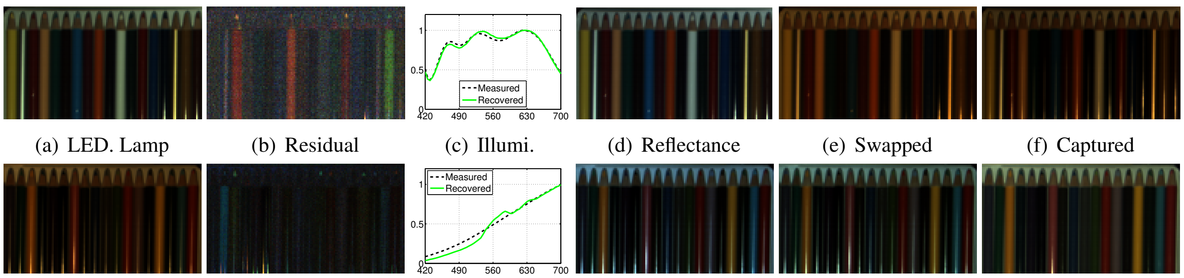

(a) LED. Lamp (b) Residual (c) Illumi. (d) Reflectance (e) Swapped (f) Captured

(g) Inc. Lamp (h) Residual (i) Illumi. (j) Reflectance (k) Swapped (l) Captured

Figure 8. Spectra separation and illumination swapping of two pencil sets. The subfigures are arranged in the same way as of Fig.7.

visual perception, although it rarely accounts for their underlying imaging mechanisms. As for the relighting results, they are pretty satisfactory even for the fruit scene in Fig.6, whose scene geometry is very complex.

5.2.3 Illumination Swapping

Given two hyperspectral images of different scenes, we can easily swap their illumination conditions on the basis of our illumination and reflectance separation results. As illustration, we have tried to swap the sushi and the currency note scene in Fig.7 and the two pencil scenes in Fig.8. Although the synthesized scene is not in complete agreement with the ground truth, the appearance discrepancy has been drastically reduced, which demonstrates the accuracy and applicability of our illumination and reflectance separation method.

6. Conclusions

We have addressed the generalized illumination and reflectance spectra separation problem of a hyperspectral image captured under arbitrary spectra illumination. We have shown that this separation problem can be handled by using low-rank matrix factorization, which is much more compact and convenient than the existing bilinear formulation. We have also proved that the separation is unique up to an unknown scale under the standard low-dimensionality assumption of reflectance. Scalable algorithms that work in the presence of model error, image noise and outliers, have been developed and shown to be successful. Experiments on both synthetic data and real images have demonstrated that, when the scene covers sufficient color variation, our separation results are sufficiently accurate. We have demonstrated two applications of spectra relighting and illumination swapping.

Our work has left out quite a few important aspects that deserve to be explored in depth. For example, highlights, being treated as outliers to the low-rank model in this paper, are known to encode some important information of the illumination. Therefore, it would be rewarding to carefully model and exploit highlights in a physically sound way, as in [17, 31]. As for the application aspect, we hope that our separation results can benefit some other potential applications, such as spectra-based material recognition and hyperspectral image compression.

Acknowledgement. This research was supported in part by the Ministry of Education, Science, Sports and Culture Grant-in-Aid for Scientific Research on Innovative Areas.

References

[1] I. Akhter, Y. Sheikh, S. Khan, and T. Kanade. Trajectory space: A dual representation for nonrigid structure from motion. *IEEE TPAMI*, 33(7):1442–1456, 2011. 3

[2] K. Barnard, V. Cardei, and B. Funt. A comparison of computational color constancy algorithms. I: Methodology and experiments with synthesized data. *IEEE TIP*, 11(9):972–984, 2002. 2

[3] D. Brainard and W. T. Freeman. Bayesian color constancy. *J. Opt. Soc. Am. A*, 14(7):1393–1411, 1997. 2

[4] C. Bregler, A. Hertzmann, and H. Biermann. Recovering non-rigid 3D shape from image streams. *In Proc. CVPR*, pages 690–696, 2000. 2, 3

[5] G. Buchsbaum. A spatial processor model for object colour perception. *J. of the Franklin Institute*, 310(1):337–350, 1980. 2, 4

[6] E. Candès, X. Li, Y. Ma, and J. Wright. Robust principal component analysis. *J. ACM*, 58(3):1–37, 2011. 2

[7] E. Candès and B. Recht. Exact matrix completion via convex optimization. *Foundations of Computational Mathematics*, 9(6):717–772, 2009. 2

[8] P. Chang and T. Hsieh. Constrained nonlinear optimization approaches to color-signal separation. *IEEE TIP*, 4(1):81–94, 1995. 1, 2, 3, 4

[9] Y. Dai, H. Li, and M. He. A simple prior-free method for non-rigid structure-from-motion factorization. *In Proc. CVPR*, pages 2018–2025, 2012. 2, 3

