Efficient Estimation of Reflectance Parameters from Imaging Spectroscopy

Lin Gu, *Student Member, IEEE*, Antonio Robles-Kelly, *Senior Member, IEEE*, and Jun Zhou, *Senior Member, IEEE*

Abstract—In this paper, we address the problem of efficiently recovering the reflectance parameters from a single multispectral or hyperspectral image. To do so, we propose a Shapelet Based Estimator (SBE) which employs shapelets to recover the shading in the image. The optimisation setting presented here is based upon a three-step process. The first of these concerns the recovery of the surface reflectance and the specular coefficients via a constrained optimisation approach. Secondly, we update the illuminant power spectrum using a simple least-squares formulation. Finally, the shading is computed directly once the updated illuminant power spectrum is at hand. This yields a computationally efficient method which achieves speed-ups of nearly an order of magnitude over its closest alternative without compromising performance. We provide results on illuminant power spectrum computation, shading recovery, skin recognition and replacement of the scene illuminant and object reflectance in real-world images.

Index Terms—Keywords: Multispectral and hyperspectral imaging, Scene analysis, Inverse methods for imaging spectroscopy

I. INTRODUCTION

The recovery of shape and photometric invariants through imaging spectroscopy can provide a means for identification and recognition purposes. Further, photometric invariants can be used in combination with the object shape and light power spectrum for advanced applications in photowork, visualisation, creation and manipulation of digital media [1]. This is illustrated in Figure 1, where we show the result of substituting the skin reflectance in a portrait.

Despite its prospective applications, the recovery of photometric invariants from imaging spectroscopy is not a straightforward task. This is since every scene comprises a rich tapestry of light sources, material reflectance and lighting effects due to object curvature and shadows. Moreover, trichromatic (*i.e.* RGB) technology does have limits in its scene analysis capabilities. For example, a camera with a RGB sensor is prone to metamerism and, therefore, can not deliver

photometric invariants characteristic to a material independent of the lighting or shape.

Along these lines, in computer vision and image processing, most of the efforts in the literature have been devoted to modelling the effects encountered on shiny or rough surfaces. For shiny surfaces, specular spikes and lobes have to be modelled. There have been several attempts to remove specularities from images of non-Lambertian objects, *i.e.* the objects whose surface brightness is the same regardless of the observer angle of view. For instance Brelstaff and Blake [2] used a simple thresholding strategy to identify specularities on moving curved objects. Other lines of research remove specularities by either using additional hardware [3], imposing constraints on the input images [4], requiring color segmentation [5] as postprocessing steps, or using reflectance models to account for the distribution of image brightness [6].

Moreover, the methods above often rely on the use of the closed form of the Bidirectional Reflectance Distribution Function (BRDF). As a result, they lack the generality required to process real-world imagery in an unsupervised or automatic way. Moreover, for rough surfaces, the main bulk of research focuses on the loss of contrast at the object limbs and an apparent flattening of perceived relief [7], [8], [9], [10]. Hence, the accurate modelling of object limb brightening for rough surface reflectance has proved to be extremely elusive, and is still an active area of research in applied optics.

The applicability of existing methods to imaging spectroscopy is further limited by the difficulty in extending these methods to images comprised of tens or even hundreds of bands. This is mainly due to the complexity regarding a closed form solution for the high-dimensional image data involved, but is also related to the fact that the BRDF parameters become wavelength dependent and, hence, are often intractable. Moreover, any pre/postprocessing and local operations on the reflectance must be exercised with caution and in relation to neighbouring spectral bands so as to prevent spectral signature variation.

To solve these problems, Healey and his colleagues [11], [12] studied photometric invariance in hyperspectral aerial imagery for material classification and mapping as related to photometric artifacts induced by atmospheric effects and changing solar illuminations. In [13], a method is presented for hyperspectral edge detection. The method is robust to photometric effects, such as shadows and specularities. In [14], 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 reflectance.

Copyright (c) 2013 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

L. Gu is with the College of Engineering and Computer Science at the Australian National University, Canberra, ACT 2601, Australia.

A. Robles-Kelly is with the National ICT Australia (NICTA) and the College and Engineering and Computer Science at the Australian National University, Canberra, ACT 2601, Australia.

J. Zhou is with the School of Information and Communication Technology at Griffith University, Nathan, QLD 4111, Australia.

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.

Fig. 1. Initial (left-hand panel) and final (right-hand panel) pseudocolour images when the skin reflectance of the subject is substituted by an alternative. The left-hand panel shows the pseudocolour image of a portrait taken using a hyperspectral camera. The skin reflectance of the subject is plotted on the top middle panel whereas the one used as an alternative is shown in the bottom middle plot. The right-hand panel, has been recovered by effecting skin recognition so as to locate those pixels whose spectra are to be replaced. The image irradiance has then been recomputed using the photometric parameters recovered by our method and the alternative reflectance. Note that a similar sequence of events would apply to the replacement of the illuminant in the scene, where the power spectrum of the light has to be recovered and substituted.

Yet, the analysis in [14] was derived from the Lambertian reflection model and, hence, is not applicable to specular reflections. Recently, Huynh and Robles-Kelly [15] proposed a method based upon an the dichromatic model [16] aimed at imaging spectroscopy. Despite effective, the method can be computationally demanding since it relies upon a dual-step gradient descent optimisation scheme governed by a regularised cost function.

Estimation of reflectance parameters is also closely related to the problem of recovering the shape of an object from its shading information. The classic approaches to shape from shading developed by Ikeuchi and Horn [17], and by Horn and Brooks [18], hinge on the compliance with the image irradiance equation and local surface smoothness. Zheng and Chellappa [19] proposed a gradient consistency constraint that penalises differences between the image intensity gradient and the surface gradient for the recovered surface. Worthington and Hancock [20] impose the Lambertian radiance constraint in a hard manner. Dupuis and Oliensis [21] provided a means of recovering probably correct solutions with respect to the image irradiance equation. Prados and Faugeras [22] have presented a shape from shading approach applicable to the case where the light source and viewing directions are no longer co-incident. Kimmel and Bruckstein have used the apparatus of level sets methods to recover solutions to the Eikonal equation [23].

Here, we present a method based upon shapelets [24] which recovers the reflection parameters, *i.e.* shading, specularities, light spectral power and surface reflectance, from the image radiance. Note that the capacity to recover the power spectrum of the illuminant, reflectance and shading is reminiscent of the method in [15]. Despite its effectiveness, the work presented in [15] can be computationally costly. Moreover, the method presented here permits the use of robust statistics [25] and makes use of an alternative formulation which yields a

different optimisation scheme. This optimisation strategy can be further simplified to yield improvements in computational efficiency of an order of magnitude over the method in [15].

The paper is organised as follows. We commence by introducing the motivation to the method presented in this paper and elaborating upon the relevant background in the next section. In Section II, we present the dichromatic reflection model and provide a link between the work in [16]. We discuss the implementation issues of our method, its initialisation and the extensions to trichromatic imagery in Section III. In Section IV, we present results for our method and, finally, we conclude on the developments presented here in Section V.

II. ESTIMATING THE REFLECTION PARAMETERS

To recover the reflection parameters, we employ the Dichromatic model in [16]. This model has long been used in physics-based vision for characterising specular reflections on non-Lambertian surfaces. The model states that the radiance for a pixel in the image can be expressed as follows

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

where $L(\lambda_i)$ is the light source spectral component as a function of the i^{th} band with wavelength λ_i , $S(\lambda_i, u)$ is the surface reflectance and $g(u)$ and $k(u)$ are the shading factor and specular coefficient at the pixel location u . These are wavelength-independent and are governed by the illumination and surface geometries.

In Equation 1, the first term on the right-hand side corresponds to the diffuse reflection, whereas the second term accounts for the specularities. We can give an intuitive interpretation to the model as follows. At specular pixels, *i.e.* where the shading factor $g(u)$ is negligible, the surface acts as a mirror, where the radiance becomes proportional to the illuminant power spectrum $L(\lambda_i)$. On the other hand, if the

radiance is diffuse, *i.e.* $k(u)$ is close to zero, then the surface appears shaded and the reflectance becomes a multiplicative term on the illuminant to determine the radiance. This can be easily verified by writing

$$\frac{I(\lambda_i, u)}{L(\lambda_i)} = r(\lambda_i, u) + k(u) \quad (2)$$

where

$$r(\lambda_i, u) = g(u)S(\lambda_i, u) \quad (3)$$

is devoid of the illuminant spectrum $L(\lambda_i)$ and proportional to the reflectance through the shading factor $g(u)$.

Note that, following the model in Equation 2, it would be straightforward to recover the specular coefficient $k(u)$ if the other parameters were at hand. Similarly, by assuming the illuminant spectrum and the specular coefficient $k(u)$ are available, one could normalise the irradiance and recover the term $r(\lambda_i, u)$. This suggests an iterative process involving the recovery of the shading factor, estimation of the specular coefficient $k(u)$ and reflectance so as to achieve a solution for the illuminant power spectrum $L(\lambda_i)$.

For now, we will assume initial estimates for the reflection model parameters are at hand. We will discuss the computation of the initial parameter values in Section III. With these initial values, we can estimate the photometric parameters via a three-step iterative process. Firstly, the shading factor is computed. Secondly, the surface reflectance and the specular coefficients are computed using a constrained optimisation approach. Finally, the illuminant power spectrum can be recovered from the above parameters that have been recovered. These three interleaved steps are applied until convergence is reached.

