



## (51) International Patent Classification:

*G06K 9/00* (2006.01) *H04N 1/62* (2006.01)  
*G06T 5/00* (2006.01)

## (21) International Application Number:

PCT/AU2012/001352

## (22) International Filing Date:

2 November 2012 (02.11.2012)

## (25) Filing Language:

English

## (26) Publication Language:

English

## (30) Priority Data:

2011904583 4 November 2011 (04.11.2011) AU

(71) Applicant: NATIONAL ICT AUSTRALIA LIMITED  
[AU/AU]; Level 5, 13 Garden Street, Eveleigh, New South  
Wales 2015 (AU).(72) Inventors: ROBLES-KELLY, Antonio; c/- Level 5, 13  
Garden Street, Eveleigh, New South Wales 2015 (AU).  
HUYNH, Cong, Phuoc; c/- Level 5, 13 Garden Street,  
Eveleigh, New South Wales 2015 (AU). GU, Lin; c/-  
Level 5, 13 Garden Street, Eveleigh, New South Wales  
2015 (AU).(74) Agent: FB RICE & CO; Level 23, 44 Market Street,  
Sydney, NSW 2000 (AU).

(81) Designated States (unless otherwise indicated, for every  
kind of national protection available): AE, AG, AL, AM,  
AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY,  
BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM,  
DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT,  
HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP,  
KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD,  
ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI,  
NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU,  
RW, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ,  
TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA,  
ZM, ZW.

(84) Designated States (unless otherwise indicated, for every  
kind of regional protection available): ARIPO (BW, GH,  
GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ,  
UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ,  
TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK,  
EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV,  
MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SL,  
TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW,  
ML, MR, NE, SN, TD, TG).

[Continued on next page]

## (54) Title: DETERMINING COLOUR VALUES IN HYPERSPECTRAL OR MULTISPECTRAL IMAGES




Fig. 1a

(57) Abstract: The disclosure concerns processing of electronic images, such as hyperspectral or multispectral images (302) to determine an output colour value for a component spectrum of a decomposed multispectral or hyperspectral image. A processor (120) first determines or accesses the component spectrum and a first colour value for the component spectrum and receives an input colour value for a pixel location in the image, wherein the component spectrum contributes to the pixel. The processor (120) then determines the output colour value for the component spectrum based on the first colour value for the component spectrum, the input colour value and a utility score and stores the output colour value for the component spectrum on a datastore (118). It is an advantage that the processor (120) receives a colour value for a pixel location but the output colour value is not exactly the received colour value but one that is determined based on a utility score.




Fig. 1b



---

**Published:**

— *with international search report (Art. 21(3))*

Title

Determining colour values in hyperspectral or multispectral images

Technical Field

- 5 The disclosure concerns processing of electronic images, such as hyperspectral or multispectral images. In particular, the invention concerns, but is not limited to, methods, software and computer systems for determining a colour value of an endmember or a material represented in an electronic image.

10 Background Art

With digital imaging replacing film-based photography processes, the accurate capture and reproduction of colours has become more challenging. The input sources, such as digital imaging sensors, as well as the output devices, such as monitors and printers, have different colour characteristics.

15

The colour vision of humans is a complex process. The light is received in the human eye by three different types of cones. Each cone has a different peak wavelength, that is the wavelength with maximal sensitivity, and different distribution of the sensitivity across the light spectrum. The brain then processes the signals from the cones to create a perception of different colours.

20

A simplified approach to colour vision and representation is the representation of the light spectrum by the three components red (R), green (G) and blue (B). Most imaging sensors produce RGB values for the captured image. However, the sensitivity of the  
25 imaging sensors is different to the sensitivity of the human eye. In addition, most output devices such as computer monitors, cannot produce arbitrary light spectra but only mixtures of the spectra of three available colours, which again have different spectra to the sensitivity of the sensor. The influence of the sensor sensitivity, the monitor colour spectrum and the perception of the human eye need to be considered in  
30 order to produce a realistic reproduction of a captured image.

35

Any discussion of documents, acts, materials, devices, articles or the like which has been included in the present specification is not to be taken as an admission that any or all of these matters form part of the prior art base or were common general knowledge  
in the field relevant to the present invention as it existed before the priority date of each claim of this application.

