

Optimal Solution of the Dichromatic Model for Multispectral Photometric Invariance

Cong Phuoc Huynh¹ and Antonio Robles-Kelly^{1,2}

¹ RSISE, Bldg. 115, Australian National University, Canberra ACT 0200, Australia

² National ICT Australia (NICTA)*, Locked Bag 8001, Canberra ACT 2601, Australia

Abstract. In this paper, we address the problem of photometric invariance in multispectral imaging making use of an optimisation approach based upon the dichromatic model. In this manner, we can cast the problem of recovering the spectra of the illuminant, the surface reflectance and the shading and specular factors in a structural optimisation setting. Making use of the additional information provided by multispectral imaging and the structure of image patches, we recover the dichromatic parameters of the scene. Thus, we formulate a target cost function combining the dichromatic error and the smoothness priors for surfaces with uniform reflectance. The dichromatic parameters are recovered through minimising this cost function in a coordinate descent manner. To this end, we make use of regular albedo patches in the scene selected automatically. The algorithm is shown to be robust, and since the cost function is convex in nature, the resulting solution is globally optimal. We illustrate the effectiveness of our method for purposes of illuminant spectrum recovery and skin recognition. We also compare our results to a number of alternatives.

1 Introduction

In computer vision, the modeling of surface reflectance is a topic of pivotal importance for surface analysis and image understanding tasks. For instance, Nayar and Bolle [1] have used photometric invariants to recognise objects with different reflectance properties. This work builds on the one reported in [2], where a background to foreground reflectance ratio is introduced. In a related development, Dror *et al.* [3] have shown how surfaces may be classified from single images through the use of reflectance properties. Moreover, although shape-from-shading usually relies on the assumption of Lambertian reflectance [4], photometric correction or specular subtraction may be applied as a preprocessing step to improve the results obtained.

In multispectral imaging, photometric invariants pose great opportunities and challenges in the areas of material identification [5] and shape analysis. This is due to the information-rich representation of the surface radiance acquired by multispectral and hyperspectral sensing devices, which acquire wavelength-indexed data in thousands of bands across a broad spectrum. As in trichromatic imaging, the multispectral reflectance

* NICTA is funded by the Australian Government as represented by the Department of Broadband, Communications and the Digital Economy and the Australian Research Council through the ICT Centre of Excellence program.

distribution is determined not only by the light source and viewer directions, but also by the material properties of the surface under study.

The main bulk of work concentrates in the effects encountered on shiny or rough surfaces. For shiny surfaces, there are specular spikes and lobes which must be modelled. There have been several attempts to remove specularities from images of non-Lambertian objects. For instance Brelstaff and Blake [6] use a thresholding strategy to identify specularities on moving curved objects. Ragheb and Hancock [7] have developed a probabilistic framework for specular subtraction which uses the Torrance and Sparrow model to account for the distribution of image brightness. Despite being effective, the application of these methods to multispectral imagery is somewhat limited since they rely on the closed form of the Bidirectional Reflectance Distribution Function (BRDF) to characterise the specular spike and lobe. Moreover, other alternatives elsewhere in the literature aimed at detecting and removing specularities either make the use of additional hardware [1], impose constraints on the input images [8] or require color segmentation [9] as postprocessing steps. Hence, they are not readily applicable to multispectral images, as there can be tens or even hundreds of bands for each pixel. Thus, any pre/postprocessing and local operations must be exercised with caution and in relation to neighbouring spectral bands so as to prevent spectral signature variation.

Specific to multispectral imagery, Healey and his colleagues [10] have addressed the problem of photometric invariance for material classification and mapping in aerial imaging as related to photometric artifacts induced by atmospheric effects and changing solar illumination. In [11], a photometrically invariant approach was proposed based on the derivative analysis of the spectra. This local analysis of the spectra was shown to be intrinsic to the surface albedo. Yet, the analysis in [11] was derived from the Lambertian reflection model and, hence, its not applicable to specular reflections.

