

A Comparative Evaluation of Spectral Reflectance Representations for Spectrum Reconstruction, Interpolation and Classification

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

¹National ICT Australia (NICTA) *, Locked Bag 8001, Canberra ACT 2601, Australia

²Research School of Engineering, Australian National University, Canberra ACT 0200, Australia

Abstract

Due to the high dimensionality of spectral data, spectrum representation techniques have often concentrated on modelling the spectra as a linear combination of a small basis set. Here, we focus on the evaluation of a B-Spline representation, a Gaussian mixture model, PCA and wavelets when applied to represent real-world spectrometer and spectral image data. These representations are important since they open up the possibility of reducing densely sampled spectra to a compact form for spectrum reconstruction, interpolation and classification. In particular, we shall perform an evaluation of these representations for the above tasks on two datasets consisting of reflectance spectra and hyperspectral images.

1. Introduction

The development of image sensor technology has made it possible to capture image data in hundreds of bands covering a broad spectrum. Traditionally, multispectral and hyperspectral imaging have been developed as a tool for aerial image classification [12] and aerial surveillance [7]. For terrestrial applications, imaging spectroscopy poses many challenges and opportunities in biological and medical applications [14]. In addition, it has been utilised for colour reproduction of artworks [5] and historical documents [16], printing [23] and material recognition [13].

Due to the high dimensional nature of 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 [17]. The main bulk of work on spectrum representation so far has concentrated on modelling spectra as a linear combination of a small basis set. Maloney validated the use of low-dimensional linear models for representing reflectance

spectra through a number of evaluations on empirical measurements of surface reflectance [19]. In [20], Marimont and Wandell computed a set of linear basis functions which yield the minimal approximating error for surface reflectance spectra and illuminant spectra. In another work, Angelopoulou *et al.* [3] used spectrophotometric data to model skin colour using several sets of basis functions including Gaussians and their first derivatives, wavelets and PCA. In [15], Huynh and Robles-Kelly employed a compact B-Spline basis to represent reflectance spectra. On the other hand, graph-based approaches usually aim at classifying rather than reconstructing hyperspectral images [4, 6].

Here, we focus on the evaluation of a number of these representations for the purposes of representing, reconstructing, interpolating and classifying spectral data. For our analysis we have used the B-Spline representation in [15], the Gaussian mixture model in [3], PCA and wavelets. These representations open up the possibility of reducing densely sampled reflectance spectra, which can consist of hundreds of spectral samples, to an efficient and compact form for reconstruction and classification.

Moreover, these representations apply equally to both spectral radiance and spectral reflectance data. For our evaluation, we use real-world spectrometer and spectral images for evaluating the performance of the representations under study when applied to spectrum reconstruction, interpolation, skin recognition and biometrics. This is important since the continuous form of the B-Spline, the Gaussian mixture and the wavelet representations permits a comparison of spectra of dissimilar lengths via a set of numerically stable algorithms. These representation provide the additional benefit of allowing operations such as absorption band detection to be effected in the continuous spectral domain using tools from functional analysis rather than discrete approximations.

In the next section, we describe the spectra representations under our study. In Section 3, we show results on real-world data pertaining reconstruction, interpolation, recognition and biometrics. Finally, in Section 4 we conclude our evaluation.

*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.

2. Spectrum Representations

2.1. B-Splines

In [15], the reflectance spectra is treated as a set of 2-D points whose coordinates correspond to wavelength-reflectance sample pairs. These are interpolated using a B-Spline via an iterative knot removal scheme which minimises the representation length, i.e. the number of knots and control points, subject to a user-defined approximation error.

To formulate the representation, the authors in [15] treated a discrete sample of a reflectance spectrum as a collection of 2-D points with two coordinates (λ_k, R_k) , where R_k is the k -th reflectance sample at the wavelength $\{\lambda_k\}, k = 1, \dots, l$. With this treatment, the parametric form of a p -th degree B-spline curve \mathcal{C} through these data points can be represented as two functions denoted as $\lambda(t)$ and $\mathcal{R}(t)$ corresponding to the wavelength and the reflectance, as follows

$$\lambda(t) = \sum_{i=0}^n N_{i,p}(t)x_i \quad \mathcal{R}(t) = \sum_{i=0}^n N_{i,p}(t)y_i, \quad (1)$$

where t is an independent parameter in a domain \mathcal{U} , $P_i = (x_i, y_i)$, where $i = 0, \dots, n$, are the coordinates of the control points and $N_{i,p}(t)$ is the i -th B-spline basis function of degree p defined over the corresponding knot vector [22].