Throughout this specification the word "comprise", or variations such as "comprises" or "comprising", will be understood to imply the inclusion of a stated element, integer or step, or group of elements, integers or steps, but not the exclusion of any other element,  
5 integer or step, or group of elements, integers or steps.

#### Disclosure of Invention

In a first aspect there is provided a computer implemented method for determining an output colour value for a component spectrum of a decomposed multispectral or  
10 hyperspectral image, the method comprising:

determining or accessing the component spectrum and a first colour value for the component spectrum;

receiving an input colour value for a pixel location in the image, wherein the component spectrum contributes to the pixel;

15 determining the output colour value for the component spectrum based on the first colour value for the component spectrum, the input colour value and a utility score; and

storing the output colour value for the component spectrum on a datastore.

20 It is an advantage that the method receives a colour value for a pixel location but the output colour value is not exactly the received colour value but one that is determined based on a utility score.

Determining the output colour value may comprise optimising a cost function.  
25

It is an advantage that the determination of the output colour value is optimised directed by the utility score. Any changes to the output colour value that negatively affect the utility score are penalised and the resulting colour value has a higher utility.

30 The input colour value may be received from a user.

The image may be decomposed into material spectra and the component spectrum may be one of the material spectra.

35 The method may comprise

determining or accessing a decomposition of the material spectrum into endmember spectra; and

determining an output colour value for each endmember spectrum based on the output colour value for the material spectrum.

5

The image may be decomposed into endmember spectra and the component spectrum may one of the endmember spectra.

10

The method may be performed, such that an output colour value for each component spectrum represented in the image is determined.

The method may comprise determining or accessing a component contribution for the pixel location and the cost function may be based on the component contribution.

15

The utility score may be any one or more of

an aesthetic score;

a score based on a predetermined scale of danger, benefit, profit or added financial or monetary value;

enhanced visualisation of multiple images; or

20

the nature of the rendering context.

The utility score may be based on a classifier.

The classifier may be any one or more of

25

a quadratic discriminant;

a support vector machine;

a linear discriminant; or

a logistic regressor.

30

The classifier may be based on one or more of:

abundance of material spectra in the image;

most abundant component spectrum;

second most abundant component spectrum;

rule of thirds;

35

whole image average;

macro feature;

domain knowledge; or  
rules and regulations pertaining the utility measure.

The utility score may be based on the contrast of the image.

5

The utility score may be based on the hazardousness of the component spectra.

The cost function may be based on a utility score for each component spectrum.

10 The method may comprise accessing a library of component spectra.

The method may comprise determining or accessing illuminant spectra and illuminant contributions.

15 The method may comprise determining a colour value for each illuminant spectrum.

The method may comprise determining a colour value for each combination of component spectra and illuminant spectra.

20 The method may comprise determining for each pixel location a pixel colour value based on the output colour value for each of the component spectra and the component contributions.

25 It is an advantage that a colour image is created for the decomposed multispectral or hyperspectral image.

The method may comprise displaying to a user an image based on the pixel colour values.

30 The method may comprise initialising the first colour value for the component spectrum with a default colour.

In a second aspect there is provided software, that when installed on a computer causes the computer to perform the method of any one or more of the claims 1 to 22.

35

In a third aspect there is provided a computer system for determining an output colour value for a component spectrum of a decomposed multispectral or hyperspectral image, the system comprising:

- 5 a processor to determine or access the component spectrum and a first colour value for the component spectrum; to receive an input colour value for a pixel location in the image wherein the component spectrum contributes to the pixel; to determine the output colour value for the component spectrum based on the first colour value for the component spectrum, the input colour value and a utility score; and to store the output colour value for the component spectrum on a datastore.

10

The computer system may comprise a port to receive wavelength indexed multispectral image data and a storage to store the image data, the component spectrum and the output colour value.

- 15 Optional features described of any aspect, where appropriate, similarly apply to the other aspects also described here.

#### Brief Description of Drawings

An example will be described with reference to

- 20 Fig. 1a illustrates a computer system for determining a colour profile.  
Fig 1b illustrates contributions to the radiance.  
Fig. 2 illustrates a method for determining a colour profile.  
Fig. 3 illustrates a data structure for the multispectral image data.  
Figs. 4a and 4b illustrate a graphical user interface.