A. Recovery of the Shading Factor

As mentioned above, we employ shapelets so as to recover the shading factor in Equation 2. The idea underpinning our method is that the shading factor $g(u)$ can be viewed as that governing the diffuse reflection from the surface. As a result, we can use Lambert's Law [26] and set

$$g(u) = \cos(\alpha, u) \quad (4)$$

where α is the angle between the light source direction and the surface normal. This angle can be computed via the inner product of the surface normal and the illuminant direction in a straightforward manner from the surface normal if the light source direction is known. Later on in this section, we discuss how the requirement pertaining the light source direction to be known can be removed. For now, we assume the light source direction is available and focus on the computation of the shading factor.

If the 3D representation of the object surface is at hand, the shading factor computation becomes the matter of a simple inner product. Unfortunately, the recovery of the 3D representation of the surface is a non-trivial task. In [24], shapelets are used as a complete orthonormal set of basis functions so that the 3D representation of the surface can be easily recovered despite the effects of noise corruption. Shapelets are a set of basis functions with finite support used in image analysis in a manner akin to Fourier or wavelet

synthesis. These basis functions or shapelets can be Hermite polynomials [27] or an over complete basis such as that in [28]. Since a complete treatise on shapelets is beyond the scope of this paper, we would like to refer the interested reader to the book of Mallat [29] or the paper by Refregier [27].

In [24], the recovery of the 3D surface is effected by correlating the slant and tilt of the field of surface normals with those of the shapelets. The main contribution in [24], hinges upon the recovery of the 3D shape from a sum of correlations between the surface gradients and the basis functions, *i.e.* the shapelets. In practice, the correlation can be computed via a convolution, whereas the tilt of the surface normal can be recovered using the direction of the image gradient on the image plane [24]. Here, we follow [24] and depart from a shapelet based upon the tangent function. As shown in [24], the 3D surface reconstruction H is given by the sum over the scales q for the shapelets $\tan(s_q)$, where s_q is the slant of the shapelet normals. Thus, the reconstruction is given by

$$\begin{aligned} H &= \sum_q |\tan(\alpha)| * |\tan(s_q)| \cdot |\cos(\tau_f) * \cos(\tau_q)| \\ &\quad + \sin(\tau_f) * \sin(\tau_q) \\ &= \sum_q |\tan(\alpha)| * |\tan(s_q)| \cdot C_{\tau_q} \end{aligned} \quad (5)$$

where τ_q and τ_f correspond to the tilt of the normals for both, the shapelet and the object under study, respectively. In Equation 5, α is the angle between the surface normal and the light direction and $*$ denotes convolution.

Unfortunately, the angle α remains a variable to be estimated and, therefore, the equation above cannot be used directly. Nonetheless, making use of Equations 2, 3 and 4, we can derive the following relation

$$\frac{1}{r(\lambda_i, u) - r(\lambda_j, u)} = \sec(\alpha, u) \frac{1}{(S(\lambda_i, u) - S(\lambda_j, u))} \quad (6)$$

Note that $\tan(\alpha)$, $\sec(\alpha)$ and $\sec(\alpha) - 1$ share an akin monotonically increasing behavior in $[0, \frac{\pi}{2}]$, whose traces are shown in Figure 2. In the figure, we also show a sample shape and the corresponding 3D reconstructions recovered using $\tan(\alpha)$ and $\sec(\alpha) - 1$. Note that the reconstructions yielded by both functions are in close accordance with each other and the ground truth. As a result, we can use $\sec(\alpha) - 1$ as an alternative to $\tan(\alpha)$ and approximate the shapelet reconstruction for each band by

$$H_{\lambda_i} \approx \sum_q \mathbf{R}_{\lambda_i} * |\tan(s_q)| \cdot C_{\tau_q} \quad (7)$$

where $\mathbf{R}_{\lambda_i}$ is a matrix whose entry indexed u is given by $\left| \frac{S(\lambda_i, u) - S(\lambda_j, u)}{r(\lambda_i, u) - r(\lambda_j, u)} - 1 \right|$.

Recall that $\tan(s_q)$ and C_{τ_q} are known and can be readily computed from the shapelet basis functions and the image gradients. Theoretically, any arbitrarily chosen reference band λ_j can be used for the difference above. Nonetheless, $r(\lambda_j, u)$ and $S(\lambda_j, u)$ may deviate from the ideal case as modelled by Equation 2 due to measurement noise. This is compounded by the fact that, in practice, the surface reconstruction can be recovered for each band in the image.

Fig. 2. Top left-hand panel: Behavior of $\tan(\alpha)$ vs $\sec(\alpha)$ and $\sec(\alpha) - 1$ in $[0, \frac{\pi}{2}]$; Top right-hand panel: Sample 3D block shape; Bottom row: Shape yielded by $\tan \alpha$ (left) and $\sec(\alpha) - 1$ (right).

To tackle both the ambiguity of the surface reconstruction to band index and noise corruption, we assume that the measurement errors are independent and identically distributed random variables governed by a Gaussian distributions with zero mean. It can be shown that under this assumption, the variance of the estimate can be much reduced through averaging [25]. Thus, instead of an arbitrary reference band, we use the mean of $r(\cdot, u)$ and $S(\cdot, u)$ given by

$$r_u^* = \frac{1}{N} \sum_{j=1}^N r(\lambda_j, u) \quad S_u^* = \frac{1}{N} \sum_{j=1}^N S(\lambda_j, u) \quad (8)$$

where N is the total number of bands in the image. For the computation of the matrix $\mathbf{R}_{\lambda_i}$ in Equation 7, we employ the resulting r_u^* and S_u^* from the equations above as an alternative to the quantities $r(\lambda_j, u)$ and $S(\lambda_j, u)$ corresponding to an arbitrary band indexed j .

Following a similar rationale regarding the surface reconstruction estimates for each band, we make use of the average

$$H^* = \frac{1}{N} \sum_{j=1}^N H_{\lambda_j} \quad (9)$$

as an alternative to H in Equation 5. With the surface reconstruction H^* at hand, the quantity $\cos(\alpha, u)$ at each pixel

u can be computed making use of the light source direction and the surface normal field over H^* .

B. Computation of the Reflectance and Specular Coefficient

We now make use of the illuminant and the shading factor to recover the the reflectance $S(\lambda_i, u)$ and the specular coefficient $k(u)$. Since the $\cos(\alpha, u)$ has been computed using shapelets as presented in the previous section, we can recovery the surface reflectance using the following equation

$$S(\lambda_i, u)^* = \frac{1}{\cos(\alpha, u)} \left(\frac{I(\lambda_i, u)}{L(\lambda_i)} - k(u)^* \right) \quad (10)$$

where $k(u)^*$ is the estimated specular coefficient given by

$$k(u)^* = \min_{k(u) \geq 0} \left\{ \sum_{j=1}^N |I(\lambda_j, u) - L(\lambda_j)(g(u)S(\lambda_j, u)^* - k(u))|^2 \right\}$$

The expressions above are consistent with both, the Dichromatic model presented earlier and the notion that the specular coefficient should be a positive real-valued quantity. In practice, the recovery of the specular coefficient can be effected making use of a constrained least squares approach. In Section III, we comment further on our implementation.

C. Illuminant Power Spectrum Computation

To estimate the illumination, we note that the light power spectrum $L(\lambda_i)$ in Equation 1 does not depend upon the pixel-wise variable u . Although the radiance sensed by the camera can be influenced by shadows and other changes such as the specular spike brightness, the spectral power distribution of the illuminant does not change. This indicates that the illuminant is a global photometric variable throughout the image. This observation is important since we can recover its spectrum by minimising the cost function given by

$$L(\lambda_i)^* = \arg \min_{L(\lambda_i)} \left\{ \sum_{u \in \mathcal{I}} |I(\lambda_i, u) - (\cos(\alpha, u)S(\lambda_i, u)^* + k(u)^*)L(\lambda_i)|^2 \right\} \quad (11)$$

where $\mathcal{I}$ is the set of all pixels in the imagery. Since the only unknown in the cost function above is the light spectrum $L(\cdot)$, the above minimisation becomes a least squares one whose solution is unique and computationally efficient.

1) *Simplification of the Model:* To simplify the method above, we note that, in real-world imagery, a number of pixels often show a large specular component. This is particularly evident at the specular spikes in the image. For these pixels, the specular coefficient is, in practice, much larger than the shading factor, *i.e.* $k(u) \gg g(u)$. This implies that, for those pixels with a large specular component, we have

$$I(\lambda_i, u) \approx k(u)L(\lambda_i) \quad (12)$$

where we have substituted the shorthand $g(u) \approx 0$ into Equation 1.

We note that, whereas accordance to the dichromatic plane implies a diffuse reflection, deviation from the dichromatic plane indicates large specular coefficients. As a result, those patches that do not conform to the dichromatic plane and deviate greatly, should conform to Equation 12. Thus, we can ignore the diffuse term $g(u)S(\lambda_i, u)$ in Equation 1 for those pixels comprising the set $\mathcal{P}$ of K -most specular patches and recover the illuminant power spectrum by minimising the expression

$$L(\lambda_i)^* = \arg \min_{L(\lambda_i)} \left\{ \sum_{u \in \mathcal{P}} |I(\lambda_i, u) - k(u)^*L(\lambda_i)|^2 \right\} \quad (13)$$

as an alternative to Equation 11.

2) *Application of Robust Statistics:* Note that, in practice, imaging spectroscopy is not devoid of noise corruption. Further, for our illumination recovery step, we use a least-squares minimisation scheme. This least-squares approach, despite effective, can be sensitive to outliers.