2 Recovery of the Reflection Model Parameters

Here, we present a structural approach based upon the processing of smooth surface-patches whose spectral reflectance is uniform over all the pixels they comprise. As mentioned earlier, the process of recovering the photometric parameters for the image is based on an optimisation method which aims at reducing the difference between the estimate yielded by the dichromatic model and the radiance image data. In this section, we commence by providing an overview of the dichromatic model as presented by Shafer [12]. Subsequently, we formulate a target minimisation function with respect to the model in [12] and an optimisation strategy based upon the radiance structure drawn from smooth image patches with uniform reflectance. Throughout the section, we also present our strategy of selecting patches used by the algorithm and describe in detail the coordinate descent optimisation procedure used here. This optimisation strategy is such that we fix in interleaved steps either the light spectrum or the surface properties (geometry and reflectance) so as to recover the optima of the remaining dichromatic parameters.

2.1 The Dichromatic Reflection Model

Throughout the paper, we employ the dichromatic model introduced by Shafer [12] so as to relate light spectral power, surface reflectance and surface radiance. Following this model, surface radiance is decomposed into a diffuse and a specular component. Let an object with surface radiance $I(\lambda, u)$ at location u and wavelength λ be illuminated by an unknown illuminant whose spectrum is $L(\lambda)$. The dichromatic model becomes

$$I(\lambda, u) = g(u)L(\lambda)S(\lambda, u) + k(u)L(\lambda) \quad (1)$$

In Equation 1, the shading factor $g(u)$ governs the proportion of diffuse light and depends solely on the surface geometry. We aim to recover the shading factor $g(u)$, the specular coefficient $k(u)$, the light spectrum $L(\lambda)$ and the spectral reflectance $S(\lambda, u)$ at location u and wavelength λ from the spectral radiance $I(\lambda, u)$ on the image.

2.2 Target Function

With the dichromatic model above, we proceed to define our target function for purposes of optimisation. At input, our algorithm takes a multispectral image whose pixel values correspond to the measurements of the spectral radiance $I(\lambda, u)$ over the bands indexed to the wavelengths $\lambda \in \{\lambda_1, \dots, \lambda_n\}$. In general, the dichromatic cost function of a multispectral image \mathcal{I} is the weighted sum of its dichromatic error and a geometric regularisation term $R(u)$ for each image location.

$$F(\mathcal{I}) \triangleq \sum_{u \in \mathcal{I}} \left[\sum_{i=1}^n [I(\lambda_i, u) - L(\lambda_i)(g(u)S(\lambda_i, u) + k(u))]^2 + \alpha R(u) \right] \quad (2)$$

In equation 2, α is a constant that acts as a balancing factor between the dichromatic error and the regularisation term $R(u)$ on the right-hand side. The wavelength-independent regularisation term $R(u)$ is related to the scene geometry and will be elaborated upon later.

For now, we focus our attention on the solution space of Equation 2. Note that minimising the cost $F(\mathcal{I})$ without further constraints is an underdetermined problem. This is due to the fact that, for an image with n spectral bands containing m pixels each, we would have to minimise over $2m + n + m \times n$ variables while having only $m \times n$ terms in the summation of Equation 2. However, we notice that this problem can be further constrained if the model is applied to smooth surfaces made of the same material, i.e. the albedo is regular across the patch under consideration. This imposes two constraints. Firstly, all locations on the surface share a common diffuse reflectance. Therefore, a regular surface P is assumed to have the same reflectance for each pixel $u \in P$, $S(\lambda_i, u) = S_P(\lambda_i)$. Note that this constraint significantly reduces the number of unknowns $S(\lambda_i, u)$ from $m \times n$ to $n \times N$, where N is the number of surface albedos in the scene. In addition, the smooth variation of the patch geometry allows us to formulate the regularisation term $R(u)$ in equation 2 as a function of the shading factor $g(u)$. In brief, smooth, regular surface patches naturally provide constraints so as to reduce the number of unknowns significantly, while providing a plausible formulation of the geometric regularisation term $R(u)$.

Following the rationale above, we proceed to impose constraints on the minimisation problem at hand. For smooth, regular surface patches, we consider the cost function for a surface patch $P \in \mathcal{I}$ and write

$$F(P) \triangleq \sum_{u \in P} \sum_{i=1}^n [I(\lambda_i, u) - L(\lambda_i)(g(u)S_P(\lambda_i) + k(u))]^2 + \alpha \sum_{u \in P} R(u)$$

As before, we have $S(\lambda_i, u) = S_P(\lambda_i)$, for all $u \in P$. Furthermore, the smoothness constraint on the patch implies that the shading factor $g(u)$ should vary smoothly across the pixels in P . This constraint can be effectively formulated by minimising the variation in the magnitude of the shading map. This, effectively, precludes discontinuities in the map of P via the regularisation term

$$R(u) \triangleq \left[\frac{\partial g(u)}{\partial x(u)} \right]^2 + \left[\frac{\partial g(u)}{\partial y(u)} \right]^2 \quad (3)$$

where the variables $x(u)$ and $y(u)$ are the x and y coordinates at pixel location u .