25

#### Best Mode for Carrying Out the Invention

- Fig. 1a illustrates a computer system 100 for determining a colour value for each component of a decomposed image of a scene 101. The colour value may be an RGB value, such as (255, 0, 0) for red, or an HSV value. In one example, the component is a material, that is the scene is decomposed into a linear combination of materials and illuminants and the computer system 100 determines a colour value for each material and for each illuminant. In a different example the component is an endmember, that is each material is further decomposed into a linear combination of endmembers and the computer system 100 determines a colour for each endmember instead of each material.

35

The difference between materials and endmembers becomes relevant where there are a number of materials and these materials can be expressed as endmembers. For instance, a set of materials, such as pills, of different kinds that can be expressed as a combination of endmembers, such as vitamins, minerals, antibiotics. The computer system 100 assigns a mapping per material (pill), that is consistent with that yielded by the linear combination of the chemicals in the pill (vitamins, minerals, etc.). The computer system 100 defines, in one example, a mapping per material or, in another example, a mapping per endmember and then express the material as a linear contribution of these endmembers.

In a different example the computer system 100 determines a colour value for each combination of materials or endmembers and illuminants.

In the example of Fig. 1 the scene comprises a rose 102 and a flower pot 103 illuminated by the sun 104 having illuminant spectrum 105. The term reflectance generally refer to the fraction of incident electromagnetic power that is reflected. Object reflectance, material reflectance and endmember reflectance refer to the fraction of power reflected from an object, a material and an endmember, respectively. Since the reflectance is dependent on the wavelength, each reflectance is represented by a reflectance spectrum.

Throughout this specification, unless stated otherwise, where reflectance spectrum is used in singular form it refers to object reflectance and where reflectance spectra is used, that is in plural form, it refers to material or endmember spectra. In the case of the plural form, typically either of material or endmember spectra may be chosen to perform that part of the method.

The objects in the scene 101 comprise multiple different materials and each material has a material reflectance spectrum. Two material reflectance spectra 106 and 107 for the blossom and the pot respectively are depicted.

Computer system 100 comprises a sensor 112 and a computer 114. In this example the sensor 112 is a hyperspectral or multispectral sensor that is able to capture an image of the scene 101. The computer 114 comprises a program memory 118(a), a data memory 118(b) and a processor 120.



When the light from the illuminant 104 hits the objects, the illuminant spectrum 105 is multiplied by the respective material reflectance spectra 106 and 107. The resulting spectra are superposed and reach the sensor 102 as a radiance spectrum 108. The sensor has a number of pixels, such as one million, and captures for each pixel location  
5 a separate sampled version of the radiance spectrum.

In one example the computer system is integrated into a handheld device such as a consumer camera and the scene 101 may be any scene on the earth, such as a tourist attraction or a person. The scene may also be illuminated by multiple artificial light  
10 sources, such as fluorescent lamps. The sensor 112 may have a number of bands that balances computational costs with accuracy. The sensor 112 may have as low as four bands and as high as hundreds.

The computer 114 receives images from the sensor 112, which are stored in local data  
15 memory 118(b) by processor 120. The processor 120 uses software stored in program memory 118(a) to perform the method shown in Fig. 2. In this sense the processor 120 performs the method of determining colour values for endmembers or materials and illuminants. The processor 120 determines material reflectance spectra and illumination spectra and decomposes the images into a combination of the material  
20 reflectance spectra weighted by material association probabilities. As mentioned above, in some examples, the computer further decomposes the images into endmember spectra weighted by endmember association probabilities.

In another example, the processor accesses predefined material or endmember  
25 reflectance spectra and predefined illumination spectra from a material or endmember library and an illumination library, respectively. These libraries are stored on data memory 118(b). In one example the endmember library consists of 297 reflectance spectra acquired using a Stellar-Net spectrometer. The endmember spectra are sampled at a spectral resolution of 1nm in the interval [430nm, 720nm], that is the visible range,  
30 and comprise nine categories, including cloth, different kinds of paint, human skin, leaves, metals, papers of different colours, plastic, porcelain and wood types.