To tackle this problem, we can employ robust statistics as an alternative to least-squares. Robust statistics provide a means to noise corruption by introducing an estimator which is robust to outliers. There are a number of robust estimators that can be used, such as maximum likelihood (M-estimators), linear combination of order statistics (L-estimators), rank transformation (R-estimators) and repeated M-estimators (RM-estimators) [30], [25].

To apply these estimators to our illuminant recovery step,

we assume that the image radiance for those pixel-sites $u \in \mathcal{P}$ can be modeled as a linear combination of the true specularities $k(u)L(\lambda_i)$ and a noise component $\eta(u)$. The image radiance, *i.e.* the observable, becomes $\mathcal{I}(\lambda_i, u) = k(u)L(\lambda_i) + \eta(u)$. With these ingredients, we can set an error penalty function $g(\cdot)$ so as to provide an alternative to Equation 13 as follows

$$\begin{aligned} L(\lambda_i)^* &= \arg \min_{L(\lambda_i)} \left\{ \sum_{u \in \mathcal{P}} g(\mathcal{I}(\lambda_i, u) - k(u)^*L(\lambda_i)) \right\} \\ &= \arg \min_{L(\lambda_i)} \{g(\eta(u))\} \end{aligned} \quad (14)$$

D. On the Light Source Direction

Recall that, in the previous sections, we assumed that the light source direction $\vec{\mathbf{L}}$ was known. In practice, this may not be the case. To relax the restriction, we note that, according to Lambert's Law, for the pixel indexed u corresponding to the point p on H^* illuminated from the direction $\vec{\mathbf{L}}$, the shading factor $g(u)$ is given by $(\vec{\mathbf{N}}_p \cdot \vec{\mathbf{L}})$, where $\vec{\mathbf{N}}_p$ is the normal to the surface H^* at p .

By substituting the shorthand $g(u) = (\vec{\mathbf{N}}_p \cdot \vec{\mathbf{L}})$ into Equation 1 we get

$$I(\lambda_i, u) = L(\lambda_i)((\vec{\mathbf{N}}_p \cdot \vec{\mathbf{L}})S(\lambda_i, u) + k(u)) \quad (15)$$

where we have used notation consistent to that introduced in Section II-A.

To remove the restriction associated to a known illuminant direction, we make use of the fact that the light direction in the scene can be measured relative to an arbitrary reference $\mathbf{e} = [e_x, e_y, e_z]^T$. This is illustrated in the left-hand panel of Figure 3. In the figure, we show the reference axis, the surface normal and the viewer and light directions making use of a sample surface H^* . Here, we let the reference be such that the light direction becomes $\vec{\mathbf{L}} \equiv [0, 0, 1]^T$. That is, we set the reference of the surface H^* such that the estimated shading factor satisfies the relation

$$g(u)^* = \cos(\alpha, u) = (\vec{\mathbf{N}}_p \cdot \vec{\mathbf{L}}) \quad (16)$$

This contrasts with shape-from-shading approaches [31], where the aim of computation is the surface given by H^* . Here, H^* is a means for the computation of $g(u)^* = \cos(\alpha, u)$ through the shorthand $g(u)^* = (\vec{\mathbf{N}}_p \cdot \vec{\mathbf{L}})$. In shape-from-shading methods, the widely used reference is such that the camera axis, *i.e.* the viewer direction, corresponds to the depth of the object. As exemplified in the right-hand panel of Figure 3. In contrast, our choice of reference $\mathbf{e}$ has the effect of rotating the surface such that its depth can no longer be measured with respect to the camera plane. As a result of the use of a reference in the direction of the light source, the depth of the surface H^* with respect to the viewer will be rotated in the direction contrary to the viewer relative to the vector $\vec{\mathbf{L}}$. This rotation, however, does not affect the angle α between the normal $\vec{\mathbf{N}}_p$ and the light source direction $\vec{\mathbf{L}}$. Since our aim of computation is not the depth of the surface H^* but the angle α , we set $\vec{\mathbf{L}} \equiv [0, 0, 1]^T$ in our computations without any loss of generality and, hence, do not require the light source direction to be known a priori.

Fig. 3. Left-hand panel: The reflection geometry making use of an arbitrary reference e ; Middle panel: Shape-from-shading treatment where the viewpoint corresponds to the z-axis of the reference; Right-hand panel: The reflection geometry when the reference is given such that light source direction becomes $\vec{L} \equiv [0, 0, 1]^T$.

E. Extension to Trichromatic Imagery

Our method can also be applied to trichromatic images captured using RGB cameras. In trichromatic imagery, each of the three colour channels has its own spectral sensitivity function, *i.e.* colour matching function, $Q_c(\cdot)$ for each of the bands λ_i , where $c = \{R, G, B\}$. As a result, the c -colour response at pixel indexed u can be written in terms of the wavelength indexed radiance as follows

$$I_c(u) = \sum_{i=1}^N I(\lambda_i, u) Q_c(\lambda_i) \quad (17)$$

where we have written I_c to imply that the quantity above corresponds to the sum over the radiance $I(\lambda_i, u)$ at the pixel u as sensed by the camera with matching function $Q_c(\cdot)$.

When we apply the equation above to the dichromatic model, we get

$$\begin{aligned} I_c(u) &= g(u) \left(\sum_{i=1}^N S(\lambda_i, u) Q_c(\lambda_i) L(\lambda_i) \right) \\ &\quad + k(u) \left(\sum_{i=1}^N L(\lambda_i) Q_c(\lambda_i) \right) \\ &= g(u) B_c(u) + k(u) L_c \end{aligned} \quad (18)$$

where

$$B_c(u) = \sum_{i=1}^N S(\lambda_i, u) Q_c(\lambda_i) L(\lambda_i) \quad L_c = \sum_{i=1}^N L(\lambda_i) Q_c(\lambda_i)$$

From Equation 18, we can conclude that our method can be applied to trichromatic imagery so as to recover the parameters corresponding to the variables $g(u)$, $B_c(u)$, $k(u)$ and L_c above. Nonetheless, there are two important differences between these parameters and their wavelength-indexed counterparts. Firstly, L_c does not account for the true illuminant power spectrum rather than its product with the colour matching functions $Q_c(\cdot)$. Secondly, the variable $B_c(u)$ is a linear combination of the product of wavelength-indexed reflectance and illuminant weighted by the colour matching function. As a result, to further recover the analogue of the reflectance for trichromatic imagery, *i.e.* the albedo $\rho_c(u)$, for the c^{th} colour channel at pixel indexed u , we use the following

equation

$$\rho_c(u) = \frac{B_c(u)}{L_c} \quad (19)$$

III. IMPLEMENTATION

Now we turn our attention to the implementation of our method. To this end, we present the corresponding pseudocode in Algorithm 1. Note that, in the figure, we have altered slightly the computation sequence with respect to that presented in Section II. The pseudocode takes at input the initial estimate of the light power spectrum $L(\cdot)^*$, the image $\mathcal{I}$, the number of scales M used for the computation of the shapelets and the threshold ϵ . It delivers at output the estimated illuminant power spectrum L_{new}^* and the values for the photometric parameters in the arrays $\mathbf{G}$, $\mathbf{K}$ and $\mathbf{S}$ whose entry indexed u is given by $g(u)^*$, $k(u)^*$ and $S(\cdot, u)^*$, respectively.

The initial estimate of the illuminant power spectrum can be computed using a number of methods elsewhere in the literature [15], [32], [33]. In our implementation, we select a set of “dichromatic” image patches making use of the procedure in [15]. This is effected by subdividing the image into a lattice, where the pixels in each rectangular patch in the lattice are then fitted to a two-dimensional hyperplane, *i.e.* a dichromatic plane [16], spanned by the surface reflectance and the illuminant power spectrum. Here, we regard a patch to be dichromatic if it contains a small percentage of pixels whose radiance is in close accordance to their projection on the dichromatic plane. For a detailed description of the dichromatic plane computation, we refer the reader to [16]. With these image patches at hand, we apply continuum removal [34] so as to recover a pixel-wise estimate of the illuminant. The initial estimate of the illuminant is given by the median over the continuum removed spectra for all the selected patches. This follows the notion that the median is, effectively, a robust estimator [25].