Thus, by making use of the set \mathcal{P} of regular patches in the image \mathcal{I} , we can recover the dichromatic model parameters by minimising the target function

$$F^*(\mathcal{I}) \triangleq \sum_{P \in \mathcal{P}} \sum_{u \in P} \left[\sum_{i=1}^n [I(\lambda_i, u) - L(\lambda_i)(g(u)S_P(\lambda_i) + k(u))]^2 + \alpha R(u) \right] \quad (4)$$

2.3 Homogeneous Patch Selection

In the previous section, we formulated the recovery of the dichromatic model parameters as an optimisation procedure over the surface patch-set \mathcal{P} . In this section, we describe our method for automatically selecting regular surface patches for the minimisation of the cost function in Equation 4. The automatic patch selection method presented here allows the application of our method to arbitrary images. It is worth noting that this contrasts with other methods elsewhere in the literature [13,14], which are only applicable to pre-segmented images.

Our patch selection strategy is performed in two steps. Firstly, we identify edge pixels by applying a traditional edge detector such as that in [15] or the one presented by Canny [16] to the mean radiance across all spectral bands. Secondly, we subdivide the mean-radiance image into patches of equal size in a lattice-like fashion. Those patches containing less than t_p edge pixels, where t_p is a global parameter, are deemed to have a regular albedo. Note that, however, not all these patches are useful for purposes of illumination spectrum recovery. This is due to the fact that perfectly diffuse surfaces do not provide any information on the illuminant spectrum. The reason being that, for every band, the radiance of perfectly diffuse patches is given by the product of both, the illuminant spectrum and the surface reflectance. In the other hand, the dichromatic model implies that the specularities have the same spectrum as the illuminant, where the specular term can be viewed as a scaling factor solely dependent on the surface geometry.

We take profit of this dependency for purposes of recovering the light spectrum. We do this based upon the observation that, for highly specular surfaces, the surface diffuse

radiance vector and the illuminant vector span a plane in the radiance space. This is a well known property in colour constancy, where a number of approaches [17,18] have employed subspace projection for purposes of light power spectrum recovery. As a result, for the recovery of the dichromatic model parameters, we use patches with the highest contrast amongst those deemed as regular albedo ones. We estimate the contrast of each patch as the variance of the average spectral radiance at each pixel.

2.4 Optimisation Procedure

Making use of the notation in Figure 1, we now present the optimisation procedure employed in our method. Here, we adopt an iterative approach so as to find the variables \mathbf{L} , \mathbf{S}_P , \mathbf{g}_P and \mathbf{k}_P which correspond to the minimum of the cost function in Equation 4. At each iteration, we minimise the cost function with respect to \mathbf{L} and the triplet $[\mathbf{S}_P, \mathbf{g}_P, \mathbf{k}_P]$ of each patch P .

The algorithm commences with initialising the unknown light spectrum $L(\lambda)$. We note that objects in the scene reflect less energy than they receive from the illuminant, i.e. the radiance of a surface cannot be greater than the irradiance of the illuminant, the highest radiance across image pixels for each wavelength can be used as a lower bound estimate of the light energy. Thus, as an initialisation for \mathbf{L} we employ the light spectrum estimated by the brightest pixel per band.

Now we elaborate on the coordinate descent approach. This is comprised by two interleaved minimisation steps. At each iteration, we can index the dichromatic variables to iteration number t and optimise, in interleaved steps, the variables \mathbf{g}_P , \mathbf{k}_P and \mathbf{S}_P with respect to the image radiance and the illuminant \mathbf{L} , and once the former are at hand, obtain a closed form solution for the latter.