Here we note that the representation presented in [15] was specifically designed for a single reflectance spectrum. In this paper, we extend the representation in [15] to perform representation and classification tasks on multiple spectra. To do this, we enforce a common set of B-Spline basis functions $N_{i,p}(t), i = 0, \dots, n$ across the input spectra. Thus, for a collection of reflectance spectra $\{\mathbf{R}_v | \mathbf{R}_v = [R_{v,1}, R_{v,2}, \dots, R_{v,l}]^T\}$, where v denotes the spectrum index, we have

$$\mathcal{R}_v(t) = \sum_{i=0}^n N_{i,p}(t)y_{v,i}, \quad (2)$$

where the expression in Equation 1 for $\lambda(t)$ remains unaffected since the spectra are sampled at the same wavelengths and $y_{v,i}$'s are the reflectance coordinates of the control points corresponding to the spectrum index v .

Following [15], we consider the optimal interpolating curve to be the one that minimises the sum of squared distances from the sampled points of the spectra subject to a pre-determined length m of the knot vector. The cost function is given by

$$K = \sum_v \sum_{k=1}^l (\mathcal{R}_v(t_k) - R_{v,k})^2, \quad (3)$$

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

The modified expression in Equation 2 does not overly affect the optimisation procedure in [15]. The optimisation of the cost function in Equation 3 can still be effected by applying the interpolation algorithm in [22] so as to obtain an initial B-Spline curve for each spectrum with a zero interpolation error. The centripetal method of Lee [18] is still used to establish an initial knot vector for each input spectrum. At this stage, we enforce a common knot vector across all the input spectra throughout the remainder of the optimisation procedure. To achieve this, we modify the global interpolation method of Piegl and Tiller in [22] by averaging the knot vectors for individual spectra.

It is worth noting that, since the B-Spline representation of each spectrum needs to be of the same length, the knot removal procedure should guarantee to reach a pre-determined number of knots. Note that there may be cases in which no knots can be removed within the interpolation error tolerance. To this end, we follow the method in [10] to pre-select a knot vector that is one-element shorter than that obtained in the previous iteration and apply a global approximation to the original spectra with this knot vector. Accordingly, this procedure guarantees a reduction in the knot vector at each iteration.

As observed by Huynh and Robles-Kelly [15], only the reflectance coordinates $y_{v,i}$ determine the characteristic shape of a spectrum. Hence, they should provide better representation and discrimination power for reflectance spectra than the knots and the wavelength coordinates. Therefore, for classification purposes, we form a feature vector \mathcal{F}_v for a spectrum v by concatenating the reflectance coordinates, i.e. $\mathcal{F}_v = [y_{v,0}, y_{v,1}, \dots, y_{v,n}]^T$.

2.2. Gaussian Mixture

We now turn our attention to the representation reported in [3]. Here, the spectrum was modelled as a linear mixture of M Gaussian basis functions. The representation in [3] treats the reflectance $\mathcal{R}(\cdot)$ as a function of the wavelength λ , where the input reflectance spectrum is given by

$$\mathcal{R}(\lambda) = \sum_{k=1}^M \frac{\gamma_k}{\sqrt{2\pi}\sigma_k} \exp\left(-\frac{(\lambda - \mu_k)^2}{2\sigma_k^2}\right) \quad (4)$$

where M is the number of mixture components and the k -th Gaussian component is associated with a mean μ_k , a standard deviation σ_k and a mixture coefficient γ_k .

Given a number of Gaussian components, we can recover the Gaussian parameters by fitting the Gaussian mixture to the input spectrum so as to minimise the fitting error with respect to Equation 4. Here, we solve this non-linear least-squares problem numerically using a Levenberg-Marquardt optimisation procedure [21]. To construct the Gaussian Mixture feature for recognition purposes, we concatenate the triplet $(\gamma_k, \mu_k, \sigma_k)$ over the mix-

ture components. We have done this since in our experiments this formulation consistently delivers better classification performance than any other combinations of γ_k , μ_k and σ_k .

2.3. Principal Component Analysis

Another alternative representation is the use of the principal components of the input reflectance spectra (PCA) [11]. Thus, we represent each spectrum as a vector composed of the reflectance values sampled over the wavelengths. Subsequently, we subtract the mean spectrum from the input spectra to obtain a zero mean vector across the spectra. Next, we assemble the mean-centred vectors into a data matrix \mathbf{D} with dimension $|v| \times l$ whose rows correspond to the spectra and columns to the wavelengths.

At this stage, the eigenvectors of the matrix $\mathbf{D}^T\mathbf{D}$ with the largest eigenvalues correspond to the directions in which the input spectra have the largest variances. Following an eigen decomposition of the matrix $\mathbf{D}^T\mathbf{D}$, we obtain l eigenvectors $e_j, j = 1, \dots, l$ sorted according to decreasing eigenvalue. Assuming that we aim to extract the top k features from the input spectra, we form an $l \times k$ matrix \mathbf{A} whose columns consist of the k eigenvectors. A representation \mathbf{F} by principal components is obtained by projecting the input spectra into the k -dimensional subspace spanned by the top eigenvectors $e_j, j = 1, \dots, k$, according to $\mathbf{F} = \mathbf{D}\mathbf{A}$.

