

A NURBS-Based Spectral Reflectance Descriptor with Applications in Computer Vision and Pattern Recognition

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

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

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

Abstract

In this paper, we present a surface reflectance descriptor based on the control points resulting from the interpolation of Non-Uniform Rational B-Spline (NURBS) curves to multispectral reflectance data. The interpolation is based upon a knot removal scheme in the parameter domain. Thus, we exploit the local support of NURBS so as to recover a compact descriptor robust to noise and local perturbation of the spectra. We demonstrate the utility of our NURBS-based descriptor for material identification. To this end, we perform skin spectra recognition making use of a Support Vector Machine classifier. We also provide results on hyperspectral imagery and elaborate on the preprocessing step for skin segmentation. We compare our results with those obtained using an alternative descriptor.

1. Introduction

The development of image sensor technology has made it possible to capture image data in hundreds of bands covering a broad spectrum of wavelengths. In contrast with trichromatic sensors, multispectral and hyperspectral sensing devices can acquire wavelength-indexed reflectance and radiance data in thousands of bands across a broad spectrum.

Thus, hyperspectral cameras, filter-sets and multiplexed light provide a means to recover an information-rich representation of the reflectance and radiance data. In [18], a method to capture hyperspectral imaging through multiplexed light is presented. Angelopoulou *et al.* [2] have used spectrophotometric data to model skin colour. In [1], filters are used to recover multispectral images for highlight removal.

Moreover, the spectral response of the material under study over a number of wavelengths permits the recovery

of the reflectance spectra as a feature for material identification and recognition. Due to the high dimensional nature of the spectral data, many classical algorithms in pattern recognition and machine learning have been naturally borrowed and adapted so as to perform feature extraction and classification [11]. Techniques such as Principle Component Analysis (PCA) [10], Linear Discriminant Analysis (LDA)[7], Projection Pursuit [9] and their kernel versions [6] treat raw spectra as input vectors in a higher-dimensional space, where the dimensionality is given by the number of bands. The task is then viewed as a mapping onto the feature space, often with reduced dimensionality, via the optimisation of a given cost function which leads to statistically optimal solutions to classification.

An alternative to the use of raw spectra as a means to classification and recognition is the use of a reflectance descriptor, robust to changes in illumination, noise and surface geometry. Recognition and identification methods have benefited from the recovery of properties which are invariant to geometric and photometric effects. Nayar and Bolle [16] have proposed a method of object recognition based on the reflectance ratio between object regions. Slater and Healey [22] used a set of Gaussian filters to derive moment invariants for recognition. Jacobs *et al.* [8] have developed a method for comparing images under variable illumination. Lin and Lee [14] have used chromaticity distributions for 3D object recognition. Lenz *et al.* [13] have used perspective projections to separate intensity from chromaticity so as to recover a three-dimensional colour descriptors.

In contrast with other approaches elsewhere in the literature, here we focus on the development of a free-form representation for arbitrary spectral data. This provides a compact representation of wavelength dependent information. To this end, we profit from the properties of NURBS curves [20]. NURBS are widely used in Computer-Aided Design and graphics but do not appear to have been used for the representation of spectra. This is somewhat surprising since NURBS allow the representation of analytical and free-form shapes using a set of control points with local support, i.e. the control points have a bounded effect on the

*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.

span of the curve. As a result, the descriptor presented here is robust to perturbations in the spectra. Moreover, the descriptor permits both, interpolation via a set of numerically stable algorithms and the representation of spectra of dissimilar lengths captured by different acquisition methods.

2. NURBS-Based Descriptor

As mentioned earlier, our NURBS-based descriptor can cope with densely sampled reflectance spectra, which could potentially consist of hundreds of data points over the visible spectrum. For purposes of classification, long feature vectors has been known to degrade performance since they incur computational cost and learning-theoretic limitations. Thus, it is desirable that our descriptor has the most discriminative power with the lowest possible dimensionality.

In this section we present a method of recovering a compact descriptor for reflectance spectra. Nonetheless, in theory, this descriptor may be of arbitrary length, we show how the descriptor may be rendered compact by means of knot removal. To do this, we commence with an introduction to NURBS curves and relate them to an interpolation step on the spectrum. We then formulate the reflectance descriptor through performing knot removal by minimising a cost function in the parameter domain of the spline curve.

2.1. Descriptor Formulation

To commence, recall a B-spline is a function that has support with respect to degree, smoothness and domain partition. These properties make B-spline a flexible tool for fitting arbitrary shapes and data. The smoothness property makes the interpolating curve robust to noise. The local support property permits the modification of the curve over a given wavelength range while keeping the rest of the spline unaffected.