In the first step, we estimate the optimal surface reflectance and geometry given the light spectrum \mathbf{L}^{t-1} . Note that, at iteration t , we can solve for the unknowns \mathbf{g}_P^t , \mathbf{k}_P^t and \mathbf{S}_P^t separately for each surface patch P . This is because these variables correspond to separate terms in Equation 4. This step is, therefore, reduced to minimising

$$\begin{aligned} F(P)|_{\mathbf{L}^{t-1}} &= \sum_{u \in P} [\|\mathbf{I}(u) - \langle \mathbf{L}^{t-1}, g(u)\mathbf{S}_P + k(u) \rangle\|^2 + \alpha R(u)] \\ &= \sum_{u \in P} [\|\mathbf{I}(u) - g(u)\mathbf{D}_P^{t-1} - k(u)\mathbf{L}^{t-1}\|^2 + \alpha R(u)] \end{aligned} \quad (5)$$

$\mathbf{I}(u)$: the spectral radiance vector at image pixel u , $\mathbf{I}(u) = [I(\lambda_1, u), \dots, I(\lambda_n, u)]^T$
\mathbf{L} : the spectral power vector of the illuminant, $\mathbf{L} = [L(\lambda_1), \dots, L(\lambda_n)]^T$
\mathbf{S}_P : the common spectral reflectance vector for each patch P , $\mathbf{S}_P = [S_P(\lambda_1), \dots, S_P(\lambda_n)]^T$
\mathbf{g}_P : the shading map of all pixels in patch P , $\mathbf{g}_P = [g(u_1), \dots, g(u_l)]^T$ with u_1, \dots, u_l being all the pixels in the patch P
\mathbf{g} : the shading map of the whole image, $\mathbf{g} = [\mathbf{g}_{P_1}, \dots, \mathbf{g}_{P_r}]^T$ where P_1, \dots, P_r are all patches in \mathcal{P}
\mathbf{k}_P : the specular map of all pixels in patch P , $\mathbf{k}_P = [k(u_1), \dots, k(u_l)]^T$
\mathbf{k} : the specular map of the whole image, $\mathbf{k} = [\mathbf{k}_{P_1}, \dots, \mathbf{k}_{P_r}]^T$

Fig. 1. Notations used in Section 2.4

where $\|\cdot\|$ and $\langle \cdot, \cdot \rangle$, respectively, denote the inner product and the L_2 -norm of the argument vectors and $\mathbf{D}_P^{t-1} \triangleq \langle \mathbf{L}^{t-1}, \mathbf{S}_P \rangle$ is the diffuse radiance component for the surface patch P .

Note that the minimisation above involves $2|P| + n$ unknowns, where $|P|$ is the number of pixels in patch P . Hence, it becomes computationally intractable when the surface area is large. In practice, the selected patches need only be large enough so as to gather useful statistics from the radiance information. Moreover, as mentioned earlier, we can further reduce the degrees of freedom of the unknowns by noting that the spectral radiance vectors at all pixels in the same patch lie in a 2-dimensional subspace Q of \mathbb{R}^n spanned by the diffuse radiance vector \mathbf{D}_P^{t-1} and the light vector \mathbf{L}^{t-1} . This is a characteristic of the dichromatic model that has been widely utilised by prior work on colour constancy [13,14].

Having all the pixel radiance vectors $\mathbf{I}(u)$ at hand, one can obtain the 2-dimensional subspace Q via Single Value Decomposition (SVD). Denote the two basis vectors resulting from this SVD operation z_1 and z_2 and, accordingly, let the subspace $Q = \text{span}(z_1, z_2)$. Since $\mathbf{D}_P^{t-1} \in Q$, we can parameterise \mathbf{D}_P^{t-1} up to scale as $\mathbf{D}_P^{t-1} = v z_1 + z_2$. Likewise, the light vector $\mathbf{L}^{t-1} \in Q$ can also be decomposed as $\mathbf{L}^{t-1} = w_1 z_1 + w_2 z_2$, where the values of w_1 and w_2 are two known scalars. Moreover, \mathbf{D}_P^{t-1} is a linear function with respect to v . We also observe that the term $R(u)$ is only dependent on $g(u)$. Therefore, the objective function in Equation 5 can be reduced to a univariate rational function of v with degree 4, and can be minimised efficiently making use of the Levenberg-Marquardt algorithm [19]. Having obtained the optimal value of v , we can recover \mathbf{g}_P , \mathbf{k}_P and \mathbf{S}_P , knowing the illuminant spectrum \mathbf{L} .