2.4. Wavelets

Similar to Angelopoulou *et al.* [2], we treat the reflectance \mathcal{R} as a one-dimensional signal with respect to the wavelength λ . Since the input consists of discrete wavelength-indexed samples of reflectance spectra, we model the fluctuation of reflectance in a range of discrete scales and offsets in the spectral domain. As a result, a reflectance spectrum $\mathcal{R}(\lambda)$ can be represented using both the approximation and wavelet terms [1], as follows

$$\mathcal{R}(\lambda) = \sum_{n=-\infty}^{\infty} \alpha_{m_0,n} \phi_{m_0,n}(\lambda) + \sum_{m=-\infty}^{m_0} \sum_{n=-\infty}^{\infty} \beta_{m,n} \psi_{m,n}(\lambda). \quad (5)$$

In Equation 5, a linear combination of the scaling functions $\phi_{m_0,n}(\lambda)$ provides a coarse approximation of the reflectance spectrum $\mathcal{R}(\lambda)$ at a chosen scale m_0 . In addition, the signal details of the reflectance spectrum is modelled by a range of wavelet functions $\psi_{m,n}(\lambda)$ with a scale from m_0 down to negative infinity. The scaling and wavelet functions above are the scaled and shifted version of a father scaling function $\phi(\lambda)$ and a mother wavelet $\psi(\lambda)$ in the wavelength domain. Further, $\alpha_{m_0,n}$ and $\beta_{m,n}$ are the approximation and wavelet (detail) coefficients corresponding to the above

functions and are computed as

$$\alpha_{m_0,n} = \sum_{\lambda \in \Lambda} \mathcal{R}(\lambda) \phi_{m_0,n}(\lambda) \quad \beta_{m,n} = \sum_{\lambda \in \Lambda} \mathcal{R}(\lambda) \psi_{m,n}(\lambda), \quad (6)$$

where Λ is the set of all discrete sampled wavelengths in the spectrum under study.

Note that we can employ a wide range of choice of the approximation and wavelet functions [1]. Regardless, these will yield a set of approximation and wavelet coefficients for each spectrum. Subsequently, we order the approximation and wavelet terms by sorting the sum of coefficients for each term over the spectra in descending order. The input reflectance spectra can then be represented by a linear combination of the functions corresponding to the largest sums of coefficients [1].

3. Experiments

In this section, we evaluate the performance of the spectral representations described in Section 2. In Sections 3.1 and 3.2, we present the accuracy of reflectance spectrum reconstruction and interpolation using spectrometer data and hyperspectral images. Subsequently, in Section 3.3, we examine the use of these representations for skin recognition and ethnic group classification.

In the following experiments, we have implemented a third-degree B-Spline representation according to [15] with a knot vector and wavelength coordinates shared across the input spectra, as described in Section 2.1. In addition, we compute the Gaussian Mixture representation [3] by means of Levenberg-Marquardt optimisation [21]. To reconstruct spectra using PCA, we employed least-square projections of the original spectra onto the subspace spanned by the principal components. To fairly compare with the B-Spline representation, we employed the Daubechies 8 (D8) wavelet [9], which is capable of encoding polynomial information in a signal up to the third degree.

We perform a comparative evaluation on the two following datasets. The first of these consists of 297 reflectance spectra acquired in house using a StellarNet spectrometer, which we name the StellarNet dataset. These spectra correspond to nine material categories, including cloth, different kinds of paint, human skin, leaves, metals, papers of different colours, plastic, porcelain and wood. The dataset includes 157 spectra of human skin and 140 of other materials. These spectra have been sampled at every $1nm$ in the $430-720nm$ wavelength range and the reflectance is normalised to the interval $[0, 1]$. The second dataset consists of multi-spectral face images captured by a hyperspectral camera under 10 artificial light sources with varying illumination directions and spectral power. The dataset comprises the frontal views of 51 human subjects in front of cluttered backgrounds, acquired by a hyperspectral camera at a spec-




Figure 1. Original and reconstructed reflectance spectra of human skin. From left to right: reconstruction results for representation lengths of 12, 21 and 30.

tral resolution of $10nm$ between $430nm$ and $720nm$. The image spectral reflectance is computed as a normalisation of the input image radiance by the illumination power spectrum measured using a white Spectralon calibration target. Note that, although the above data consists of spectra spanning only the visible wavelengths, we could clearly apply the above spectrum representations to non-visible spectral segments.