The canonical illuminant library is acquired with similar settings and consists of four light sources (tungsten, incandescent, fluorescent and sunlight). The spectra for the end  
35 members and canonical light sources are normalised so that the integrated reflectance across the visible range is unity. The local storage 118(b) stores 77 multispectral

images representing three categories (landscapes, people and still life) acquired under a wide variety of lighting conditions spanning from outdoor settings to indoor locations. The imagery is composed of spectra sampled at intervals of 10 nm in the range between 430 nm and 650 nm, that is 23 bands.

5

To normalise the spectral images and the reflectance spectra to the same spectral resolution, the B-Spline representation in [4] is employed to re-interpolate the endmember reflectance spectra and canonical light power spectra at 27, 36, 45, 54, 63 and 72 equally spaced wavelengths in the 430-720nm region, yielding libraries of spectra sampled at between 4nm and 11nm steps. Next, from these libraries are selected a number of endmembers equal to the number of sampling bands. From the 297 spectra in the initial library, those are selected that yield the lowest reconstruction error following a down-sampling operation on them so as to arrive at the same spectral resolution as the spectral images. For the imagery, the pixel spectra are up-sampled to achieve the spectral resolution corresponding to the endmember library size. Note that this is important since otherwise the problem of recovering mixture coefficients may become under-constrained and has trivial or non-unique solutions for  $\rho_{m,v}$ .

10

15

A reconstruction error is the error between the pseudocolour imagery computed using the colorimetric standard and that yielded by the method described below. This error is monotonically decreasing with respect to the number of endmember spectra used. As the endmember library increases in size, the reduction in error becomes less apparent. This suggests a trade-off between complexity and performance. In one example the endmember library comprises 54 spectra and the number of end members per object material is 4.

20

25

It is noted that no libraries are necessary for the method to perform. If no libraries are available, the proposed method determines the reflectance spectra and illumination spectra as described further below. The result is a decomposed representation comprising reflectance spectra, that is material or endmember reflectance spectra, illumination spectra and respective contributions for each pixel.

30

The processor 120 stores the decomposed representation instead of or in addition to the originally stored image data in data memory 118(b). For storing the illumination spectrum, the endmembers or any other spectrum the computer may employ the method

35

described in WO 2009/152583 "Compact Representation of a Reflectance Spectrum" which is incorporated herein by reference.

5 The software provides a user interface 122 that can be presented to the user on a touch screen 122. In most examples, the user interface 122 accepts RGB data from the processor 120 to be displayed to the user. However, the proposed method is not limited to a certain number of colour channels and the acronym RGB may be replaced by any number of different colour channels accepted by the user interface 122. The processor 120 converts the decomposed image data into RGB values and sends this data to the  
10 user interface 122. The user interface 122 displays to the user the RGB image data. The processor 120 may access a default colour profile, that is a mapping of the decomposed image data to RGB colour values, stored on memory 118(b). The user interface accepts input from the user, such as colour values for certain pixel locations. Upon receiving such user input, the processor 120 optimises the colour profile,  
15 converts the image data to RGB using the updated colour profile and sends the RGB data to the user interface 112.

In this example the memory 118(b) is local to the computer 114, but alternatively could be remote to the computer 114.

20 The processor 120 solves the following problem. A multispectral or hyperspectral image with the spatial domain  $\mathcal{I}$  consisting of wavelength-indexed radiance spectra in the visible spectrum is captured by sensor 112. The aim of computation is the recovery of a colour image which corresponds to the input imagery where the RGB colour of the materials in the scene is determined by user input. This is achieved using two libraries.  
25 The first of these is given by the set  $S$  of primordial materials or endmembers which account for object materials in the scene. These end members can be any man-made or natural materials such as plastic, metals and skin.

30 The second library includes "canonical" illuminant power spectra corresponding to light sources which commonly occur in real world environments. Similarly to the endmembers above, these are lights that can be viewed as widespread in natural and man-made environments, such as sunlight, neon tubes and tungsten lighting. To each of the end members or canonical light sources corresponds a colour mapping function  
35 that yields a tristimulus RGB value for each of the entries in the libraries.

To recover the colour mappings and to provide a means to their computation the processor employs the dichromatic model introduced by Shafer [7]. The model decomposes the scene radiance  $I(\lambda, v)$  at pixel  $v$  and wavelength  $\lambda$  into a diffuse and a specular component as follows.