First, we require some formalism. Since spectra are a function of wavelength λ , we restrict our analysis to the two-dimensional case. A p -degree B-Spline curve \mathcal{C} in \mathbb{R}^2 composed of n segments is a function in the parameter domain \mathcal{U} of the univariate variable t given by the linear combination $\mathcal{C}(t) = \sum_{i=0}^n N_{i,p}(t)P_i$, where $P_i = (x_i, y_i)$ are the 2D control points, and $N_{i,p}(t)$ are the p -degree B-spline basis functions defined on the parameter domain [20]. The coordinates (x, y) of a point on the curve are expressed in the parametric form

$$x(t) = \sum_{i=0}^n N_{i,p}(t)x_i \quad (1)$$

$$y(t) = \sum_{i=0}^n N_{i,p}(t)y_i \quad (2)$$

A B-Spline is characterised not only by the control points but also by a knot vector $U = \{u_0, \dots, u_m\}$, where $m =$

$n + p + 1$. With these ingredients, we can define the i^{th} B-spline basis function $N_{i,p}(t)$ of degree p as follows

$$N_{i,0}(t) = \begin{cases} 1 & \text{if } u_i \leq t < u_{i+1} \\ 0 & \text{otherwise} \end{cases}$$

$$N_{i,p}(t) = \frac{t - u_i}{u_{i+p} - u_i} N_{i,p-1}(t) + \frac{u_{i+p+1} - t}{u_{i+p+1} - u_{i+1}} N_{i+1,p-1}(t)$$

Note that the basis function $N_{i,p}(t)$ is a piecewise polynomial assuming non-zero values only on the interval $[u_i, u_{i+p+1})$. Therefore, it only affects the shape of the spline in a local section governed by the parameter t .

To formulate the descriptor, we treat a spectrum as a collection of spectral samples with two coordinates (λ_k, R_k) , where R_k is the k^{th} reflectance sample at the wavelength λ_k . Thus, the parametric form of a B-spline curve through these data points is obtained by representing the wavelength and reflectance as two functions of t , which we denote $\lambda(t)$ and $\mathcal{R}(t)$, respectively. The descriptor is then governed by a knot vector and a control point-set which minimise a cost function defined by squared differences between the measured reflectance R_k and the one computed in the parameter domain $\mathcal{R}(t)$. We depart from an initial interpolation of the sampled reflectance spectrum so as to arrive at the curve that minimises the cost function through a knot removal algorithm.

Once the control points and knot vector that minimise the cost function are at hand, we proceed to construct the NURBS descriptor. We do this by selecting the most discriminative components from the control point-set and the knot vector U . We observe that the knots and x -coordinates of each control point, i.e. the x_i variables in Equation 1, vary in a specific neighbourhood of the parameter domain where they govern the local support of the NURBS. This is reflected in the wavelength domain and, therefore, their variations across the spectra may not be as related to the spectrum shape as the y -coordinates of the control points, i.e. the y_i variables in Equation 2. In other words, the variations in the y -coordinates of the control points primarily determine the general shape of the spectrum. Thus, the y -coordinates should provide better discrimination between spectra than the knots and x -coordinates. To minimise the complexity of the descriptor in terms of vector length and computational cost, we limit our descriptor \mathcal{Y} to the y -coordinates of the control points $\mathcal{Y} = [y_0, y_1, \dots, y_n]^T$.

2.2. Target Function

As mentioned earlier, in statistical learning, a descriptor is desired to retain statistical information that gives high discriminative power. Hence, we devise a target function to optimise the choice of the interpolating B-spline curve

for each reflectance spectrum. Given a surface with reflectance R_k at each wavelength $\{\lambda_k\}, k = 1, \dots, l$, we aim to recover an interpolating curve \mathcal{C} with control points $P_i = (x_i, y_i)$ and a knot vector U such that

$$\lambda(t_k) = \sum_{i=0}^n N_{i,p}(t_k)x_i \quad (3)$$

$$\mathcal{R}(t_k) = \sum_{i=0}^n N_{i,p}(t_k)y_i \quad (4)$$

where the parameter $t_k \in \mathcal{U}$ corresponds to the k^{th} wavelength, i.e. $\lambda_k = \lambda(t_k) \forall k$.

The cost of interpolating the points (λ_k, R_k) using the B-spline curve above is then given by

$$K = \alpha \sum_{k=1}^l (\mathcal{R}(t_k) - R_k)^2 + (1 - \alpha)|U| \quad (5)$$

where $|\cdot|$ denotes the length of the vector argument and α is a constant between zero and unity.