In the second step of each iteration t , we solve for \mathbf{L}^t given \mathbf{g}_P^t , \mathbf{k}_P^t and \mathbf{S}_P^t . Since the second term $R(u)$ in Equation 4 is independent of the illuminant, whereas the first term is a quadratic function of \mathbf{L} . Therefore, the illuminant spectrum can be solved in closed form as

$$\mathbf{L}^t(\lambda) = \frac{\sum_{P \in \mathcal{P}} \sum_{u \in P} I(\lambda, u)(g^t(u)S_P^t(\lambda) + k^t(u))}{\sum_{P \in \mathcal{P}} \sum_{u \in P} (g^t(u)S_P^t(\lambda) + k^t(u))^2}$$

The closed form solution above can be easily obtained by differentiating Equation 4 with respect to the illuminant irradiance and equating the resulting expression to zero.

3 Experiments

Here, we present results on the recovery of the illuminant spectra and skin recognition. To this end, we have acquired an image database of 51 human subjects, each captured under one of 10 light sources with varying directions and spectral power. The imagery has been acquired using a pair of OKSI Turnkey Hyperspectral Cameras. The light sources are divided into two rows, one placed above the camera system and one at the same height as the cameras. The light positions are evenly distributed in each row so that the illumination directions are towards the centre of the scene. The camera system is equipped with Liquid Crystal Tunable Filters which allow multi-spectral images to be resolved up to $10nm$ in both the visible ($430-720nm$) and near infrared ($650-990nm$) wavelength ranges. To obtain the ground truth for the illuminant spectrum, we have

Table 1. The accuracy of our illuminant estimation method, in degrees, compared to Finlayson & Schaefer’s method [13]

Number of patches	Visible spectrum		Near-infrared spectrum	
	Our method	Finlayson & Schaefer	Our method	Finlayson & Schaefer
1	10.87 ± 5.70	17.34 ± 9.04	5.98 ± 4.38	24.67 ± 9.46
2	10.77 ± 5.22	11.20 ± 6.18	5.06 ± 3.69	6.81 ± 5.57
4	10.94 ± 4.92	11.09 ± 5.32	4.47 ± 3.03	6.21 ± 4.89
5	11.26 ± 4.35	11.59 ± 4.16	4.21 ± 2.86	5.96 ± 4.15
10	11.26 ± 4.35	11.59 ± 4.16	4.21 ± 2.86	5.96 ± 4.15
20	11.51 ± 4.04	11.92 ± 3.65	4.08 ± 2.59	5.86 ± 3.81
30	11.62 ± 3.93	12.03 ± 3.52	4.08 ± 2.57	5.87 ± 3.65

measured the average radiance reflected from a white calibration target, i.e. a LabSphere Spectralon, illuminated by the light sources under consideration.

3.1 Illumination Spectrum Recovery

For our experiments on illumination spectra recovery, we compare the performance of our method with that yielded by the method of Finlayson and Schaefer [13]. We choose the work in [13] since it is based on the dichromatic model and can be adapted from trichromatic imagery to multispectral data in a straightforward manner.

In our experiment, we run our illuminant spectrum estimation algorithm on 1, 2, 4, 5, 10, 20 and 30 most contrasting patches of size 16×16 pixels from each image. The accuracy of light spectrum recovery is measured as the Euclidean angle between the estimated and ground truth spectrum in n dimensions, where n is the number of sampled bands for both spectra. Table 1 shows the mean and standard deviations of this angle, in degrees, over all images in the database versus the number of selected patches, in both the visible and infrared spectrum.

As shown in Table 1, our algorithm consistently achieves a higher accuracy than the alternative method. This is across all the number of patches for both spectral ranges. In general, the reported results match the expectation that both methods gain better accuracy with an increasing number of patches. However, it is noticeable that with a single patch, our method significantly outperforms the alternative method. This observation is consistent with the well-known fact that the dichromatic plane method and its variants require at least two homogeneous surfaces to compute the intersection between the dichromatic planes. Moreover, the deviation angle of the estimated illuminant decreases slowly with the growing number of patches, which indicates that our method is stable and does not require many homogeneous surfaces to achieve a reasonable estimate.