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

where  $L(\cdot)$  is the illuminant power spectrum,  $S(\lambda, v)$  is the object reflectance, the shading factor  $g(v)$  governs the proportion of diffuse light reflected from the object and the factor  $k(v)$  models surface irregularities that cause specularities in the scene.

Once the image data is decomposed into endmembers  $S$  and illuminant spectra  $L$ , equation (1) is used to generate the RGB values from the decomposed image data. The image radiance is expressed using the mixture coefficient  $\rho_{m,v}$  of the material  $m$  in the composition of the pixel  $v$  and the mixture coefficient  $\alpha_l$  of the canonical illuminant  $l$ , as follows:

$$I_c(v) = \left( \sum_{\ell=1}^M \alpha_\ell g_c(L_\ell(\cdot)) \right) \left( g(v) \sum_{m=1}^N \rho_{m,v} f_c(S_m(\cdot)) + k(v) \right) \quad (2)$$

where  $c \in \{R, G, B\}$   $S_m(\cdot)$  is the vectorial representation of the reflectance spectrum of the endmember with the index  $m$  in the material library  $S$ ,  $L_\ell$  is the power spectrum of the  $\ell$ -th entry in the library of  $M$  canonical illuminants and  $g_c$  and  $f_c$  are the colour mapping functions that assign a tristimulus to each combination of canonical illuminant power and end member reflectance.  $g_c$  and  $f_c$  together are referred to as colour profile.

Fig. 1b illustrates the contribution of the illumination spectrum 105 and the endmember spectra 106 and 107 to the radiance in equation (2).

In a different example, equation (2) is replaced by

$$I_c(v) = g(v) \sum_l \sum_m \alpha_l \rho_{m,v} g_c(L_l(\cdot)) f_c(S_m(\cdot)) + k(v) \sum_l \alpha_l g_c(L_l(\cdot)) \quad (3)$$

which can be further expressed as follows

$$I_c(v) = g(v) \sum_l \sum_m \beta_{l,m,v} \mu_c(L_l(\cdot), S_m(\cdot)) + k(v) \sum_l \alpha_l g_c(L_l(\cdot)) \quad (4)$$

where

$$\beta_{m,v,l} = \alpha_l \rho_{m,v}$$

$$\mu_c(L_l(\cdot), S_m(\cdot)) = g_c(L_l(\cdot)) f_c(S_m(\cdot))$$

The difference between equations (2) and (4) is that in equation (2) each endmember and each illuminant has one individual colour mapping function. In equation (4) each combination of endmembers and illuminants has one colour mapping function.

5 Fig. 2 illustrates a method 200 for determining a colour profile. The first step is to retrieve 202 wavelength indexed hyperspectral or multispectral image data of a scene. In this example, the image data is retrieved from the memory 118(b). In other examples, the image data is received directly from the sensor 112. The image data is represented by pixel location that can be one or more individual pixels and for each  
10 pixel the data contains a number of radiance values, one radiance value for each wavelength of the light spectrum. In one example, the light spectrum covers the visible spectrum (430-720nm). In another example, the light spectrum also covers wavelengths that are not visible to the human eye, such as infrared light (650-990nm). The different wavelengths may be equally spaced, such as by 10nm, or may be closer to  
15 each other over parts of the spectrum that are more significant for the discrimination of different materials. For instance, the spectrum of Chlorophyll has a distinctive peak in the green light spectrum.

Fig. 3 illustrates a data structure 301 for the multispectral image data. The data  
20 structure comprises layers, one for each wavelength. Each layer represents the radiance values for one wavelength. One pixel 302 is highlighted. The values of pixel 302 for different wavelengths, that is the radiance values from lower layers at the same location as pixel 302, represent a sampled radiance spectrum also referred to as the image spectrum. This image spectrum is a mixture of the illumination spectrum and the  
25 reflectance spectra of the different materials present in the part of the scene that is covered by pixel 302. It is noted here that some pixel areas cover only one single material while other pixel areas cover parts of the scene with multiple materials. Further, each material may be a mixture of multiple endmembers.