Our pseudocode commences by computing the parameters s_k and C_{τ_k} for the shapelets as described in [24]. In Lines 3 and 4 we initialise S_u^* and $S(\lambda_j, u)$ such that, for the first iteration of the algorithm $S(\lambda_i, u) - S_u^* \equiv 1$. In Line 5, we initialise the specular coefficient $k(u)$. To this end, we

```

1: Input:  $L(\cdot)^*$ ,  $\mathcal{I}$ ,  $M$  and  $\epsilon$ 
2: Compute  $s_k$  and  $C_{\tau_k}$  for  $k = 1, 2, \dots, M$ 
3: Set  $S_u^* \equiv 0 \quad \forall u \in \mathcal{I}$ 
4: Set  $S(\lambda_j, u) \equiv 1 \quad \forall u \in \mathcal{I}$  and  $k = 1, 2, \dots, M$ 
5: Set  $k(u)^* = \min_i \left\{ \frac{I(\lambda_i, u)}{L(\lambda_i)} \right\}$ 
6: For  $u \in \mathcal{I}$ 
7:  $r_u^* = \frac{1}{N} \sum_{j=1}^N \frac{I(\lambda_j, u)}{L(\lambda_j)} - k(u)^*$ 
8:  $\mathbf{R}_{\lambda_i}(u) = \left| \frac{S(\lambda_i, u) - S_u^*}{\frac{I(\lambda_i, u)}{L(\lambda_i)} - k(u)^* - r_u^*} \right|$ 
9: For  $i = 1, 2, \dots, N$ 
10:  $H_{\lambda_i} = \sum_{k=1}^M \mathbf{R}_{(\lambda_i)}^* |\tan(s_k)| \cdot C_{\tau_k}$ 
11: Compute  $H^* = \frac{1}{N} \sum_{j=1}^N H_{\lambda_j}$ 
12: For  $u \in \mathcal{I}$ 
13:  $g(u)^* = (\nabla H^*|_u, [0, 0, 1]^T)$ 
14: For  $i = 1, 2, \dots, N$ 
15:  $S(\lambda_i, u)^* = \frac{1}{g(u)^*} \left( \frac{I(\lambda_i, u)}{L(\lambda_i)} - k(u)^* \right)$ 
16: Set  $L_{new}(\lambda_i)$  via Section II-C
17: Compute  $S_u^* = \frac{1}{N} \sum_{j=1}^N S(\lambda_j, u) \quad \forall u \in \mathcal{I}$ 
18: For  $u \in \mathcal{I}$ 
19: Compute  $k(u)^* = \min_{k(u) \geq 0} \left\{ \sum_{j=1}^N |I(\lambda_j, u) - L_{new}(\lambda_j)(g(u)^* S(\lambda_j, u)^* - k(u))|^2 \right\}$ 
20: If  $|L(\cdot)^* - L_{new}| \leq \epsilon$  then
21: Return  $L_{new}^*$ ,  $\mathbf{G}$ ,  $\mathbf{K}$  and  $\mathbf{S}$ 
22: Else
23:  $L(\cdot)^* \leftarrow L_{new}$ 
24: Go to Line 6

```

Algorithm 1: Pseudocode for our method

note that we can always bound $r(\lambda_i, u)$ to be between 0 and 1. We can give a physical interpretation to this assumption. The bounded nature of $r(\lambda_i, u)$ reflects the notion that the radiant power of light per unit of surface area, as sensed by the camera, can always be scaled so its maximum value is unity. This, together with the fact that, from Equation 2, the specular coefficient $k(u)$ is an additive constant to $r(\lambda_i, u)$, we initialise, in Line 5, $k(u)$ by setting it to the minimum value of the reflectance over the wavelength domain, *i.e.*

$$k(u)^* = \min_i \left\{ \frac{I(\lambda_i, u)}{L(\lambda_i)} \right\}$$

The initialisation of these parameters permits the computation of $\mathbf{R}_{\lambda_i}(u)$ in Line 8. In Lines 7 and 17, we compute r_u^* and S_u^* as presented in Section II-A. In Algorithm 1, we have used the fact that, making use of H^* , we can express the normal $\mathbf{N}_p = \nabla H^*|_{u \sim p}$, where the expression $u \sim p$ implies that the pixel indexed $u \in \mathcal{I}$ corresponds to the point $p \in H^*$. Note that the pseudocode is such that, if the change in the estimated illuminant power spectrum is less or equal to the threshold ϵ , it returns the photometric parameters. Otherwise

it goes to Line 6 so as to iterate accordingly. The optimisation in Line 19, which yields the updated specular coefficient has been effected using a constrained least squares based upon a Newton algorithm [35].

IV. EXPERIMENTS

In this section, we illustrate the effectiveness of our approach for the recovery of photometric parameters for a number of sample applications. We commence by briefly describing the data sets used throughout the section. We then provide a quantitative analysis on the recovery of the illuminant power spectrum and the shading factor for both imaging spectroscopy data and trichromatic images. We show results on skin recognition using the recovered reflectance and, for trichromatic imagery, Lambertian reflectance correction. We finish the section by presenting results for illuminant and surface reflectance replacement.

A. Data sets

For our experiments, we have used four hyperspectral and two trichromatic image data sets. The first two hyperspectral data sets are the Human Face Vis and Human Face NIR imagery. Both of these two data sets comprise images of 51 human subjects, each captured under one of 10 directional light sources with varying directions and spectral power. The light sources are divided into two rows. The first of these is placed above the camera system and the second one at the same height as the camera. The main direction of the lights is adjusted so as to point towards the center of the scene. The imagery has been acquired using a pair of cameras equipped with Liquid Crystal Tunable Filters which allow multi-spectral images to be resolved up to 10nm. The subjects in Human Face Vis Database are captured in the visible (430–720nm) wavelength ranges. The subjects in Human Face NIR Database are captured in the near infrared (650–990nm) wavelength ranges. Samples of these spectral images, rendered in pseudocolour, are shown in Figure 4.

The third hyperspectral data set comprises imagery depicting landscapes. This Landscape data set is composed of 14 scenes which vary from urban to wild fields containing sky, trees, vegetation and human made objects. All of these scenes are purely illuminated by sunlight. In every image, we have placed, at acquisition time, a white calibration target at the corner of the frame so as to capture the illuminant power spectrum ground truth. It is worth noting that, so as to avoid any unfair advantage of our method or the alternatives when processing the Landscape imagery, this calibration target has been cropped from the imagery at processing time. We show three sample images in the Landscape data set in Figure 5.

The fourth hyperspectral data set comprises imagery obtained using a camera equipped with an Acousto-Optic Tunable Filter. This filter allows tuning to be effected at steps of 1nm. The data set comprises images corresponding to five categories of toys and a set of s. Each toy sample was acquired over ten views by rotating the object in increments of 10° about its vertical axis whereas the imagery was captured only in two different views, heads and tails. In our

Fig. 4. Sample images of three subjects in the first two of our multi-spectral data sets. Note that, for our two first data sets (Human Face Vis and Human Face NIR), the 50 subjects used are the same, with the difference being the spectral range comprised by the imagery (visible for the Vis and near-infrared for the NIR database). These two spectral ranges have been acquired separately following [36].

Fig. 5. Sample images for three scenes in the Landscape data set.

data set, there are in total of 62 toys and 32 coins, which, over multiple viewpoints yielded 684 hyperspectral images. Each image is comprised of 51 bands for those wavelengths ranging from 550 to 1000 nm over 9nm steps. For purposes of photometric calibration, we have also captured an image of a white Spectralon calibration target so as to recover the power spectrum of the illuminant across the scene. In Figure 6, we show images for one object in each of the five categories of toys and two of the coins at 670nm. In the figure, each column corresponds to one of our six categories whereas each row shows the sample object over four different views.

As mentioned above, the other two data sets are trichromatic. The first of these is the Mondrian image data set described in [37]¹. This is a colour constancy data set containing 743 images of objects illuminated under 11 different illuminants whose ground truth is available. Figure 7 shows images for three sample objects in the data set captured under different illumination conditions. Lastly, we have used the MSRA salient object image data set [38] for our Lambertian reflectance correction results².

B. Illuminant Power Spectrum Recovery

We first turn our attention to the recovery of the illuminant power spectrum. To this end, we have used Human Face Vis, Human Face NIR, Coins and Toys, Landscape and the Mondrian data sets. To compute the accuracy of the recovered

illuminant with respect to the ground truth, we have employed two similarity measures. The first of these is the Spectral Information Divergence (SID) [39], which is an information theoretic measure which assesses similarity making use of the probabilistic discrepancy between the spectra under consideration. The SID is widely used in processing hyperspectral image data. We have also employed the Euclidean angle, which has also been widely used in the literature and is given by

$$d_{angle} = \cos^{-1} \left(\frac{\sum_{k=1}^N L(\lambda_k) L_{gt}(\lambda_k)}{\sqrt{(\sum_{k=1}^N L(\lambda_k)^2)} \sqrt{(\sum_{k=1}^N L_{gt}(\lambda_k)^2)}} \right) \quad (20)$$

Where $L(\lambda_i)$ is the estimated illumination and $L_{gt}(\lambda_i)$ is the ground truth illuminant.

Making use of the measures above, we present the illuminant estimation result yielded by our Shapelet Based Estimator (SBE), its simplification as presented in Section II-C1 and its extension to robust statistics given in Section II-C2. For the use of robust statistics, the regularisers used here are the Tukey bi-weight function [30], the Andrews M-estimator [40], the Cauchy regulariser [25] and the Huber kernel [25]. These are used as error penalty functions $g(\cdot)$ in Equation 14. In Figure 8, we show the qualitative shapes and formal expressions for the penalty functions used in our experiments. Note that these act as robust regularisers assigning a large penalty to those estimates far away from the sensed radiance and a non-linear penalty to those closer to the image values $I(\lambda_i, u)$. Note that, in the equations in Figure 8, we have used the variable κ as a bandwidth parameter. In all our experiments, we have set κ to be the equivalent of 10% of the dynamic range of the camera.

¹The data set is accessible at <http://www2.cs.sfu.ca/~colour/data/>

²The imagery can be downloaded from http://research.microsoft.com/en-us/um/people/jiansun/salientobject/salient_object.htm

Fig. 6. Sample images for 6 objects in our data set. We show the radiance corresponding to the band centered at $649nm$. Each column corresponds to one type of objects whereas the rows show different views.

Fig. 7. Three sample objects in the Mondrian data set [37]. Each column shows an object captured under different light settings.

Also, recall that the initial estimate of the illuminant can influence the results yielded by our SBE method. This is since, as discussed in Section III, our method requires initial estimates of the photometric parameters. As a result, we have compared our use of the median for the dichromatic patches with the results yielded by two alternative initialisation schemes. The first of these employs a uniform initialisation of the illuminant power spectrum, this is $L(\lambda_i) \equiv v$, where v is the mean over the set of maximum radiance values for each band in the image. The second scheme employs the maximal radiance in each band for any of the pixels in the set of dichromatic patches. Thus, in the former case we set the global maximum for each band as the initial illuminant estimation without using the homogeneous patches, whereas

for the latter scheme we employ an approximation of the light continuum [41].