3.1. Reflectance Spectrum Reconstruction

We first turn our attention to the reconstruction results for the samples in the StellarNet dataset. For this experiment, we have re-sampled the original spectra at an interval of $5nm$ in the visible range between 430 and $720nm$. In Figure 1, we present reflectance spectra of human skin reconstructed with the above representations, overlaid on top of the original spectrum. The columns show the reconstructed spectra with representation lengths of 12, 21 and 30. In general, for all the representations, the reconstruction performance improves with an increase in the representation length. Among these, the B-Splines and PCA achieve the highest reconstruction accuracy. These two representations are also much better than the Gaussian Mixture and wavelets at capturing absorption bands. The evidence for this capability is discernible at the $550nm$ wavelength of the skin sample. In contrast, the Gaussian Mixture and wavelets tend to smooth out the original curves, thus not being able to reproduce exactly the original spectra.

To provide a quantitative study, we employ two measurements for the reconstruction error. The first of these is the mean of the absolute difference between the reconstructed and the original reflectance spectra per band, *i.e.* the average absolute reconstruction error (ARE). The second one is the relative reconstruction error (RRE), which we quantify as the ratio of the absolute error to the original reflectance value. With these measurements, we examine the variation of the reconstruction error as the representation length ranges between 9 and 30, in increments of 3. Equivalently, the Gaussian mixture representation in [3] consists of be-

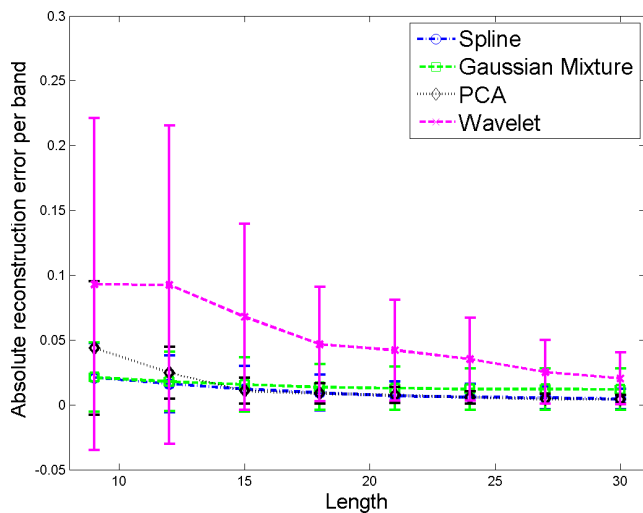


Figure 2. Absolute reconstruction error (ARE) for the reflectance spectra in the StellarNet dataset with respect to the representation length. The plot shows errors for the B-Spline representation [15], the Gaussian Mixture in [3], PCA and the Daubechies 8 (D8) wavelet.

tween 3 and 10 components.

Figure 2 shows the variation of the absolute reconstruction error of the StellarNet dataset with respect to the representation length. In particular, we observe that the B-Spline representation and PCA yield similar mean errors and consistently outperform the other two when used with a length of 15 or more. The absolute reconstruction error for all the representations, except the Gaussian Mixture, diminishes steadily and stabilises with an increasing number of control points. However, the Gaussian Mixture representation does not improve the reconstruction accuracy with six or more components. We also note that the B-Splines outperform PCA with a short representation length (of 12 or less) and yields similar performance to PCA with longer representation lengths. These observations imply that the B-Spline representation offers the best overall balance between reconstruction accuracy and compactness.

	B-Spline	Gaussian Mixture	PCA	Wavelet
Absolute error	$(0.42 \pm 0.30) \times 10^{-2}$	$(2.72 \pm 2.85) \times 10^{-2}$	$(0.27 \pm 0.19) \times 10^{-2}$	$(1.02 \pm 0.55) \times 10^{-2}$
Relative error	$4.51 \pm 2.31(\%)$	$16.61 \pm 11.53(\%)$	$2.15 \pm 3.30(\%)$	$15.86 \pm 6.39(\%)$

Table 1. Absolute and relative error of reconstructing the hyperspectral image database with a representation length of 18. The representations used include a third-degree B-Spline representation with 18 control points, a Gaussian Mixture with 6 components, a representation consisting of 18 principal components resulting from PCA and the 18 highest coefficients of a Daubechies 8 wavelet.

In Table 1, we compare the reconstruction performance of the four representations on the hyperspectral image dataset described earlier when the representation length is set to 18. Here, PCA is the best performer, followed by B-Splines. These two representations significantly outperform the other two, with errors several times as small as those delivered by the wavelet and the Gaussian mixture representations. Although PCA achieves a higher reconstruction accuracy than B-Splines, it cannot be utilised for spectrum interpolation purposes.

3.2. Reflectance Spectrum Interpolation

Now we demonstrate the capability of some of the above representations for interpolating reflectance spectra in the spectral domain. We note that PCA cannot be used for spectrum interpolation at arbitrary wavelengths. In addition, the results in Figures 1 and 2 suggest that the wavelet representation may be of limited practical interest for spectrum interpolation. Therefore, we limit the comparison of spectrum interpolation performance to only the B-Spline and Gaussian Mixture representations.