30 In one example spectral data is acquired by heterogeneous sensing devices, which vary in spectral resolution and range. For instance, spectrometers acquire hundreds of bands in the visible and infrared regions, whereas hyperspectral and multispectral imagers are limited to a much coarser spectral resolution due to their optical design. In this example, the endmember library is collected by a spectrometer with a high spectral  
35 resolution, while the spectral imagery is acquired by a hyperspectral camera with a lower spectral resolution. To overcome this barrier, it is necessary to re-interpolate

spectral data originally collected from these different sources so as to provide consistent spectral resolutions and ranges

Furthermore, operations such as an object material recovery process require spectra of equal lengths and resolutions. In some examples representations of reflectance spectra include basis functions. These are, in essence, spectral descriptors which allow for the interpolation of spectral data by employing either a B-Spline [4] or a Gaussian Mixture representation [1]. In one example the B-Spline representation is employed to normalise the length and resolution of image and endmember spectra due to its superior performance to the Gaussian Mixture representation for material classification tasks [4].

In order to prepare the data for the later steps, the method 200 determines 204 the illumination spectrum. The method for determining the illumination spectrum is described in WO 2011/026167, which is incorporated herein by reference. This method determines the optical invariants of equation (1) including the illumination spectrum  $L$  and object reflectance spectrum  $S$ .

Once the reflectance spectrum  $S$  is known, the method determines 206 the material or endmember spectra and the respective contributions. The algorithm for endmember extraction is described in Australian application AU 2011211336, which is incorporated herein by reference.

To recover colour imagery dependent on user-defined preferences, the heterogeneous nature of the scene is dealt with by imposing consistency over object materials.

In some examples, there are variations in the actual reflectance of pixels belonging to the same material in the scene. This does not necessarily indicate a change of the object material under consideration rather than small variations in composition. For instance, the spatial variations on the enamel of a mug entirely made of porcelain may result in changes in the endmember association to each pixel and adversely affect the colour finally produced. Hence, it would be rather undesirable to have objects partitioned or fragmented into inconsistent endmembers as the colour image is produced. It is the aim to impose endmember consistency between pixels sharing the same material in the scene using the reflectance spectra obtained by the algorithm in WO 2011/026167. Instead of solving the problem at the pixel level, the problem is

extended to material clusters, where the endmember decomposition occurs per material rather than per pixel.

The method 200 imposes consistency of endmember composition by recovering their mixture coefficients for each material from the mean reflectance spectrum of the material. The method 200 assigns a partial membership  $P(\omega|v)$  of a material cluster  $\omega$  with mean reflectance  $S_\omega(\cdot)$  to each pixel  $v$  in the image. Taking the set of all the material clusters  $\Omega$  in the scene into account, the recovery of the mixture coefficients comprises two successive optimisation problems. The first problem considers the clustering of image pixels based on their spectral reflectance. Here, the method 200 employs an affinity metric  $a(\omega, v)$  to preselect  $k$  endmembers per material. This is done using the affinity between the pixel reflectance spectrum  $S(\cdot, v)$  and the mean spectrum  $S_\omega(\cdot)$  of a material cluster. Mathematically, this affinity measure can be defined by their Euclidean angle

$$a(\omega, v) = 1 - \frac{\langle S(\cdot, v), S_\omega(\cdot) \rangle}{\|S(\cdot, v)\| \|S_\omega(\cdot)\|} \quad (5)$$

With this metric, the first optimisation problem aims to minimise the total expected affinity for the entire image as

$$A_{Total} = \sum_{v \in \mathcal{I}, \omega \in \Omega} P(\omega|v) a(\omega, v) \quad (6)$$

subject to the law of total probability  $\sum_{\omega \in \Omega} P(\omega|v) = 1 \forall v \in \mathcal{I}$ .