In Table I, we show the illuminant recovery performance for the three methods using the three initialisation schemes mentioned earlier when processing the hyperspectral data sets and the Mondrian trichromatic one. In the table, the best performance per dataset have been typeset bold. Also, note that we provide error measurements for the near-infrared (NIR), *i.e.* $650nm - 990nm$, and the visible (VIS) bands, *i.e.* $430nm - 720nm$ of the human subject data set. This is to provide a consistent wavelength range to that used in NIR face identification systems [36] which employ NIR illumination as a means to recognition. From Table I, we observe that, in general, the NIR bands for the human subjects allow more

Fig. 8. Penalty functions $g(\eta(u))$ and their respective qualitative shapes for the robust regularisers used in our experiments.

accurate illuminant recovery than their visible counterparts. Further, the median of for the dichromatic patches outperforms the alternative in initialisation schemes regardless for all the data sets except the NIR face imagery.

Following the results in Table I, from now on, we compare the performance of our method and the alternatives when the median for the dichromatic patches is used for initialisation purposes. For all our experiments, we compare the results of our SBE method, its simplification presented in Section II-C1 (SBE Simplified) and the application of robust statistics as per Section II-C2 with a number of alternatives. The first of these is the method in [15]. We have chosen this alternative since its based upon a structural optimization approach which recovers the spectra of the illuminant, the surface reflectance and the shading and specular factors based upon the dichromatic model. Thus, it shares the multi-spectral nature of our method. The second alternative is given by the method in [16]. We also compare our method with the White Patch [42], Grey-World [43] and a Grey Edge [44] methods presented. For Grey Edge method, we only use the 1st order.

In order to provide a fair comparison between the methods under consideration, we employ the methodology in [44]. The methodology in [44] employs the mean, median, trimean, the mean for the best 25% and the worst 25% so as to better summarise non-Gaussian error distributions. This is particularly important since, as we will see later on, the error measurements across our data sets are not always normally distributed.

In Tables II-VI, we present the performance measures, yielded by both, Eculidean angle and SID, for the illumination

recovered using our method and the alternatives on the data sets under consideration. For the sake of clarity, in the tables, we have used bold font for the best Eculidean angle performance per data set for each of the error measurements. Note that, our method and its variants outperform the alternatives across the human face (VIS and NIR) data sets. They are also often in par with other algorithms when applied to the Landscapes and the Mondrian data sets.

The effectiveness of the simplified method can be attributed to the fact that it employs a subset of pixels whose image radiance is correlated to the illuminant for the purposes of recovering the light source power spectrum. Thus, estimating the illuminant by employing a subset of pixels which are less prone to noise corruption. This is as the assumption of a large specular coefficient indirectly implies that the illuminant is almost directly proportional to the image radiance for a pixel whose intensity is large and, hence, has a high SNR with respect to noise and the strength of light impinging on the sensor. It is worth noting in passing that this is somewhat akin to the white patch method. Nonetheless, our approach does not employ the brightest pixel in the image, but rather the ones such that $I(\lambda_i, u) \approx k(u)L(\lambda_i)$.

Note that this simplification does not affect the recovery of the reflectance, shading factor or specular coefficient. Also, note that the robust statistics application improves, in general, the results in the median and the mean of best 25%. Further, the robust statistics do improve the results for the overall mean on the Human Face NIR and Mondrian data set. This implies that the error distribution yielded by the robust statistics is less

TABLE I
AVERAGE PERFORMANCE, MEASURED USING THE EUCLIDEAN ANGLE, YIELDED BY OUR SBE ILLUMINANT RECOVERY METHOD AND ITS SIMPLIFIED ANALOGUE PRESENTED IN SECTION II-C1 (SBE SIMPLIFIED) ON OUR HYPERSPECTRAL DATA SETS AND THE MONDRIAN IMAGERY.

Data set	Initialisation Method	SBE	SBE (simplified)
Human Face Vis Database	Light continuum approx.	7.1123 ± 2.3802	4.8431 ± 4.2883
	Global radiance max.	7.8727 ± 3.0814	4.8985 ± 4.3139
	Dichromatic patch median	6.5048 ± 1.9252	3.1547 ± 4.2929
Human Face NIR Database	Light continuum approx.	3.3550 ± 0.9547	2.5624 ± 1.2623
	Global radiance max.	3.1537 ± 1.0777	2.5632 ± 1.1195
	Dichromatic patch median	3.0898 ± 1.0486	3.1291 ± 1.1475
Landscape Database	Light continuum approx.	19.0984 ± 8.6352	15.3212 ± 5.1458
	Global radiance max.	19.5287 ± 10.9002	13.4850 ± 5.6958
	Dichromatic patch median	19.3174 ± 7.0929	11.2873 ± 5.9089
Toys and coins	Light continuum approx.	11.7577 ± 4.7089	9.3796 ± 10.4101
	Global radiance max.	15.9850 ± 4.6914	9.8006 ± 10.6195
	Dichromatic patch median	12.0502 ± 4.1861	7.9062 ± 10.2213
Mondrian data set	Light continuum approx.	22.4693 ± 10.7677	6.2118 ± 9.3858
	Global radiance max.	16.0533 ± 10.6651	5.7872 ± 9.1088
	Dichromatic patch median	14.8967 ± 10.5720	5.7857 ± 9.1126

skewed and, hence, conforms better to a Gaussian distribution.

To finalise the section, in Table VII, we show the average processing timings (in seconds) per image in each of the data sets when applying our method, its simplified variant, robust statistics and the alternative [15]. For purposes of comparison, we have used the method in [15] since this is the only one of the alternatives which can recover the shading factor and the reflectance together with the specular coefficient and the illuminant power spectrum. This is as the other alternatives can just estimate the illuminant. Following the timings in the table, we note that our simplified method is approximately 10.21 times faster than the alternative. This is due to both, our much simpler optimisation approach and the use of the shapelets, which permits convergence in an average of 3.2 iterations. Moreover, the application of the robust estimators has a marginal increase in the processing time with respect to our SBE (simplified) method. This is particularly evident on the hyperspectral data sets, with an increase of 0.6 seconds, which is 1.01% of the computation time.

C. Shading Factor Recovery and Lambertian Reflectance Correction

We now show qualitative results on the recovery of the shading factor on sample human subjects from the data set used in the previous section. We also show sample results of Lambertian reflectance correction on trichromatic images taken from the MSRA data set [38]. This implies that, for the multi-spectral data, we are computing the shading factor $g(u)$ whereas for the trichromatic images the task involves removing specularities from shiny surfaces and for limb brightening effects to be compensated on rough objects.

In our experiments, we compare the method as presented in Section II against the results yielded by the alternative in [15] and the simplified variant of our method. Note that our two sets of experiments, *i.e.* those on multi-spectral images and the ones on trichromatic data, are quite related since the

shading factor here is the cosine of the angle between the surface normal and the light source direction. This is precisely Lambert's Law [31], which has been widely used in shape-from-shading approaches.

In Figure 9, we show the shading map, *i.e.* the value of $g(u)$ for every pixel in the detail for the imagery corresponding to three subjects in the data set. In the figure, the left-hand panel shows the pseudocolour image of the input multi-spectral imagery whereas the other three columns show the results yielded by the alternative in [15], our method and its simplified version, respectively. In Figure 10, we show a similar sequence for our Lambertian reflectance correction corresponding to three sample images taken from the MSRA data set [38]. For both figures, the shading factor is expected to appear as a diffuse, matte analogue of the input multispectral or trichromatic image.

From the images, we can observe that the secularities have been better suppressed by our method as compared to the alternative. Moreover, the shading factor recovered by our method is smoother and conveys a better shading cue. This is particularly evident in the forehead of the subject in the top row of Figure 9 and the image in the top row of Figure 10. Also note that the background for the shaded images delivered by the alternative, as shown in the second column of both figures, is somewhat noisy and distorted. Further, for the trichromatic imagery, the method in [15] has been duely affected by the textures in the MSRA images.

D. Skin Recognition

We now turn our attention to the invariance properties 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. For purposes of skin recognition, we

TABLE II
PERFORMANCE, IN BOTH EUCLIDAN ANGLE AND SID (IN PARENTHESIS), YIELDED BY OUR METHOD AND THE ALTERNATIVES WHEN APPLIED TO THE HUMAN FACE VIS DATA SET.