To illustrate the statistics in Table 1, we plot in Figure 2 both the visible and infrared spectra of the ground truth and the illuminant spectra of a sample light estimated from the image of a human subject in the database. All of the spectra are normalised by the highest spectral power. The ground truth spectra, the one estimated by our method and that recovered by Finlayson and Schaefer’s algorithm, are drawn in red, blue and green, respectively. This illustrates the robustness our method, even when a single patch is used. Note that, this contrasts with the results delivered by the alternative method.

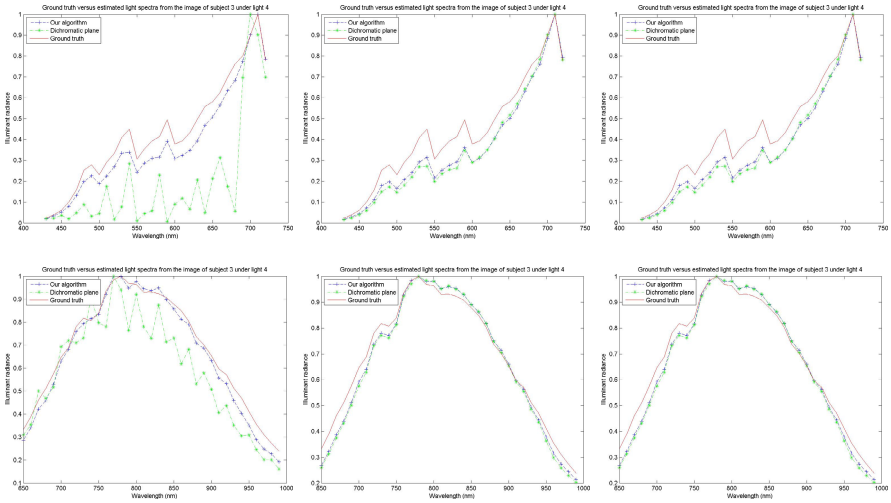


Fig. 2. Ground truth illuminant spectra and that delivered by our method and the alternative as estimated from the image of one subject in the database for both, the visible (top row) and near-infrared ranges (bottom row). From left to right: the spectra estimated from 1, 5 and 10 patches in the image.

3.2 Skin Recognition

We now turn our attention to the illumination invariance of the spectral reflectance recovered by our algorithm and its applications to recognition tasks. Specifically, we focus on using the surface reflectance extracted from an image for skin segmentation. This task can be viewed as a classification problem where the skin and non skin spectra comprise positive and negative classes. To perform this task, we select skin and non skin regions from the image of a subject taken under one of the illumination conditions in the data set and store the normalised reflectance spectra of these pixels as training data. Subsequently, we train a Support Vector Machine (SVM) classifier with a polynomial kernel of degree 2 on the training set. The resulting SVM model is applied for classifying skin versus non skin pixels in the remaining images of the subject, which have been acquired under diverse illuminant spectral conditions. In this manner, we can assert how robust are both, the illuminant spectrum recovered by our algorithm at training time and the normalised reflectance yielded by the method for skin recognition.

In our experiments, we also compare the results yielded by the SVM trained on the normalised reflectance spectra delivered by our method with those recovered using alternative recognition methods. To this end, we have used raw radiance spectra and principal components resulting from Linear Discriminant Analysis (LDA) projection of the original radiance spectra. The raw spectra and the LDA components are then used as input to SVMs whose parameters have been recovered using cross validation.

The top row of Figure 3 contains the training images for a number of sample subjects. Each training image shows the pixel regions from which skin and non-skin training data are collected. These have all been illuminated by the light source placed in a



Fig. 3. Skin segmentation results. Top row: sample training images for skin recognition; Second row: Skin probability map yielded by the recovered spectral reflectance on the image of the respective subject taken under another lighting condition; Third row: Skin probability map recovered making use of the normalised raw radiance data; Fourth row: Skin probability map recovered making use of 20 top principal components resulted from LDA projection on raw radiance data.

high-oblique position in front of the subject. For purposes of training, in all cases, we have employed 5 skin and 18 non-skin rectangular regions, of sizes 25×17 and 24×16 , respectively. The second row shows the skin probability maps yielded by our method on images of the same subjects under a different illuminant in the dataset. The third and fourth rows, respectively, show the skin probability map for the input image used in the second row as recovered by making use of the normalised raw radiance spectra and the 20 top principal components obtained from LDA as the classification feature vectors. In these skin maps, lighter pixels are classified as more likely to be skin.