Note that the first term in Equation 5 above is the weighted sum of squared errors in the 2D space defined by the wavelength-reflectance pairs, whereas the second term is the weighted number of knots. Thus, the optimal interpolating curve minimises the sum of squared distances $(\mathcal{R}(t_k) - R_k)^2$, while penalising a large number of knots. This imposes a balance in the resulting curve between describing the general shape of the reflectance data, while minimising the number of knots required to describe it. This trade-off is governed by α . A small value of α favours a short descriptor over one that respects the general shape of the original data. Also, note that, as the number of knots decreases, the interpolating curve becomes smoother. Thus, an appropriate choice of α makes the descriptor less susceptible to noise while preventing over-smoothing and loss of detail.

2.3. Minimisation of the Cost Function

This section presents our knot removal method which aims at minimising the interpolation cost introduced in Equation 5. We depart from an initial approximation of the NURBS curve to the reflectance spectrum under study. To this end, we apply the curve interpolation algorithm in [20], which employs the centripetal method of Lee [12] to recover parameter values for every control point.

With this initial approximation at hand, we proceed to remove knots sequentially using a knot removal method akin to that in [23]. The algorithm is a two-pass process. In the first pass, removable knots are identified. In the second pass, knots are sequentially removed and new control points are computed. Although being effective, this algorithm does not automatically determine the best knot to remove at each

pass, but rather assumes the knot to be removed to be designated as input. Our knot removal algorithm computes the potential cost reduction for removable knots. Once these knots and their contributions to the cost function are at hand, we employ Tiller's algorithm [23] to remove them.

Thus, our algorithm selects amongst the candidate knots the one that yields the maximum cost reduction. Also, note that, following the strategy above, the parameter t_k should be recovered for every wavelength λ_k . This is not a straightforward task since the function $\lambda(t_k)$ is expressed as a linear combination of the basis functions $N_{i,p}$ given in Equation 3. Nonetheless, this equation may not be solved analytically. We adopt a numerical approach in order to find an approximate solution and reduce the computational cost involved. In fact, it is reasonable to assume that the wavelength $\lambda(t_k)$ is an increasing function in the parameter domain. Therefore, for a given wavelength λ_k , we can perform a binary search for t_k such that $\lambda_k \sim \lambda(t_k)$.

Algorithm 1 *KnotRemoval*($Q, p, \alpha, target$)

Require: $Q, p, \alpha, target$

Q : The given data points
 p : The degree of basis functions
 α : The balance factor
 $target$: The target number of knots
 U_0, P_0 : the returned knots and control points

- 1: $(U_0, P_0) \leftarrow Interpolate(Q, p)$
- 2: **while true do**
- 3: $Reduction_{max} \leftarrow -1$ //Highest cost reduction
- 4: **for all** $u \in U_0$ **do**
- 5: $(flag, U_1, P_1) \leftarrow RemoveKnot(u)$ //flag is **true** if knot u is removable
- 6: **if** $flag$ **then**
- 7: $SSE_{old} \leftarrow SSE(U_0, u)$
- 8: $SSE_{new} \leftarrow SSE(U_1, u)$
- 9: $Reduction \leftarrow (1 - \alpha) + \alpha(SSE_{old} - SSE_{new})$
- 10: **if** $Reduction > Reduction_{max}$ **then**
- 11: $Reduction_{max} \leftarrow Reduction$
- 12: $candidate \leftarrow u$
- 13: **end if**
- 14: **end if**
- 15: **end for**
- 16: **if** $Reduction_{max} < 0$ **then**
- 17: **return** U_0, P_0
- 18: **end if**
- 19: Remove the *candidate* one time
- 20: Remove a data point corresponding to *candidate*
- 21: $Q' \leftarrow$ The remaining data points.
- 22: $(U_0, P_0) \leftarrow Interpolate(Q', p)$
- 23: **end while**
- 24: **return** U_0, P_0

2.4. Implementation Issues

The knot removal process can be viewed as a greedy approach which is reminiscent of a gradient descent method. In practice, this is an iterative method in which, at every iteration, we locate the knot that maximises the reduction in the cost K . The knot removal algorithm is summarised in Algorithm 1, where the sum of squared errors before and after knot removal are denoted as SSE_{old} and SSE_{new} , respectively. The $RemoveKnot(\cdot)$ procedure implements the knot removal algorithm in [23].

As a result of the local support property of B-spline curves, the removal only affects the curve partition in the neighbouring sections of the knot. Thus, for the sake of efficiency, the change in sum of squared errors can be computed as the change within the neighbourhood of the removal candidate u , to profit from the local support of the NURBS. To do this, we use the span of the NURBS [20] and employ lists to back-track their effect across the spline.