To obtain input for interpolation, we down-sampled the spectral domain of both the StellarNet dataset and the multi-spectral image database at every $5nm$, resulting in 59 bands per spectrum in the visible range. Next, we derive B-Spline and Gaussian Mixture representations for the down-sampled reflectance spectra. We then interpolate the down-sampled reflectance spectra into the original wavelengths by evaluating the B-Spline and the Gaussian Mixture functions.

We now assess the variation of the interpolation error for various representation lengths between 9 and 30, in increments of 3. These representation lengths are equivalent to 9 to 30 B-Spline control points and 3 to 10 Gaussian components. In Figure 3, we plot the means and standard deviations of the interpolation errors for the B-Spline and the Gaussian Mixture representation as separate curves. Similar to the previous experiment, these representations achieve almost identical interpolation accuracy with a short representation length, *i.e.* 6. As the length increases, the B-Spline representation appears to be superior to the Gaussian Mixture in terms of both interpolation accuracy and stability. Moreover, the interpolation error for the B-Spline representation decreases at a faster rate than that for the Gaussian Mixture. This is evident due to the increasing gap between

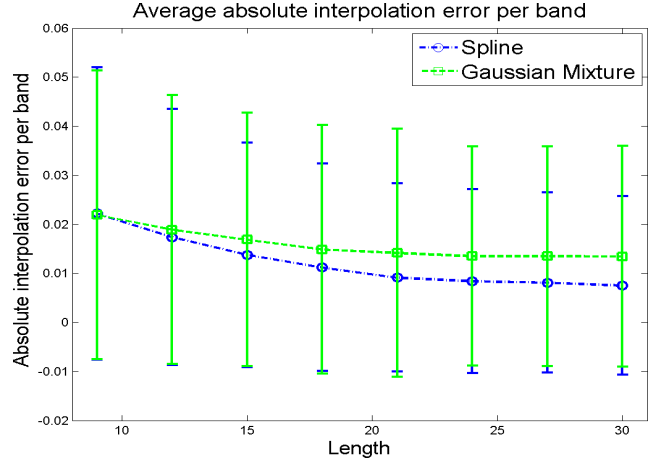


Figure 3. Absolute interpolation error versus the representation length for the B-Spline and the Gaussian Mixture representations, evaluated on the reflectance spectra in the StellarNet dataset.

Length	Descriptor	
	B-Spline	Gaussian Mixture
9	$(2.22 \pm 2.98) \times 10^{-2}$	$(2.19 \pm 2.94) \times 10^{-2}$
12	$(1.73 \pm 2.61) \times 10^{-2}$	$(1.89 \pm 2.74) \times 10^{-2}$
15	$(1.38 \pm 2.29) \times 10^{-2}$	$(1.69 \pm 2.58) \times 10^{-2}$
18	$(1.12 \pm 2.12) \times 10^{-2}$	$(1.49 \pm 2.54) \times 10^{-2}$
21	$(0.91 \pm 1.92) \times 10^{-2}$	$(1.42 \pm 2.52) \times 10^{-2}$
24	$(0.84 \pm 1.88) \times 10^{-2}$	$(1.35 \pm 2.23) \times 10^{-2}$
27	$(0.81 \pm 1.83) \times 10^{-2}$	$(1.35 \pm 2.24) \times 10^{-2}$
30	$(0.75 \pm 1.82) \times 10^{-2}$	$(1.35 \pm 2.25) \times 10^{-2}$

Table 2. Absolute interpolation error for the StellarNet dataset using B-Splines and Gaussian Mixtures.

the two curves as the representation length increases.

Table 2 shows the numerical errors corresponding to the graph plotted in Figure 3. As shown, errors are in the order of 10^{-2} , with that for the B-Spline representation below 10^{-2} using 21 or more control points. It is consistent with Figure 3 that the Gaussian Mixture representation tends to reach its performance limit at a length of 24, *i.e.* 6 Gaussian components. On the other hand, the B-Spline representation still improves its interpolation accuracy beyond 24 control points. With 30 control points, its absolute interpolation error is just above half of that of the Gaussian Mixture.

Next, we examine the interpolation performance of B-Splines and Gaussian Mixtures on the multi-spectral image database. To obtain the input for this experiment, we

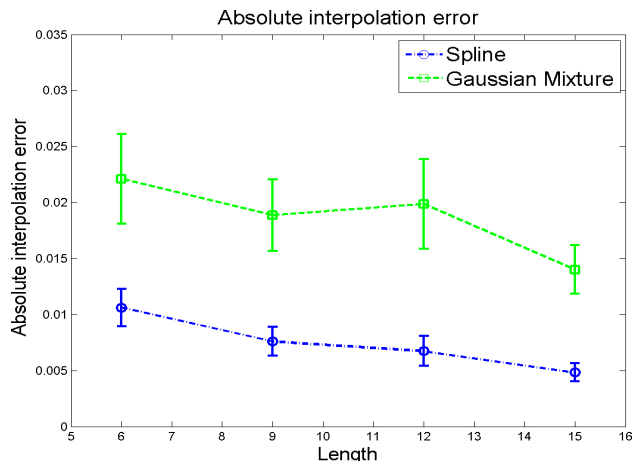


Figure 4. Absolute interpolation error versus the representation length for the hyperspectral image database. The plot shows the performance of a third-degree B-Spline and a Gaussian Mixture representation with lengths between 6 and 15.