Method	Mean	Median	Trimean	Best 25 %	Worst - 25 %
SBE	6.5048 (0.0188)	5.9019 (0.0143)	6.1173 (0.0154)	4.3625 (0.0077)	9.5392 (0.0375)
SBE (simplified)	3.1547 (0.0065)	2.6684 (0.0027)	2.5816 (0.0027)	1.4149 (0.0008)	6.0543 (0.0197)
Andrews	3.2598 (0.0083)	2.4896 (0.0024)	2.4900 (0.0026)	1.4350 (0.0008)	6.5455 (0.0270)
Cauchy	3.2539 (0.0083)	2.4871 (0.0024)	2.4810 (0.0026)	1.4370 (0.0008)	6.5386 (0.0269)
Huber	3.2617 (0.0083)	2.5147 (0.0024)	2.4910 (0.0026)	1.4362 (0.0008)	6.5475 (0.0269)
Tukey	3.2597 (0.0083)	2.4896 (0.0024)	2.4900 (0.0026)	1.4351 (0.0008)	6.5449 (0.0270)
Method in [15]	6.2868 (0.0277)	5.5591 (0.0186)	6.1183 (0.0245)	3.0727 (0.0062)	10.2407 (0.0587)
Method in [16]	5.9688 (0.0249)	4.9071 (0.0135)	5.5439 (0.0201)	3.2769 (0.0068)	9.8436 (0.0559)
Grey World	11.5688 (0.0629)	11.5026 (0.0603)	11.5347 (0.0612)	9.5857 (0.0414)	13.6424 (0.0877)
White Patch	5.0173 (0.0146)	4.0269 (0.0061)	4.4411 (0.0092)	1.7788 (0.0013)	9.4410 (0.0398)
Grey Edge	7.0479 (0.0265)	6.8091 (0.0194)	6.7336 (0.0202)	3.3973 (0.0044)	11.2162 (0.0622)

TABLE III
PERFORMANCE, IN BOTH EUCLIDAN ANGLE AND SID (IN PARENTHESIS), YIELDED BY OUR METHOD AND THE ALTERNATIVES WHEN APPLIED TO THE HUMAN FACE NIR DATA SET.

Method	Mean	Median	Trimean	Best 25 %	Worst - 25 %
SBE	3.0898 (0.0039)	2.9558 (0.0032)	3.0289 (0.0035)	2.0577 (0.0016)	4.3342 (0.0072)
SBE (simplified)	3.1291 (0.0089)	1.8283 (0.0012)	2.0243 (0.0021)	0.7820 (0.0002)	7.4315 (0.0301)
Andrews	1.8867 (0.0021)	1.6584 (0.0009)	1.7640 (0.0016)	0.7306 (0.0002)	3.4097 (0.0058)
Cauchy	1.8917 (0.0021)	1.6775 (0.0009)	1.7522 (0.0015)	0.7291 (0.0002)	3.3965 (0.0057)
Huber	1.9102 (0.0022)	1.6921 (0.0011)	1.8114 (0.0017)	0.7309 (0.0002)	3.4188 (0.0058)
Tukey	1.8866 (0.0021)	1.6587 (0.0009)	1.7645 (0.0016)	0.7306 (0.0002)	3.4091 (0.0058)
Method in [15]	6.6822 (0.0207)	6.3695 (0.0160)	6.4752 (0.0178)	3.7635 (0.0053)	10.0058 (0.0425)
Method in [16]	8.7401 (0.0352)	9.2414 (0.0353)	9.2357 (0.0356)	4.9264 (0.0106)	11.7016 (0.0595)
Grey World	4.1531 (0.0078)	3.6962 (0.0053)	3.8341 (0.0059)	2.8260 (0.0031)	6.0701 (0.0161)
White Patch	2.8205 (0.0054)	2.2525 (0.0024)	2.2746 (0.0026)	0.9411 (0.0003)	5.7469 (0.0164)
Grey Edge	2.2417 (0.0022)	2.0003 (0.0017)	2.0951 (0.0018)	1.5053 (0.0009)	3.3173 (0.0043)

select skin and non-skin regions from one of the images for each subject in our human subject data set. We use this

single image per subject, taken under one of the illuminants, and store the normalised reflectance spectra as training data.

TABLE IV
PERFORMANCE, IN BOTH EUCLIDAN ANGLE AND SID (IN PARENTHESIS), YIELDED BY OUR METHOD AND THE ALTERNATIVES WHEN APPLIED TO THE LANDSCAPE DATA SET.

Method	Mean	Median	Trimean	Best 25 %	Worst - 25 %
SBE	19.3174 (0.0197)	16.1788 (0.0019)	17.1320 (0.0021)	10.8625 (0.0010)	29.0911 (0.0722)
SBE (simplified)	11.2873 (0.0101)	7.8968 (0.0019)	9.0438 (0.0021)	4.3386 (0.0010)	19.3380 (0.0343)
Andrews	12.7993 (0.0120)	12.2962 (0.0019)	11.4875 (0.0021)	4.2221 (0.0010)	20.4111 (0.0419)
Cauchy	12.8096 (0.0120)	12.2861 (0.0019)	11.4544 (0.0021)	4.2556 (0.0010)	20.3243 (0.0419)
Huber	12.8188 (0.0120)	12.2089 (0.0019)	11.4109 (0.0021)	4.2753 (0.0010)	20.2923 (0.0419)
Tukey	12.7996 (0.0120)	12.2963 (0.0019)	11.4870 (0.0021)	4.2238 (0.0010)	20.4118 (0.0419)
Method in [15]	10.9602 (0.0085)	10.0906 (0.0019)	10.0430 (0.0021)	6.9585 (0.0010)	16.0119 (0.0284)
Method in [16]	14.5406 (0.0131)	13.4594 (0.0019)	13.2821 (0.0021)	8.2529 (0.0010)	21.5987 (0.0462)
Grey World	20.6015 (0.0217)	22.8744 (0.0019)	20.8913 (0.0021)	10.5402 (0.0010)	30.5261 (0.0799)
White Patch	10.0340 (0.0080)	11.1583 (0.0019)	10.2222 (0.0021)	4.0532 (0.0010)	15.8473 (0.0264)
Grey Edge	14.1116 (0.0127)	12.6478 (0.0019)	13.6811 (0.0021)	6.8654 (0.0010)	22.9733 (0.0447)

Fig. 9. Shading factor recovery results. First column: pseudo-color image for sample input imagery; Second column: results yielded by [15]; Third and fourth columns: results for our SBE method and its simplified variant, respectively.

Subsequently, we train a Support Vector Machine (SVM) classifier with a polynomial kernel of degree 2 on the training

TABLE V
PERFORMANCE, IN BOTH EUCLIDAN ANGLE AND SID (IN PARENTHESIS), YIELDED BY OUR METHOD AND THE ALTERNATIVES WHEN APPLIED TO THE TOYS AND COINS DATA SET.

Method	Mean	Median	Trimean	Best 25 %	Worst - 25 %
SBE	12.0502 (0.1158)	13.6316 (0.1260)	12.5501 (0.1116)	4.3165 (0.0163)	18.2270 (0.2105)
SBE (simplified)	7.9062 (0.0587)	4.8601 (0.0276)	5.2016 (0.0269)	3.4501 (0.0123)	17.7344 (0.1689)
Andrews	8.0873 (0.0671)	4.4454 (0.0189)	4.7704 (0.0213)	2.9906 (0.0065)	19.9429 (0.2194)
Cauchy	7.9747 (0.0658)	4.4025 (0.0185)	4.7333 (0.0216)	2.8468 (0.0061)	19.7823 (0.2158)
Huber	8.0114 (0.0662)	4.3813 (0.0189)	4.7336 (0.0218)	2.8983 (0.0062)	19.8272 (0.2165)
Tukey	8.0853 (0.0671)	4.4454 (0.0189)	4.7671 (0.0213)	2.9879 (0.0065)	19.9388 (0.2193)
Method in [15]	3.4282 (0.0069)	3.2775 (0.0045)	3.2580 (0.0046)	1.9365 (0.0020)	5.3213 (0.0160)
Method in [16]	53.2769 (0.0670)	75.8116 (0.0074)	58.5687 (0.0102)	3.4888 (0.0020)	78.4717 (0.2386)
Grey World	20.2554 (0.2581)	22.6857 (0.3006)	21.3653 (0.2775)	11.2977 (0.0981)	25.7157 (0.3634)
White Patch	9.7244 (0.0841)	5.1664 (0.0240)	7.0590 (0.0494)	3.3204 (0.0117)	21.4815 (0.2360)
Grey Edge	9.8812 (0.0842)	6.4993 (0.0386)	8.2428 (0.0644)	2.6969 (0.0058)	21.2989 (0.2246)

TABLE VI
PERFORMANCE, IN BOTH EUCLIDAN ANGLE AND SID (IN PARENTHESIS), YIELDED BY OUR METHOD AND THE ALTERNATIVES WHEN APPLIED TO THE MONDRAIN DATA SET.

Method	Mean	Median	Trimean	Best 25 %	Worst - 25 %
SBE	14.8967 (0.1230)	12.3843 (0.0486)	13.0833 (0.0623)	4.3854 (0.0106)	29.6882 (0.3620)
SBE (simplified)	5.7857 (0.0455)	2.6024 (0.0027)	2.9151 (0.0038)	0.7123 (0.0002)	17.0204 (0.1747)
Andrews	5.6729 (0.0468)	2.4052 (0.0023)	2.6890 (0.0035)	0.4797 (0.0001)	17.2213 (0.1812)
Cauchy	5.6169 (0.0443)	2.5118 (0.0024)	2.8415 (0.0039)	0.5697 (0.0002)	16.6817 (0.1706)
Huber	5.6201 (0.0450)	2.4426 (0.0024)	2.8068 (0.0036)	0.5466 (0.0001)	16.7923 (0.1735)
Tukey	5.6726 (0.0468)	2.4013 (0.0023)	2.6888 (0.0035)	0.4802 (0.0001)	17.2194 (0.1812)
Method in [15]	11.4822 (0.0923)	9.2743 (0.0335)	9.4839 (0.0409)	2.5933 (0.0031)	24.1779 (0.2894)
Method in [16]	9.1814 (0.0913)	3.7366 (0.0059)	5.9405 (0.0255)	0.3594 (0.0001)	25.9029 (0.3292)
Grey World	7.3296 (0.0404)	4.1031 (0.0071)	5.5799 (0.0175)	0.6015 (0.0002)	18.6296 (0.1316)
White Patch	6.8212 (0.0273)	5.4641 (0.0123)	5.5606 (0.0137)	1.8248 (0.0015)	14.3614 (0.0818)
Grey Edge	4.6137 (0.0141)	3.6339 (0.0048)	3.7481 (0.0071)	0.3642 (0.0001)	10.7983 (0.0446)

set. The resulting SVM model is applied to classify 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 the robustness of the

TABLE VII
AVERAGE TIMING (IN SECONDS) PER IMAGE IN EACH DATA SET FOR OUR SBE METHOD, ITS SIMPLIFIED VERSION (SBE SIMPLIFIED), THE APPLICATION OF ROBUST REGULARISERS AND THE ALTERNATIVE IN [15].