The knot removal algorithm terminates when removing any knot cannot further reduce the interpolation cost. However, the number of knots should be imposed as a hard constraint since descriptors should be normalised in length before they are input to classifiers. Hence, the knot removal method in Algorithm 1 is applied recursively by resampling data points in the parameter domain. The pseudocode for the recursion on the knot removal algorithm is shown in Algorithm 2. The resampling operation allows further knots to be removed by reducing the number of curve sections without changing the distribution of the control points from the original data. Note that smoothing becomes most acute when the parameter values for the resampling operation are given by the midpoints of the knot spans. Thus, we sample parameter values near the resulting knots to preserve the shape of the original curve.

Algorithm 2 $IterKnotRemoval(Q, p, \alpha, target)$

Require: $Q, p, target$

Q : The given data points

p : The degree of basis functions

$target$: The target number of knots

α : The balance factor

U_2, P_2 : The final knots and control points

1: $samples \leftarrow Q$

2: **while** $|U_2| > target$ **do**

3: $(U_1, P_1) \leftarrow Interpolate(samples, p)$

4: $(U_2, P_2) \leftarrow KnotRemoval(samples, p, \alpha, target)$

5: Select Parameters (t_1, \dots, t_l)

6: $samples \leftarrow Resample(p, U_2, P_2, t_1, \dots, t_l)$

7: **end while**

8: **return** U_2, P_2

To illustrate the behaviour of our algorithm, Figure 1

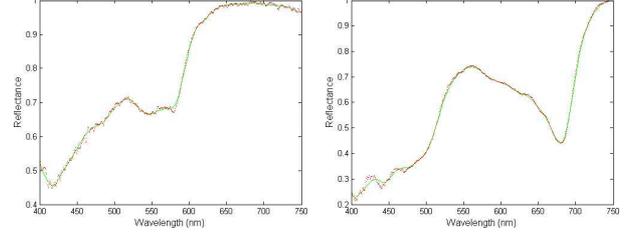


Figure 1. Result of knot removal on spectral reflectance samples of human skin (left) and a leaf (right).

shows the resulting NURBS curves achieved by removing knots from the reflectance spectra of human skin and a leaf. The original reflectance data, plotted in red, are sampled every $4nm$. In the figure, we have performed iterative knot removal with B-Spline basis functions of degree 3. This yielded 34 knots and 30 control points per curve. Note that, while this number of knots is considerably less than the number of original data points, the resulting NURBS curve, plotted in green, still aligns well with the original data. Thus, our cost-optimal knot removal algorithm is able to perform dimensionality reduction while respecting the global shape of the reflectance spectra. This is evident throughout the whole spectrum between 400 and $750nm$.

3. Applications

In this section we illustrate the utility of the NURBS-based descriptor for recognition and segmentation in hyperspectral images. Firstly, we perform skin spectra recognition on spectrometer data. We then expand this experiment to perform skin segmentation in hyperspectral images. Recognition and classification of the spectra is effected using a soft-margin Support Vector Machine (SVM) [3, 5], with a polynomial kernel of degree 3.

In our experiments, we compare our results with those yielded by descriptors recovered through Gaussian Mixture Regression [2]. The Gaussian Mixture model yields a better approximation to the raw spectral data than that of Principal Component Analysis (PCA). In fact, the sum of five Gaussian Mixtures gives an average root mean squared (RMS) error of 1.0% over 46 skin spectra, as compared to a much noisier approximation yielded by PCA with an error of 4.65%. Thus, through comparing the NURBS-based descriptor with the Gaussian Mixture descriptor, we are, by transitivity, comparing it with PCA.

In [2], a linear combination of M Gaussian basis functions with means μ_k , standard deviations σ_k and mixture coefficients β_k is fitted to each reflectance spectrum. The fitting problem is treated as a nonlinear least-squares optimisation one, which can be solved numerically using a Levenberg-Marquardt optimisation procedure [15]. To construct the feature vector, we concatenate the triplet $(\beta_k, \mu_k, \sigma_k)$ for every Gaussian. We have done this since,

| Feature combination | Accuracy (%) |
|--------------------------------|-----------------------------------|
| y -coordinate | 91.27 ± 1.6 |
| x -coordinate | 84.04 ± 5.02 |
| Knots | 88.83 ± 3.38 |
| x and y -coordinate | 90.86 ± 2.15 |
| Knots and y -coordinate | 91.14 ± 2.91 |
| Knots and x -coordinate | 87.21 ± 4.60 |
| Knots, x and y -coordinate | 91.04 ± 2.61 |

Table 1. The mean and standard deviation of the classification accuracy over 20 random tests with various feature combinations from a NURBS descriptor consisting of 31 knots and 27 control points.

in our experiments, these triplets consistently deliver better classification performance than other tuple combinations involving β_k , μ_k or σ_k .