Length	Descriptor	
	B-Spline	Gaussian Mixture
6	$(1.06 \pm 0.17) \times 10^{-2}$	$(2.21 \pm 0.40) \times 10^{-2}$
9	$(0.76 \pm 0.13) \times 10^{-2}$	$(1.89 \pm 0.32) \times 10^{-2}$
12	$(0.67 \pm 0.13) \times 10^{-2}$	$(1.99 \pm 0.40) \times 10^{-2}$
15	$(0.48 \pm 0.08) \times 10^{-2}$	$(1.40 \pm 0.22) \times 10^{-2}$

Table 3. The absolute interpolation error for the hyperspectral reflectance image database using a third-degree B-Spline and a Gaussian Mixture representation with various lengths between 6 and 15.

down-sampled the original spectral resolution from $10nm$ to $20nm$. This reduces the original number of spectral bands from 30 to 15. As before, we extracted a third-degree B-Spline and a Gaussian Mixture representation from the spectrally down-sampled images. These representations are then used for an interpolation over the original wavelengths.

In Figure 4, we show the absolute interpolation errors for both the B-Spline and the Gaussian Mixture representations with lengths of 6, 9, 12 and 15. Again, the B-Spline representation outperforms the Gaussian mixture in both accuracy and stability. It is evident that the error difference is significant, being several standard deviations. In fact, Table 3 shows that the error for the B-Spline is consistently less than half of that for Gaussian Mixture. While the B-Spline representation exhibits a steady reduction in error when the representation length increases, the Gaussian Mixture suffers from a fluctuation at the length of 12.

3.3. Skin Recognition and Biometrics

In this section we illustrate the utility of the above representations as descriptors for skin recognition and skin biometrics. We commence by performing skin spectrum recognition on the StellarNet dataset and later study the behaviour

of these descriptors in recognising skin types on the multi-spectral images. As before, we study the recognition accuracy with the representation length varying between 9 and 30, in increments of 3.

3.3.1 Skin Spectrum Recognition

To perform the recognition experiments on the StellarNet dataset, we have repeated the experiment in 20 trials. The data for this experiment has been sub-sampled at a $5nm$ sampling interval in the wavelength domain, resulting in 59 bands per spectrum. For each trial, we randomly selected 50 skin and 50 non-skin spectra for training and use the remaining ones for testing. The performance of the descriptors involved in this experiment is computed as the average over the 20 trials. The skin spectra recognition task is effected using a soft-margin Support Vector Machines (SVM) classifier [8] with a third-degree polynomial kernel function. The other parameter values of the SVM classifier are selected by a four-fold cross validation procedure.

In Table 4, we report the skin recognition accuracy on the StellarNet dataset. From the second to the sixth columns, we show the means and standard deviations yielded by the B-Spline descriptor, the Gaussian Mixture descriptor, PCA, a Daubechies 8 (D8) wavelet and the Raw input reflectance. Note that the Raw reflectance feature consists of 59 spectral bands.

The table shows that the PCA delivers the highest overall recognition rates among the five features, followed closely by the D8 wavelet, raw reflectance and B-Splines. However, the difference in recognition performance between PCA, the D8 wavelet, raw reflectance and B-Splines is not significant, being within a standard deviation. On the other hand, the Gaussian Mixture descriptor is significantly outperformed by the former ones, especially with 12 or more elements in the representation. Although having a comparable or better recognition rate, PCA lacks the capability of spectrum interpolation, which is accommodated by the other representations. The performance of PCA can be explained by the fact that, when applied to the spectra, they may remove sampling noise.

Furthermore, the B-Spline descriptor does not suffer as much fluctuation of recognition accuracy as the descriptor length varies. This is the case for the Gaussian Mixture, PCA and the D8 wavelet. This implies that B-Splines can be employed as a descriptor with a short length, to achieve a high recognition accuracy with a low computational overhead. With 30 elements, the B-Spline descriptor delivers a recognition accuracy equivalent to the use of the full raw reflectance spectrum using nearly twice as many bands, i.e. 59 bands.