[10] M. Drew and G. Finlayson. Analytic solution for separating spectra into illumination and surface reflectance components. *J. Opt. Soc. Am. A*, 24(2):294–303, 2007. 1, 2, 3

[11] M. Ebner. *Color constancy*. John Wiley & Sons, 2007. 2

[12] A. Gijsenij and T. Gevers. Color constancy using natural image statistics and scene semantics. *IEEE TPAMI*, 33(4):687–698, 2011. 2

[13] A. Gijsenij, T. Gevers, and J. Van De Weijer. Computational color constancy: Survey and experiments. *IEEE TIP*, 20(9):2475–2489, 2011. 2

[14] P. Gotardo and A. Martinez. Computing smooth time-trajectories for camera and deformable shape in structure from motion with occlusion. *IEEE TPAMI*, 33(10):2051–2065, 2011. 3

[15] J. Ho, B. Funt, and M. Drew. Separating a color signal into illumination and surface reflectance components: theory and applications. *IEEE TPAMI*, 12(10):966–977, 1990. 1, 2, 3

[16] C. Huynh and A. Robles-Kelly. A solution of the dichromatic model for multispectral photometric invariance. *International Journal of Computer Vision*, 90(1):1–27, 2010. 2

[17] A. Ikari, R. Kawakami, R. Tan, and K. Ikeuchi. Separating illumination and surface spectra from multiple color signals. *Digitally Archiving Cultural Objects*, pages 297–321, 2008. 2, 4, 8

[18] I. Jolliffe. *Principal Component Analysis*. Wiley Online Library, 2005. 2

[19] D. Judd, D. Macadam, and G. Wyszecki. Spectral distribution of typical daylight as a function of correlated color temperature. *J. Opt. Soc. Am.*, 54(8):1031–1040, 1964. 1

[20] E. Land and J. McCann. Lightness and retinex theory. *J. Opt. Soc. Am*, 61(1):1–11, 1971. 2, 4

[21] J. Lee, Y. Matsushita, B. Shi, I. Kweon, and K. Ikeuchi. Radiometric calibration by rank minimization. *IEEE TPAMI*, 35(1):144–156, 2013. 2

[22] Z. Lin, M. Chen, and Y. Ma. The augmented lagrange multiplier method for exact recovery of corrupted low-rank matrices. *arXiv preprint arXiv:1009.5055*, 2010. 2

[23] Z. Lin, R. Liu, and Z. Su. Linearized alternating direction method with adaptive penalty for low-rank representation. *In Proc. NIPS*, pages 612–620, 2011. 2

[24] L. Maloney. Evaluation of linear models of surface spectral reflectance with small numbers of parameters. *J. Opt. Soc. Am. A*, 3(10):1673–1683, 1986. 1, 2, 3, 4

[25] J. Parkkinen, J. Hallikainen, and T. Jaaskelainen. Characteristic spectra of Munsell colors. *J. Opt. Soc. Am. A*, 6(2):318–322, 1989. 1, 2, 3, 4

[26] A. Robles-Kelly and C. Huynh. *Imaging Spectroscopy for Scene Analysis*. Springer, 2012. 2

[27] J. Romero, A. Garcia-Beltran, and J. Hernandez-Andres. Linear bases for representation of natural and artificial illuminants. *J. Opt. Soc. Am. A*, 14(5):1007–1014, 1997. 1

[28] D. Slater and G. Healey. What is the spectral dimensionality of illumination functions in outdoor scenes? *In Proc. CVPR*, pages 105–110, 1998. 1

[29] R. Tan, K. Nishino, and K. Ikeuchi. Illumination chromaticity estimation using inverse-intensity chromaticity space. *In Proc. CVPR*, pages 673–680, 2003. 2

[30] C. Tomasi and T. Kanade. Shape and motion from image streams under orthography: a factorization method. *IJCV*, 9(2):137–154, 1992. 2

[31] S. Tominaga. Surface identification using the dichromatic reflection model. *IEEE TPAMI*, 13(7):658–670, 1991. 8

[32] L. Wu, A. Ganesh, B. Shi, Y. Matsushita, Y. Wang, and Y. Ma. Robust photometric stereo via low-rank matrix completion and recovery. *In Proc. ACCV*, pages 703–717, 2010. 2

[33] Z. Zhang, Y. Matsushita, and Y. Ma. Camera calibration with lens distortion from low-rank textures. *In Proc. CVPR*, pages 2321–2328, 2011. 2

[34] Y. Zheng, G. Liu, S. Sugimoto, S. Yan, and M. Okutomi. Practical low-rank matrix approximation under robust L1-norm. *In Proc. CVPR*, pages 1410–1417, 2012. 2, 5

[35] Y. Zheng, S. Sugimoto, S. Yan, and M. Okutomi. Generalizing Wiberg algorithm for rigid and nonrigid factorizations with missing components and metric constraints. *In Proc. CVPR*, pages 2010–2017, 2012. 2