3.1. Skin spectra recognition

We first turn our attention to skin spectra recognition. To this end, we have acquired in house 297 surface reflectance spectra using a StellarNet spectrometer. These surface reflectance spectra correspond to nine material categories, including cloth, paints, human skin, leaves, metal, coloured paper, plastic, porcelain and wood. In our dataset there are 157 spectra of human skin and 140 of the other materials above. All the spectra have been sampled in $0.5nm$ intervals between $430-720nm$ with a normalised reflectance between $[0, 1]$. The performance of the descriptors is given by the average performance over 20 random tests. In each random test, we train a classifier on a randomly selected training set comprised by 50 skin and 50 nonskin spectra. The remaining spectra are used for testing.

At this point, we turn our attention to the discriminative power of $\mathcal{Y} = [y_0, y_1, \dots, y_n]^T$ as compared to the x -coordinates of the control points and the knot vector. In Table 1, the recognition rate using \mathcal{Y} is compared with other choices of descriptor formulation. For the results shown here, we have used 31 knots and 27 control points with basis functions of degree 3. Our choice of descriptor \mathcal{Y} achieves a higher and more stable recognition accuracy than the knots and x -coordinates. Note that combining the knots and x , y -coordinates does not necessarily improve the classification performance.

We now compare the performance of the NURBS descriptor with that yielded by the Gaussian Mixtures over various descriptor lengths. Figure 2 shows the mean and standard deviation of recognition accuracy when using the two descriptors with varying lengths between 9 and 30, and raw spectral reflectance consisting of 291 bands. Overall, the NURBS descriptor performs better and is more stable than the Gaussian Mixture descriptor, corresponding to a recognition rate with a higher mean and a lower range of standard deviation as compared to that yielded by the Gaus-

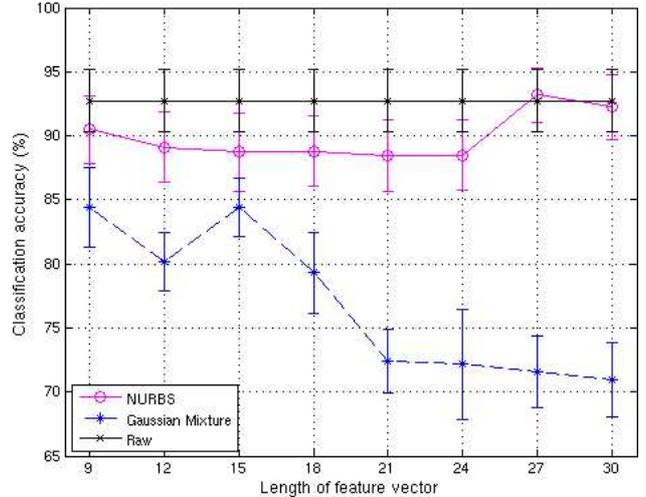


Figure 2. Classification performance on spectrometer data for the NURBS and Gaussian Mixture descriptors with varying lengths and raw spectral data.

sian Mixture descriptor. As shown in the figure, the NURBS descriptor is able to perform well, even with a descriptor length shorter than 15, achieving a mean classification accuracy between 88.43% and 93.20%. On the other hand, the Gaussian Mixture descriptor obtains the peak performance of 84.44% with 5 mixture components (corresponding to a length of 15) and declines rapidly as the order of the descriptor increases. This is consistent with the results reported in [2], where 5 mixture components give the lowest average root mean squared reconstruction error with respect to the original spectra. This also shows that using more than 5 mixture components not only adds computational complexity, but also degrades the classification accuracy severely. The NURBS descriptor does not suffer from this problem and achieves a performance comparable to the raw spectra. It is worth stressing that raw spectra have a length of 291, whereas our descriptor is never longer than 30 elements, which means our descriptor is much more computationally efficient.

3.2. Skin Segmentation

We now focus our attention on segmentation of hyperspectral imagery. Our hyperspectral images were taken using an OKSI Turnkey Hyperspectral Camera System which acquires wavelength resolved images in $10nm$ intervals over the visible spectrum. Our imagery depicts 7 subjects: two Caucasian, two Indian and three Asian. Figure 4 shows the images of these subjects at $670nm$. For purposes of validation, we have performed denoising on the spectral data. Each spectrum is smoothed by a 2^{nd} order Savitsky-Golay [17] filter with window size of 5 bands. This reduces the effects of noise and provides us with filtered spectral data which can then be used for our experiments.

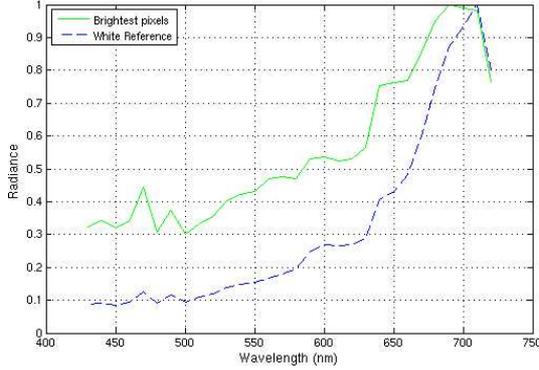


Figure 3. Comparison of the normalised lower bound $\hat{L}(\lambda)$ (solid line) and the true spectrum of the illuminant (dashed line).