Length	Descriptor				
	B-Spline	Gaussian Mix.	PCA	Wavelet	Raw
9	93.25 ± 2.51	90.66 ± 2.11	91.75 ± 2.85	93.53 ± 2.54	-
12	92.69 ± 2.85	81.98 ± 2.93	97.06 ± 2.43	95.79 ± 2.06	-
15	94.16 ± 2.12	83.43 ± 4.05	97.87 ± 1.37	94.01 ± 2.27	-
18	94.14 ± 1.54	81.65 ± 3.25	95.10 ± 3.12	94.42 ± 1.98	-
21	93.83 ± 2.35	83.55 ± 3.30	95.25 ± 3.69	95.03 ± 1.67	-
24	93.86 ± 1.75	83.60 ± 2.95	94.29 ± 3.22	95.08 ± 2.55	-
27	94.54 ± 2.17	82.13 ± 4.28	95.89 ± 2.19	94.62 ± 1.94	-
30	94.67 ± 2.44	85.99 ± 2.47	96.47 ± 2.37	93.50 ± 2.79	-
59	-	-	-	-	94.44 ± 2.72

Table 4. The skin spectrum recognition accuracy (in %) on the StellarNet dataset for several spectral reflectance descriptors, with respect to the length of the feature vector. The descriptors being compared are the B-Spline descriptor [15]; the Gaussian Mixture descriptor [3]; the principal components recovered by PCA; the Daubechies 8 (D8) wavelet [1] and the raw spectra.

3.3.2 Skin Biometrics

Next, we illustrate the utility of the above representations for ethnic group identification using skin reflectance. Here, we perform this experiment on the hyperspectral image database using a length of 15 for all the descriptors. To classify the subjects in our dataset into the main ethnic groups (Caucasian, Indian and Oriental), we use training data from two typical representatives randomly selected from each ethnic group. To this end, we use frontally-illuminated images of these training subjects from each of which we select 5 rectangular skin regions with an average size of 25×17 pixels. This results in three training classes (one for each ethnic group) which we use to train three Support Vector Machines (SVM) classifiers according to a one-versus-all strategy. Each of these uses a first-degree polynomial kernel with its parameters tuned through a four-fold cross validation procedure. During the testing phase, the three SVM classifiers assign an ethnic group association probability to the skin pixels in the test image.

In Figure 5, we present the probability of pixels being associated to the three skin groups as 2D maps for a sample image. On the maps, the values of the red, green and blue channels of a pixel are proportional to the association probability of that pixel to the Caucasian, Indian and Oriental groups, respectively. As shown in the figure, the B-Spline, PCA and the raw reflectance feature produce similar skin group association maps. These three features yield a higher percentage of pixels correctly assigned to the groundtruth ethnic groups than the Gaussian Mixture and the wavelet feature. Here, the Gaussian Mixture and the wavelet descriptors suffer from confusion between the Caucasian and Oriental skin groups when being used to recognise the ethnic group of the Caucasian subject.

To elaborate further on the quantitative results of this experiment, in Table 5, we report the percentage of pixels classified into each ethnic group as a confusion matrix between ethnicities. Overall, there is no clear best performer.

Actual \ Classified		Classified		
		C	I	O
B-Spline	C	71.18	3.45	25.37
	I	2.74	72.71	24.55
	O	20.17	10.70	69.13
Gauss Mixture	C	47.59	4.88	47.53
	I	1.40	74.02	24.59
	O	11.36	12.65	75.99
PCA	C	64.57	3.78	31.65
	I	2.15	75.31	22.54
	O	15.42	11.06	73.52
Wavelet	C	62.52	5.04	32.44
	I	10.35	67.39	22.26
	O	29.37	10.36	60.26
Reflectance	C	64.78	3.79	31.43
	I	2.18	75.53	22.29
	O	15.51	11.16	73.33

Table 5. Percentage of image pixels classified into the Caucasian (C), Indian (I) and Oriental (O) ethnic groups. The diagonal elements for each sub-table spanning columns 3–5 present the overall classification rates for ethnic groups at the pixel level, whereas the off-diagonal elements show the “confusion” between these groups per descriptor under study.

While the B-Spline descriptor leads the recognition rate for the Caucasian group, the raw reflectance spectrum and the Gaussian Mixture are the best performers for the Indian and Oriental groups, respectively. On the other hand, the Daubechies wavelet feature delivers the lowest recognition rate. It is worth stressing that, despite PCA and the raw reflectance spectra deliver comparable performance to the B-Spline and Gaussian mixture representations, the latter allow spectrum reconstruction, interpolation and classification. In this regard, PCA does not allow the input data to be reconstructed while the raw spectra can be computationally costly for recognition purposes due to its full dimensionality.




Figure 5. Skin type association maps of a Caucasian subject. The red, green and blue values correspond to the likelihood that each pixel is classified into the Caucasian, Indian and Oriental groups, respectively. First column: input images rendered in pseudo colour; Remaining columns, from left to right: the skin type association maps for the B-Spline descriptor, the Gaussian Mixture, PCA, the Daubechies 8 wavelet and the raw reflectance spectrum.