Since the formulation in equation (6) often favours hard assignment of pixels to their closest materials, the problem is restated subject to the maximum entropy criterion [5]. The entropy of the material association probability distribution at each pixel is defined as

$$\mathcal{P} = - \sum_{v \in \mathcal{I}} \sum_{\omega \in \Omega} P(\omega|v) \log P(\omega|v) \quad (7)$$

With the affinity metric in equation (5), the problem becomes that of finding a set of object material spectra and a distribution of material association probabilities  $P(\omega, v)$  for each pixel  $v$  so as to minimise  $A_{Entropy} = A_{Total} - \mathcal{L}$ .

where

$$\mathcal{L} = T\mathcal{P} + \sum_{v \in \mathcal{I}} \varrho(v) \left( \sum_{\omega \in \Omega} P(\omega|v) - 1 \right) \quad (8)$$

5

in which  $T \geq 0$  and  $\varrho(v)$  are Lagrange multipliers. Note that  $T$  weighs the level of randomness of the material association probabilities whereas  $\varrho(v)$  enforces the total probability constraint for every image pixel  $v$ .

- 10 The optimisation approach to the problem is somewhat similar to an annealing soft-clustering process. At the beginning, this process is initialised assuming all the image pixels are made of the same material. As the method progresses, the set of materials grows. This, in essence, constitutes several “phase transitions”, at which new materials arise from the existing ones. This phenomenon is due to the discrepancy in the affinity  
15  $a(\omega, v)$  between the pixel reflectance  $S(., v)$  and the material reflectance spectrum  $S_\omega(.)$ .

The method 200 then derives the optimal set of scene materials so as to minimise the cost function above. The method 200 computes the derivatives of  $A_{Entropy}$  with respect to the material reflectance spectrum  $\omega(.)$  and equates it to zero, which yields

20

$$S_\omega(.) \propto \sum_{v \in \mathcal{I}} P(\omega|v) \frac{S(., v)}{\|S(., v)\|} \quad (9)$$

- In Equation (9), the probability  $P(\omega|v)$  is required. To compute this probability, the  
25 method 200 employs deterministic annealing. A major advantage of the deterministic annealing approach is that it avoids being attracted to local minima. In addition, deterministic annealing converges faster than stochastic or simulated annealing [6].

- The deterministic annealing approach casts the Lagrangian multiplier  $T$  as a system  
30 temperature. In one example the initial and terminal temperatures for the deterministic annealing were set to 0.02 and 0.00025, respectively. At each phase of the annealing process, where the temperature  $T$  is kept constant, the algorithm proceeds as two interleaved minimisation steps so as to arrive at an equilibrium state. These two minimisation steps are performed alternately with respect to the material association  
35 probabilities and the end members.



For the recovery of the pixel-to-material association probabilities, the method fixes the material reflectance spectrum  $S_\omega(\cdot)$  and seeks for the probability distribution which minimises the cost function  $A_{Entropy}$ . This is achieved by setting the partial derivative of  $A_{Entropy}$  with respect to  $P(\omega | v)$  to zero. Since  $\sum_{\omega \in \Omega} P(\omega | v) = 1$ , it can be shown that

5 the optimal material association probability for a fixed material set  $\Omega$  is given by the Gibbs distribution

$$P(\omega | v) = \frac{\exp\left(\frac{-a(\omega, v)}{T}\right)}{\sum_{\omega' \in \Omega} \exp\left(\frac{-a(\omega', v)}{T}\right)} \forall \omega, v \quad (10)$$

10

The second problem is concerned with the decomposition of the mean material reflectance spectra into endmember proportions. With the material reflectance  $S_\omega(\cdot)$  in

15 hand, the method 200 recovers the endmember proportions  $\rho_{m,v}$  contributing to the composition of each material  $\omega$ . The problem is formulated as a minimisation of the cost function

$$\epsilon_\omega = \int_V \left( S_\omega(\lambda) - \sum_{m=1}^N \rho_{m,\omega} S_m(\lambda) \right)^2 d\lambda \quad (11)$$

20

subject to the constraints  $\rho_{m,v} \geq 0$  and  $\sum_{m=1}^N \rho_{m,\omega} = 1 \forall \omega \in \Omega$ .

The cost function in Equation 10 can be minimised by a least-squares solver. The

25 weight of the endmember  $m$  for pixel  $v$  is then given by

$$\rho_{m,v} = \sum_{\omega \in \Omega} P(\omega | v) \rho_{m,\omega} \quad (12)$$

With the material clusters obtained, the method 200 proceeds to the recovery of the colour values. In one example the user can choose pixels in the image  $\mathcal{I}$  which the

30 user would like to edit and assign a colour to the corresponding material. In another example the problem of recovering the relevant user profile becomes that of finding a strategy to update the function  $f_c(\cdot)$  for each endmember so as to recover the colours of the materials under consideration. Another part to this problem is to update the function  $g_c(\cdot)$  for each illuminant. This update is reflected in the temperature of the

35 light and relates to the colour constancy problem in computer vision.