To recover the illumination-invariant reflectance at individual pixels from the radiance of a hyperspectral image without prior knowledge of the light spectrum we have adopted a method akin to the continuum removal in [4].

Let an object with surface radiance $I(\lambda, u)$ at pixel u and wavelength λ be illuminated by an unknown illuminant whose spectrum is $L(\lambda)$. The dichromatic model [21] has long been used to decompose surface radiance into a diffuse and a specular component such that

$$I(\lambda, u) = g(u)L(\lambda)S(\lambda, u) + k(u)L(\lambda) \quad (6)$$

Using this model, we aim to recover the spectral reflectance $S(\lambda, u)$ at location u and wavelength λ from the measured spectral radiance $I(\lambda, u)$ on the image. In Equation 6, the shading factor $g(u)$ governs the proportion of diffuse light reflected from the object and depends solely on the surface geometry. On the other hand, the factor $k(u)$ accounts for the specularities on the image.

To recover the light spectrum $L(\lambda)$, we employ a variant of the continuum removal technique widely used in remote sensing [4]. The rationale behind the method used here is that, since objects in the scene cannot reflect more energy than what they receive from the light source, the radiance at the brightest pixel in the image can be treated as a lower bound on the illuminant. If the sample set over the spectra is large, the measured lower bound is expected to approximate the true illuminant. Thus, in a hyperspectral image with millions of pixels, we can use the lower bound $\hat{L}(\lambda) = \max_u I(\lambda, u)$ as an alternative to the light spectrum $L(\lambda)$. Figure 3 shows the lower bound $\hat{L}(\lambda)$ recovered from the hyperspectral image, in comparison with the illuminant spectrum as measured by a spectrometer. Both spectra are normalised to the range $[0, 1]$. From the figure, we can conclude that the lower bound $\hat{L}(\lambda)$ is approximately proportional to the true illumination spectrum $L(\lambda)$, i.e. $\hat{L}(\lambda) = \rho L(\lambda)$, where ρ is a constant.

With the approximate light spectrum $\hat{L}(\lambda)$, we have

$$\frac{I(\lambda, u)}{\hat{L}(\lambda)} \simeq \frac{1}{\rho} (g(u)S(\lambda, u) + k(u)) \quad (7)$$

We can assume that the spectral reflectance $S(\lambda, u)$ at pixel u is normalised to $[0, 1]$, i.e. $\max_\lambda S(\lambda, u) = 1$ and $\min_\lambda S(\lambda, u) = 0$. From Equation 7, we obtain

$$\begin{aligned} k(u) &\simeq \rho \min_\lambda \frac{I(\lambda, u)}{\hat{L}(\lambda)} \\ g(u) &\simeq \rho \left(\max_\lambda \frac{I(\lambda, u)}{\hat{L}(\lambda)} - \min_\lambda \frac{I(\lambda, u)}{\hat{L}(\lambda)} \right) \\ S(\lambda, u) &\simeq \frac{1}{g(u)} \left(\rho \frac{I(\lambda, u)}{\hat{L}(\lambda)} - k(u) \right) \end{aligned}$$

Later, we use the spectral reflectance $S(\lambda, u)$ recovered from hyperspectral images to classify image pixels into skin and nonskin using a method akin to that in Section 3.1. Note that ρ can be further removed by means of normalisation of the spectra.

For purposes of skin segmentation on our imagery, we label 1259 skin and 1256 nonskin pixels in the image of the Caucasian C_1 , 1141 skin and 1177 nonskin pixels in the image of the Indian I_1 and 983 skin and 994 nonskin pixels in the image of the Asian A_1 . Subsequently, we use the labelled pixels from each image as the training data to perform skin pixel recognition on the remaining images. We also perform experiments in which the training set contains the labelled pixels and the spectrometer data collected in Section 3.1. This illustrates the ability of our NURBS descriptor to generalise well to spectral data measured by various devices, with different spectral resolutions and spectrum lengths. For our NURBS descriptor, we have used 20 y -coordinates, whereas for the Gaussian Mixture descriptor we have employed, as before, 5 components. We have also compared the performance of the NURBS descriptor with raw pixel spectra consisting of 30 bands.