4. Conclusions

In this paper, we have evaluated four spectral reflectance representations for the purposes of spectrum reconstruction, interpolation, recognition and biometrics on real-world spectrometer and image data. For purposes of reconstruction, PCA achieves the highest reconstruction accuracy, followed by the B-Splines. Note that PCA cannot be utilised for spectrum interpolation purposes. Further, we have compared the B-Spline and the Gaussian mixture representation for the spectrum interpolation task. Of these two, the B-Spline delivers better reconstruction. For skin recognition, PCA slightly outperforms the other representations, followed closely by the wavelets and the B-Splines. Finally, for biometrics, our results show that there is no clear best approach, with the wavelets delivering the worst performance amongst the representations under study.

References

- [1] P. S. Addison. *The illustrated wavelet transform handbook: introductory theory and applications in science, engineering, medicine, and finance*. IoP, 2002. 3, 7
- [2] E. Angelopoulou. Objective Colour from Multispectral Imaging. In *European Conference on Computer Vision*, pages 359–374, 2000. 3
- [3] E. Angelopoulou, R. Molana, and K. Daniilidis. Multispectral Skin Color Modeling. In *Computer Vision and Pattern Recognition*, volume 2, pages 635–642. IEEE Computer Society, 2001. 1, 2, 3, 4, 7
- [4] J. Bai, S. Xiang, and C. Pan. A Graph-Based Classification Method for Hyperspectral Images. *IEEE Transactions on Geoscience and Remote Sensing*, 51(2):803–817, 2013. 1
- [5] V. Bochko, N. Tsumura, and Y. Miyake. Spectral color imaging system for estimating spectral reflectance of paint. *Journal of Imaging Science and Technology*, 51(1):70–78, 2007. 1
- [6] G. Camps-Valls, T. V. B. Marsheva, and D. Zhou. Semi-supervised graph-based hyperspectral image classification. *IEEE T. Geoscience and Remote Sensing*, 45(10):3044–3054, 2007. 1
- [7] C.-I. Chang and S.-S. Chiang. Anomaly detection and classification for hyperspectral imagery. *IEEE Transactions on Geoscience and Remote Sensing*, 40(6):1314–1325, 2002. 1
- [8] C. Cortes and V. Vapnik. Support-Vector Networks. *Machine Learning*, 20(3):273–297, Sept. 1995. 6
- [9] I. Daubechies. *Ten Lectures on Wavelets*. Society for Industrial and Applied Mathematics, 1992. 3
- [10] C. De Boor. *A Practical Guide to Splines*. 1978. 2
- [11] R. O. Duda, P. E. Hart, and D. G. Stork. *Pattern Classification*. Wiley-Interscience, 2nd edition, 2001. 3
- [12] Dundar, M.M. and Landgrebe, D.A. Toward an Optimal Supervised Classifier for the Analysis of Hyperspectral Data. *IEEE Transactions on Geoscience and Remote Sensing*, 42(1):271–277, January 2004. 1
- [13] Z. Fu and A. Robles-Kelly. Discriminant absorption feature learning for material classification. *IEEE Transactions on Geoscience and Remote Sensing*, 49(5):1536–1556, 2011. 1
- [14] N. Gat. Imaging spectroscopy using tunable filters: a review. In *Proceedings of SPIE 4056, Wavelet Applications VII*, volume 50, pages 50–64. SPIE, 2000. 1
- [15] C. P. Huynh and A. Robles-Kelly. A NURBS-based Spectral Reflectance Descriptor with Applications in Computer Vision and Pattern Recognition. *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, 2008. 1, 2, 3, 4, 7
- [16] S. J. Kim, F. Deng, and M. S. Brown. Visual enhancement of old documents with hyperspectral imaging. *Pattern Recognition*, 44(7):1461 – 1469, 2011. 1
- [17] D. Landgrebe. Hyperspectral image data analysis. *IEEE Signal Processing Magazine*, 19, 2002. 1
- [18] E. T. Y. Lee. Choosing nodes in parametric curve interpolation. *Computer-Aided Design*, 21(6):363–370, 1989. 2
- [19] L. T. Maloney. Evaluation of linear models of surface spectral reflectance with small numbers of parameters. *J. Opt. Soc. Am. A*, 3(10):1673–1683, Oct 1986. 1
- [20] D. H. Marimont and B. A. Wandell. Linear Models of Surface and Illuminant Spectra. *Journal of The Optical Society of America A*, 9(11):1905–1913, Nov 1992. 1
- [21] D. W. Marquardt. An Algorithm for Least-Squares Estimation of Nonlinear Parameters. *SIAM Journal on Applied Mathematics*, 11(2):431–441, 1963. 2, 3
- [22] L. Piegel and W. Tiller. *The NURBS book*. Springer-Verlag, London, UK, 1995. 2
- [23] S. Yamamoto, N. Tsumura, T. Nakaguchi, and Y. Miyake. Development of a Multi-spectral Scanner using LED Array for Digital Color Proof. *Journal of Imaging Science and Technology*, 51(1):61–69, 2007. 1