Figure 3 demonstrates that the recovered skin reflectance spectra are invariant to illuminant variations and light source direction changes. Indeed, the skin segmentation maps in the second row show a substantial proportion of skin pixels correctly classified. In addition, non-skin face details such as eye brows and mouth are correctly distinguished from skin. On the contrary, using normalised raw radiance data and LDA components for skin recognition, as shown in the third and fourth rows, causes much more false negatives in skin areas and false positives in other materials. The poor classification results in these cases can be explained by the variation induced by the illuminant spectrum and the surface geometry.

4 Conclusions

This paper has discussed an optimisation approach to the photometric invariance problem from a single multispectral image. The recovery process is based on optimising a cost function that imposes the smoothness constraint on uniform dichromatic reflectance surfaces. The method serves several purposes, including scene illuminant

estimation, surface geometry recovery and illumination-invariant reflectance extraction. Firstly, the resulting illuminant spectra are shown to outperform the competitive method making use of the dichromatic plane hypothesis [13]. This is particularly evident when the number of homogeneous surface patches is limited. Secondly, we have shown how the invariance of the reflectance spectra recovered by our algorithm can be used for purposes of skin recognition.

References

1. Nayar, S., Bolle, R.: Reflectance based object recognition. *International Journal of Computer Vision* 17(3), 219–240 (1996)
2. Nayar, S.K., Bolle, R.M.: Computing reflectance ratios from an image. *Pattern Recognition* 26, 1529–1542 (1993)
3. Dror, R.O., Adelson, E.H., Willsky, A.S.: Recognition of surface reflectance properties from a single image under unknown real-world illumination. In: *Proc. of the IEEE Workshop on Identifying Objects Across Variations in Lighting* (2001)
4. Horn, B.K.P., Brooks, M.J.: The variational approach to shape from shading. *CVGIP* 33(2), 174–208 (1986)
5. Chang, J.Y., Lee, K.M., Lee, S.U.: Shape from shading using graph cuts. In: *Proc. of the Int. Conf. on Image Processing*, vol. I, pp. 421–424 (2003)
6. Brelstaff, G., Blake, A.: Detecting specular reflection using lambertian constraints. In: *Int. Conf. on Comp. Vision*, pp. 297–302 (1988)
7. Ragheb, H., Hancock, E.R.: A probabilistic framework for specular shape-from-shading. *Pattern Recognition* 36(2), 407–427 (2003)
8. Lin, S., Shum, H.: Separation of diffuse and specular reflection in color images. In: *IEEE Conf. on Computer Vision and Pattern Recognition*, vol. I, pp. 341–346 (2001)
9. Klinker, G., Shafer, S., Kanade, T.: A physical approach to color image understanding. *Intl. Journal of Computer Vision* 4(1), 7–38 (1990)
10. Suen, P.H., Healey, G.: Invariant mixture recognition in hyperspectral images. In: *Int. Conference on Computer Vision*, pp. 262–267 (2001)
11. Angelopoulou, E.: Objective colour from multispectral imaging. In: *European Conf. on Computer Vision*, pp. 359–374 (2000)
12. Shafer, S.A.: Using color to separate reflection components. *Color Research and Applications* 10(4), 210–218 (1985)
13. Finlayson, G.D., Schaefer, G.: Convex and Non-convex Illuminant Constraints for Dichromatic Colour Constancy. *CVPR* 1, 598–604 (2001)
14. Tominaga, S., Wandell, B.A.: Standard surface-reflectance model and illuminant estimation. *Journal of the Optical Society of America A* 6, 576–584 (1989)
15. Marr, D., Hildreth, E.: Theory of edge detection. In: *Proceedings of Royal Society, London*, vol. B–207, pp. 187–217 (1980)
16. Canny, J.: A computational approach to edge detection. *IEEE Trans. Pattern Anal. Mach. Intell.* 8(6), 679–698 (1986)
17. Klinker, G.J., Shafer, S.A., Kanade, T.: The measurement of highlights in color images. *International Journal of Computer Vision* 2, 7–32 (1988)
18. Lee, H.C.: Method for computing the scene-illuminant chromaticity from specular highlights. *Journal of the Optical Society of America A* 3, 1694–1699 (1986)
19. Marquardt, D.W.: An algorithm for least-squares estimation of nonlinear parameters. *SIAM Journal on Applied Mathematics* 11(2), 431–441 (1963)