In Figure 5, we show segmentation results on our hyperspectral images. The first two columns show the normalised shading factor $k(u)$ and the specularity term $g(u)$ at each pixel in the hyperspectral images. The next columns show, from left-to-right, the segmentation results using the NURBS, Gaussian Mixture descriptor and raw spectral data, in which whiter pixels indicate a higher likelihood of being skin. The last three columns show that our descriptor clearly distinguishes between skin and non-skin pixels. Conversely, the raw spectral data and Gaussian Mixture descriptor produce false positives in the background and on clothes.

In Table 2, we quantify the skin segmentation performance for our descriptor and the alternatives in terms of the classification rate (CR), the correct detection rate (CDR) and false detection rate (FDR). The correct detection rate is the percentage of skin pixels correctly classified. The false detection rate is the percentage of nonskin pixels incorrectly classified. The classification rate is the overall

percentage of skin and nonskin pixels correctly classified. From the table, we observe that the NURBS descriptor achieves a similar classification rate across training sets, including either image pixels from C_1 , I_1 , A_1 , or combinations of images and spectrometer data ($C_1 + Spec$, $I_1 + Spec$ or $A_1 + Spec$). This implies that the descriptor is able to generalise well with respect to training skin type. Moreover, from the table, it becomes evident that this generalisation ability does not hold for the raw spectral data or the Gaussian Mixture descriptor. Secondly, combining spectrometer data with image pixel data yields a similar overall classification rate. This result shows that the descriptor is applicable to both high resolution data from the spectrometer, and low resolution, more noise-prone data from hyperspectral cameras.

Further, the NURBS descriptor consistently outperforms the alternatives. Note that although the Gaussian Mixture descriptor achieves a slightly better CDR , it suffers from a high false detection rate. In addition, the NURBS descriptor achieves a significantly lower false detection rate, with an average of 1.86%, while maintaining a competitive CDR of 88.58%. Compared with recent results on skin segmentation [19], the NURBS descriptor for hyperspectral imagery delivers a significant performance improvement over trichromatic features. The survey in [19] reported classification rates of no more than 90%, with false detection rates not lower than 10%. In addition, the NURBS descriptor only requires a training set 10 to thousand times smaller than those reported in [19].

4. Conclusions

In this paper, we presented a descriptor comprised by the y -coordinates of the control points for the NURBS curves interpolated on reflectance spectra. We have presented a method to recover the knots and control points making use of a knot removal algorithm in the parameter domain. We demonstrate the utility of this feature for skin spectra recognition and segmentation. Our experimental results show that the NURBS descriptor is more stable and provides a margin of performance improvement over the Gaussian Mixture representation of reflectance spectra and raw spectral data.

References

- [1] E. Angelopoulou. Specular highlight detection based on the fresnel reflection coefficient. In *Int. Conf. on Computer Vision*, 2007.
- [2] E. Angelopoulou, R. Molana, and K. Daniilidis. Multispectral skin color modeling. In *Computer Vision and Pattern Recognition*, pages 635–642, 2001.
- [3] B. E. Boser, I. Guyon, and V. Vapnik. A training algorithm for optimal margin classifiers. In *Computational Learning Theory*, pages 144–152, 1992.
- [4] R. Clark and T. Roush. Reflectance spectroscopy: Quantitative analysis techniques for remote sensing applications. *Journal of Geophysical Research*, 89:6329–6340, 1984.
- [5] C. Cortes and V. Vapnik. Support-vector networks. *Machine Learning*, 20(3):273–297, 1995.
- [6] M. Dundar and D. Landgrebe. Toward an optimal supervised classifier for the analysis of hyperspectral data. *IEEE Transaction on Geoscience and Remote Sensing*, 42(1):271–277, 2004.
- [7] K. Fukunaga. *Introduction to Statistical Pattern Recognition*. Academic Press, NY, 2nd edition, 1990.
- [8] D. W. Jacobs, P. N. Belhumeur, and R. Basri. Comparing images under variable illumination. In *Computer Vision and Pattern Recognition*, page 610, 1998.
- [9] L. Jimenez and D. Landgrebe. Hyperspectral data analysis and feature reduction via projection pursuit. *IEEE Transaction on Geoscience and Remote Sensing*, 37(6):2653–2667, 1999.
- [10] I. Jolliffe. *Principal Component Analysis*. Springer, 2002.
- [11] D. Landgrebe. Hyperspectral image data analysis. *IEEE Signal Processing Magazine*, 19, 2002.
- [12] E. T. Y. Lee. Choosing nodes in parametric curve interpolation. *Computer-Aided Design*, 21(6):363–370, 1989.
- [13] R. Lenz, P. Latorre, and P. Meer. The hyperbolic geometry of illumination-induced chromaticity changes. In *Computer Vision and Pattern Recognition*, 2007.
- [14] S. Lin and S. W. Lee. Using chromaticity distributions and eigenspaces for pose, illumination, and specular invariance 3d object recognition. In *Computer Vision and Pattern Recognition*, pages 426–431, 1997.
- [15] D. W. Marquardt. An algorithm for least-squares estimation of nonlinear parameters. *SIAM Journal on Applied Mathematics*, 11(2):431–441, 1963.
- [16] S. K. Nayar and R. M. Bolle. Reflectance based object recognition. *Int. J. Comput. Vision*, 17(3):219–240, 1996.
- [17] S. J. Orfanidis. *Introduction to signal processing*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1995.
- [18] J. I. Park, M. H. Lee, M. Grossberg, and S. Nayar. Multispectral imaging using multiplexed illumination. In *Int. Conf. on Computer Vision*, 2007.
- [19] S. L. Phung, A. Bouzerdoum, and D. Chai. Skin segmentation using color pixel classification: Analysis and comparison. *IEEE Trans. Pattern Anal. Mach. Intell.*, 27(1):148–154, 2005.
- [20] L. Piegl and W. Tiller. *The NURBS book*. Springer-Verlag, London, UK, 1995.
- [21] S. A. Shafer. Color research and applications. 10:210–218, 1985.
- [22] D. Slater and G. Healey. Object recognition using invariant profiles. In *Computer Vision and Pattern Recognition*, pages 827–832, 1997.
- [23] W. Tiller. Knot-removal algorithm for NURBS curves and surfaces. *Computer-Aided Design*, 24(8):445–453, 1992.