Method	Human Face Vis	Human Face NIR	Landscape	Toys and Coins	Mondrain
SBE	158.55 ± 1.86	180.99 ± 2.68	161.32 ± 4.32	128.35 ± 1.49	61.26 ± 0.48
SBE (simplified)	59.20 ± 0.46	64.87 ± 0.78	60.39 ± 1.48	43.21 ± 0.20	61.52 ± 6.10
Andrews	59.51 ± 0.61	65.32 ± 0.94	60.55 ± 1.38	43.82 ± 0.30	70.12 ± 0.39
Cauchy	59.46 ± 0.54	65.17 ± 0.70	59.90 ± 0.89	43.68 ± 0.22	70.30 ± 2.66
Huber	59.46 ± 0.58	65.23 ± 1.21	59.92 ± 0.70	43.74 ± 0.25	70.11 ± 0.42
Tukey	59.43 ± 0.56	65.19 ± 0.74	60.28 ± 1.06	43.72 ± 0.22	70.14 ± 0.47
Method in [15]	597.06 ± 5.49	669.25 ± 4.99	502.32 ± 5.08	499.56 ± 6.87	683.43 ± 1.70

Fig. 10. Lambertian reflectance correction results. First column: sample trichromatic input imagery; Second column: results yielded by [15]; Third and fourth columns: results for our SBE method and its simplified variant, respectively.

TABLE VIII
SKIN SEGMENTATION ACCURACY YIELDED BY THE REFLECTANCE
RECOVERED USING OUR SBE METHOD, THE ALTERNATIVE IN [15] AND
THE RAW IMAGE RADIANCE

	CR	FDR
SBE	94.87% $\pm$ 6.94%	2.05% $\pm$ 1.87%
SBE (simplified)	97.72% $\pm$ 1.84%	2.18% $\pm$ 1.27%
Alternative in [15]	86.61% $\pm$ 15.88%	8.12% $\pm$ 8.74%
Raw radiance	71.23% $\pm$ 9.17%	12.48% $\pm$ 11.64%

illuminant spectrum recovered by our algorithm at training time and the normalised reflectance yielded by the method at testing.

In the left-hand column of Figure 11, we show a sample band for the images used for training corresponding to two subjects in the data set. The second column shows one of the testing images of the same subjects under a diverse illuminant condition. The skin probability maps yielded by the SVM applied to the reflectance recovered by the alternative in [15] are shown in the third column, whereas those yielded using the reflectance computed by our method and its simplified version are shown in the two right-most columns. In these skin probability maps, lighter pixels are more likely to be skin.

Figure 11 provides a proof of concept that the recovered skin reflectance spectra are invariant to illuminant variations and light source direction changes. Indeed, the skin segmentation maps for our method 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, the alternative has produced much more false negatives in skin areas and false positives in other materials.

We provide a more quantitative analysis in Table VIII. In the table, we show the performance of the above skin segmentation scheme in terms of the classification rate (*CR*) and false detection rate (*FDR*) over 5-fold cross validation. In the table, we also show the accuracy when the raw radiance is used as an alternative to the reflectance recovered by our SBE method or the alternative in [15]. The classification rate is the overall percentage of skin pixels classified correctly. The false detection rate is the percentage of non-skin pixels incorrectly classified as skin. The table shows the *CR* and *FDR* over all the subjects in the data set when testing is done on the image illuminated by the frontal light source. As expected, the reflectance recovered by our method achieves the highest skin recognition rates.

E. Reflectance and Illuminant Replacement

Finally, we illustrate the utility of our method for tasks relevant to image editing. In this section, we show how the photometric parameters recovered from imaging spectroscopy can be used to replace the reflectance of an object in the scene or change the illuminant power spectrum. This benefits from the ability of combining spatial and compositional information so as to reconstruct an image using image formation process. To replace the reflectance of an object or change the illuminant in the scene, we commence by recovering the photometric

parameters using our method. Once the reflectance, illuminant power spectrum, shading factor and specular coefficients are at hand, we recover the new image by evaluating the model in Equation 2.

In this set of experiments, we have used a spectrometer to obtain the reflectance of some naturally occurring materials, such as leaves, human skin and terracotta. These reflectance spectra are used as an alternative to that of the object under consideration so as to recover the edited image. For the illuminants, we have used a spectrometer equipped with an integrating sphere to acquire the power spectrum of light sources commonly found in the real world, such as sunlight, tungsten illuminants, Xenon-Krypton lights, etc.

We first illustrate how the skin recognition results shown in the previous section can be used for skin reflectance substitution. In Figure 12, we show three pseudocolour images corresponding to imagery from the human subjects data set. For these subjects, we have performed skin recognition and replaced the skin spectra with an alternative skin reflectance acquired with the spectrometer.

In a similar setting, we have performed reflectance recognition on the clothing of the three subjects. Here, we have used the reflectance extracted by our method and applied the object material recognition algorithm in [45] to the cloth under consideration so as to select the pixels whose reflectance is to be replaced. The results of these cloth reflectance replacement operations are shown in the second row of Figure 13. Note that, for the second subject, the creases are much more evident with the alternative reflectance. We would like to emphasise that no further editing or retouching has been effected. The reason for the creases to appear more evident when an alternative reflectance is used resides in the fact that these are indeed visible in some bands on the multi-spectral imagery. These bands have a marginal effect when the pseudocolour image is produced due to the dark (black) colour of the t-shirt. Nonetheless, with the alternative reflectance, whose colour is beige, these creases appear on the pseudocolour image. This is consistent with the notion that our method can deliver a separate output for the shape, the reflectance and other photometric parameters and, hence, the reflectance substitution does not affect the creases in the t-shirt.

Finally, we show results for illuminant power spectrum replacement on the three subjects used in our experiments. In Figure 14 we show, in the second row the pseudocolour for the images in the top row when the spectrum of the illuminant is replaced by that corresponding to a strip light. In the third row, we show the results of the cloth, skin and illuminant power spectrum replacement on the imagery. Note that the temperature of the imagery changes, tending to the greenish hue commonly induced by neon tubes.

V. CONCLUSIONS

In this paper, we have presented an efficient method Shapelet Based Estimator for the recovery of the dichromatic photometric parameters to tackle the photometric invariance problem from a single multispectral image. The recovery process is based upon shapelets and permits the use of an efficient

Fig. 11. Left-hand column: sample training images for two subjects in our data set; Second column: sample testing multi-spectral image; Third column: probability skin maps yielded by the reflectance extracted using the method in [15]; Fourth and fifth columns: skin probability maps recovered making use of the reflectance computed using our method and its simplified version, respectively.

Fig. 12. Skin reflectance substitution. Top panel: Pseudocolour imagery for the input images; Bottom panel: Result of replacing the skin reflectance of the subject for an alternative in our database.

Fig. 13. Top panel: Pseudocolour imagery for the input images; Bottom panel: Result of replacing the reflectance of the subject's clothing for an alternative in our database.

three-step process which can be solved via a constrained least-squared optimisation. We have also presented a simplified version of the method which employs highly specular pixels so as to recover the illuminant power spectrum. This simplification has proved to be effective in our experiments and provides a means to avoiding the use of least-squares optimisation. We have explored the use of robust statistics and performed a number of experiments on illuminant power spectrum recovery, shading factor computation and skin recognition. We have also compared our method against an alternative and shown simple examples of re-illumination and material substitution which illustrate the potential of the method for digital media generation, image editing and retouching.

VI. ACKNOWLEDGMENT

The work presented here was done when Jun Zhou was with the National ICT Australia (NICTA).

REFERENCES

- [1] S. J. Kim, S. Zhuo, F. Deng, C. W. Fu, and M. Brown, "Interactive visualization of hyperspectral images of historical documents," *IEEE Transactions on Visualization and Graphics*, vol. 16, no. 6, pp. 1441–1448, 2010.
- [2] G. Brelstaff and A. Blake, "Detecting specular reflection using lambertian constraints," in *Int. Conference on Computer Vision*, 1988, pp. 297–302.
- [3] S. Nayar and R. Bolle, "Reflectance based object recognition," *International Journal of Computer Vision*, vol. 17, no. 3, pp. 219–240, 1996.
- [4] S. Lin and H. Shum, "Separation of diffuse and specular reflection in color images," in *IEEE Conf. on Computer Vision and Pattern Recognition*, 2001.
- [5] G. Klinker, S. Shafer, and T. Kanade, "A physical approach to color image understanding," *Intl. Journal of Computer Vision*, vol. 4, no. 1, pp. 7–38, 1990.
- [6] H. Ragheb and E. R. Hancock, "A probabilistic framework for specular shape-from-shading," *Pattern Recognition*, vol. 36, no. 2, pp. 407–427, 2003.
- [7] B. van Ginneken, M. Starvridi, and J. J. Koenderink, "Diffuse and specular reflectance from rough surfaces," *Applied Optics*, vol. 37, no. 1, pp. 130–139, 1998.
- [8] T. K. Leung and J. Malik, "Recognizing surfaces using three-dimensional textons," in *Proc. of the IEEE Int. Conf. on Compl. Vis.*, 1999, pp. 1010–1017.
- [9] S. K. Nayar, K. Ikeuchi, and T. Kanade, "Surface reflection: Physical and geometrical perspectives," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 13, no. 7, pp. 611–634, 1991.
- [10] J. E. Harvey, C. L. Vernold, A. Krywonos, and P. L. Thompson, "Diffracted radiance: A fundamental quantity in a non-paraxial scalar diffraction theory," *Applied Optics*, vol. 38, no. 31, pp. 6469–6481, 1999.
- [11] G. Healey and D. Slater, "Invariant recognition in hyperspectral images," in *IEEE Conf. on Computer Vision and Pattern Recognition*, 1999, pp. 1438–1043.
- [12] P. H. Suen and G. Healey, "Invariant mixture recognition in hyperspectral images," in *Int. Conference on Computer Vision*, 2001, pp. 262–267.
- [13] H. M. G. Stokman and T. Gevers, "Detection and classification of hyperspectral edges," in *British Machine Vision Conference*, 1999.
- [14] E. Angelopoulou, "Objective colour from multispectral imaging," in *European Conf. on Computer Vision*, 2000, pp. 359–374.
- [15] C. P. Huynh and A. Robles-Kelly, "A solution of the dichromatic

Fig. 14. Top panel: Pseudocolour imagery for the input images; Second row: Result of replacing the illuminant in the scene. Bottom row: Result of replacing the skin spectra, cloth reflectance and scene illuminant for alternatives in our database.

- model for multispectral photometric invariance,” *International Journal of Computer Vision*, vol. 90, no. 1, pp. 1–27, 2010.
- [16] S. A. Shafer, “Using color to separate reflection components,” *Color Research and Applications*, vol. 10, no. 4, pp. 210–218, 1985.
- [17] K. Ikeuchi and B. Horn, “Numerical shape from shading and occluding boundaries,” *Artificial Intelligence*, vol. 17, no. 1-3, pp. 141–184, 1981.
- [18] B. K. P. Horn and M. J. Brooks, “The variational approach to shape from shading,” *CVGIP*, vol. 33, no. 2, pp. 174–208, 1986.
- [19] Q. Zheng and R. Chellappa, “Estimation of illuminant direction, albedo, and shape from shading,” *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 13, no. 7, pp. 680–702, 1991.
- [20] P. L. Worthington and E. R. Hancock, “New constraints on data-closeness and needle map consistency for shape-from-shading,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 21, no. 12, pp. 1250–1267, 1999.
- [21] P. Dupuis and J. Oliensis, “Direct method for reconstructing shape from shading,” in *Proc. of the IEEE Conf. on Comp. Vision and Pattern Recognition*, 1992, pp. 453–458.
- [22] E. Prados and O. Faugeras, “Perspective shape from shading and viscosity solutions,” in *IEEE Int. Conf. on Computer Vision*, 2003, pp. II:826–831.
- [23] R. Kimmel and A. M. Bruckstein, “Tracking level sets by level sets: a method for solving the shape from shading problem,” *Computer vision and Image Understanding*, vol. 62, no. 2, pp. 47–48, July 1995.
- [24] P. Kovesi, “Shapelets correlated with surface normals produce surfaces,” in *Int. Conference on Computer Vision*, 2005, pp. 994–1001.
- [25] P. Huber, *Robust Statistics*. Wiley, 1981.
- [26] B. K. P. Horn, “Shape from shading: A method for obtaining the shape

of a smooth opaque object from one view,” in *MIT AI-TR*, 1970.

- [27] A. Refregier, “Shapelets - I. A method for image analysis,” *Monthly Notice of the Royal Astronomical Society*, vol. 338, no. 1, pp. 35–47, 2003.
- [28] E. P. Simoncelli, W. T. Freeman, E. H. Adelson, and D. J. Heeger, “Shiftable multiscale transforms,” *IEEE Trans. Inf. Theor.*, vol. 38, no. 2, pp. 587–607, Sep. 2006.
- [29] S. G. Mallat, *A Wavelet Tour of Signal Processing*. Academic Press, 1999.
- [30] J. W. Tukey, *Exploratory Data Analysis*. Addison-Wesley, 1977.
- [31] B. K. P. Horn and B. Schunk, “Determining optical flow,” *Artificial Intelligence*, vol. 17, pp. 185–204, 1981.
- [32] G. D. Finlayson and G. Schaefer, “Convex and non-convex illuminant constraints for dichromatic colour constancy,” in *IEEE Conf. on Computer Vision and Pattern Recognition*, 2001, pp. 1:598–604.
- [33] G. D. Finlayson, C. Fredembach, and M. S. Drew, “Detecting illumination in images,” in *IEEE International Conference on Computer Vision*, 2007.
- [34] Z. Fu, A. Robles-Kelly, T. Caelli, and R. Tan, “On Automatic Absorption Detection for Imaging Spectroscopy: A Comparative Study,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 45, pp. 3827–3844, 2007.
- [35] T. F. Coleman and Y. Li, “On the convergence of interior-reflective newton methods for nonlinear minimization subject to bounds,” *Mathematical Programming*, vol. 67, pp. 189–224, 1994.
- [36] S. Z. Li, R. Chu, S. Liao, and L. Zhang, “Illumination invariant face recognition using near-infrared images,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 29, pp. 627–639, 2007.
- [37] K. Barnard, L. Martin, B. Funt, and A. Coath, “A data set for colour research,” *Color Research and Application*, vol. 27, no. 3, pp. 147–151, 2002.
- [38] T. Liu, J. Sun, N.-N. Zheng, X. Tang, and H.-Y. Shum, “Learning to detect a salient object,” in *Computer Vision and Pattern Recognition*, 2007, pp. 1–8.
- [39] C.-I. Chang, “An information-theoretic approach to spectral variability, similarity, and discrimination for hyperspectral image analysis,” *IEEE Trans. Inf. Theor.*, vol. 46, no. 5, pp. 1927–1932, Sep. 2000.
- [40] S. Andrews, I. Tsochantaris, and T. Hofmann, “Support vector machines for multiple-instance learning,” in *Advances in Neural Information Processing Systems*, 2003, pp. 561–568.
- [41] R. N. Clark, A. J. Gallagher, and G. Swayze, “Material absorption band depth mapping of imaging spectrometer data using a complete band shape least-squares fit with library reference spectra,” in *Proceedings of the Second Airborne Visible/Infrared Imaging Spectrometer Workshop*, 1990, pp. 176–186.
- [42] E. Land, “The retinex theory of color vision,” *Scientific American*, vol. 237, no. 6, pp. 108–128, Sep. 1977.
- [43] G. Buchsbaum, “A spatial processor model for object colour perception,” *Journal of the Franklin Institute*, vol. 310, no. 1, pp. 1–26, 1980.
- [44] J. van de Weijer, T. Gevers, and A. Gijsenij, “Edge-based color constancy,” *IEEE Transactions on Image Processing*, vol. 16, no. 9, pp. 2207–2214, 2007.
- [45] Z. Fu and A. Robles-Kelly, “Discriminant absorption-feature learning for material classification,” *Geoscience and Remote Sensing, IEEE Transactions on*, vol. 49, no. 5, pp. 1536–1556, 2011.

Antonio Robles-Kelly Antonio Robles-Kelly received the B.Eng. degree in electronics and telecommunications in 1998. In 2001 he received the William Gibbs/Plessey Award to the best research proposal and, in 2003, the Ph.D. degree in computer science from the University of York. He remained at the University of York until December 2004 as a Research Associate under the Mathematics for IT (MathFIT) initiative of the Engineering and Physical Sciences Research Council (EPSRC). In 2005, he took a research scientist appointment with NICTA.

After working on surveillance systems with query capabilities, in 2006 he was appointed project leader of the Spectral Imaging project and later promoted to Principal Researcher and Research Leader. He is also an adjunct Associate Professor at the ANU. He has been a Postdoctoral Research Fellow of the ARC and currently serves as an Associate Editor of the IET Computer Vision and Pattern Recognition journals. His research interests are in the areas of computer vision, statistical pattern recognition, computer vision and image processing. Dr. Robles-Kelly has been a Chair, co-chair, and technical committee member of several mainstream computer-vision and pattern-recognition conferences.

Jun Zhou Jun Zhou (M06, SM12) received the B.S. degree in computer science and the B.E. degree in international business from Nanjing University of Science and Technology, Nanjing, China, in 1996 and 1998, respectively, the M.S. degree in computer science from Concordia University, Montreal, Canada, in 2002, and the Ph.D. degree from the University of Alberta, Edmonton, Canada, in 2006. He joined the School of Information of Communication Technology at Griffith University, Nathan, Australia, as a Lecturer in June 2012. Previously, he had been a

Research Fellow in the Research School of Computer Science in the Australian National University, Canberra, Australia, and a Researcher in the Canberra Research Laboratory, NICTA, Australia. His research interests include pattern recognition, computer vision, and machine learning with human in the loop, with their applications to spectral imaging and environmental informatics.

Lin Gu Lin Gu received the B.Eng. degree from Shanghai University, China in 2009. Currently, he is working towards his PhD degree at the Australian National University and the National ICT Australia(NICTA). His research interests include image spectroscopy, shape recovery, colour science and pattern recognition.