| | NURBS | | | Gaussian Mixture | | | Raw spectral data | | |
|--------------|-------|--------|--------|------------------|--------|--------|-------------------|--------|--------|
| | CR(%) | CDR(%) | FDR(%) | CR(%) | CDR(%) | FDR(%) | CR(%) | CDR(%) | FDR(%) |
| C_1 | 98.08 | 83.78 | 1.03 | 92.49 | 88.95 | 7.13 | 94.82 | 76.38 | 3.80 |
| I_1 | 96.71 | 98.99 | 3.41 | 87.54 | 97.66 | 12.98 | 86.86 | 85.06 | 13.15 |
| A_1 | 97.22 | 91.18 | 2.28 | 87.74 | 92.85 | 12.39 | 95.25 | 73.55 | 3.56 |
| $C_1 + Spec$ | 98.08 | 81.11 | 0.91 | 91.47 | 89.15 | 8.22 | 94.07 | 86.84 | 5.26 |
| $I_1 + Spec$ | 97.81 | 90.88 | 1.87 | 87.76 | 97.78 | 12.76 | 97.74 | 92.58 | 1.97 |
| $A_1 + Spec$ | 97.59 | 85.53 | 1.64 | 86.71 | 92.36 | 13.44 | 94.81 | 68.04 | 3.67 |
| Average | 97.58 | 88.58 | 1.86 | 88.95 | 93.12 | 11.15 | 93.92 | 80.41 | 5.23 |

Table 2. The classification rate (CR), correct detection rate (CDR) and false detection rate (FDR) on test images using the NURBS descriptor (with 20 y -coordinates), Gaussian Mixture descriptor (with 5 mixture components) and raw spectral data (with 30 spectral bands), by training an SVM on image pixels of the Caucasian subject C_1 , the Indian subject I_1 and the Asian subject A_1 and on combinations of image pixels and the spectrometer data.

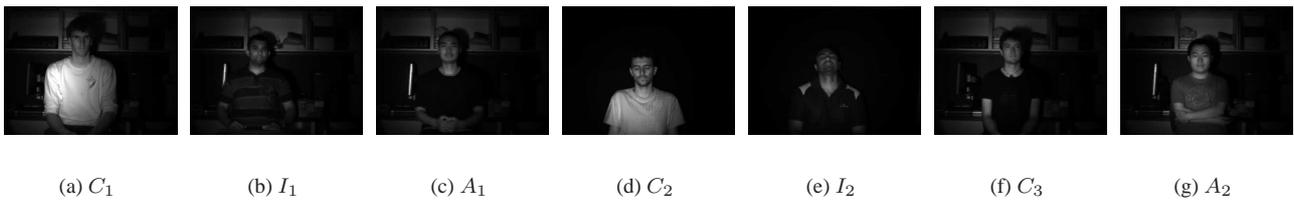


Figure 4. Human subjects used for skin segmentation.



Figure 5. Results of skin segmentation using training pixels from the image of subject C_1 . From left to right, the columns show the shading map, specular map, segmentation result using the NURBS descriptor, Gaussian Mixture descriptor, and raw spectral data. The test subjects (from the top to bottom row) are I_1 , A_1 , C_2 , I_2 , C_3 , A_2 .