In a different example the problem is to update the function  $\mu(L_i(\cdot), S_m(\cdot))$  from equation (4) for each combination of illumination spectra and endmembers, such as rose petals in daylight. In this example, for four illuminants and 54 endmembers there are  $4 * 54 = 216$  colour mapping functions to determine.

Fig. 4a illustrates a graphical user interface (GUI) 400 showing a scene 410. The scene 410 comprises a rose 412 having petals 414 and a pot 416. The GUI 400 comprises a colour picker tool 420 and a colour palette 430. In this example a first colour 432 and a second colour 434 are used.

In the scene 410, the petals 414 and the pot 416 both have the same first colour 432 while the material of the petals 414 and the pot 416 are different. In order to display a colour image from the multispectral image data the method initialises a user colour profile with the default human colour matching function studied by Stiles and Burch [8]. This allows the colour  $I_c(v)$  to be available at all the pixels. In this example the first colour 432 is red. The user wishes to replace the colour of the petals 414 by a second colour 434. In this example the second colour 434 is blue. The user selects the second colour 434 and then directs the colour picker tool 420 over the petals 414 using an input device, such as a computer mouse or a touch screen. Once the tip of the colour picker tool 420 is located at a location of a pixel belonging to the petals 414, the user selects this pixel by depressing a mouse button, the enter button on a keyboard or tapping the touch screen.

Fig. 4b illustrates the desired outcome after the method 200 has finished. The colour of the petals 414 is now the second colour 434 while the colour of the pot 416 is still the first colour 432. This means, the method 200 has applied the second colour 434 only to the material of the petals 414 and not to the material of the pot 416.

Once a user finishes editing, the system computes the dichromatic model parameters and the proportions of endmembers and updates the colour mapping functions as described later. The application development hinges on the underlying assumption that each user has a particular colour profile. This is in general true, since trichromatic cameras fix the material-colour relationship making use of sensor-dependent colour

matching functions. Thus, the proposed method provides a means for the user to create and update a profile or profiles of his/her preferences.

To compute the colour mapping functions from user input, the method 200 requires the user-provided colour at a number of image pixels. In Equation (4) the pixel colour is a linear combination of  $MN + M$  terms involving functions  $f_c(\cdot)$  and  $g_c(\cdot)$ . Therefore, to recover the colour values  $C'_m$ , at least  $MN + M$  pixel spectra are required to solve for the endmembers and canonical light sources when the user edits the material colours. The colour picker tool 420 allows the user to edit the colour of a single pixel at a time. Therefore, the system is designed to automatically select more pixels sharing the same material as that of the selected one using a similarity threshold. This similarity measure between two pixels  $u$  and  $v$  is reminiscent of that used in bilateral filtering, which is given by

$$d(u, v) = \exp \left\{ - \left( \frac{\angle(S(\cdot, u), S(\cdot, v))}{\sigma_r} \right)^2 - \left( \frac{\|u - v\|}{\sigma_s} \right)^2 \right\}$$

where  $\angle(S(\cdot, u), S(\cdot, v))$  is the Euclidean angle between the reflectance spectra,  $\|u - v\|$  is the spatial distance between  $u$  and  $v$ , and  $\sigma_r$  and  $\sigma_s$  signify the kernel widths in the spectral reflectance and spatial domains.

Once the user has selected a user colour and a particular pixel location the GUI 400 sends this information to the method 200 which receives 206 the information.

A straightforward solution for determining the colour mapping functions would be to solve a linear system of equations so as to recover the colour mapping functions based on, for instance, a least squares criterion. The main drawback of this treatment is that given a set of hyperspectral images, when the user edits the material in one of them, such an action will change the appearance of the other imagery.

Instead, the method 200 moderates the effect of changes in the colour of library endmembers for the material being edited. This is done using an utility score, such as an aesthetic score, of the rest of the imagery. The idea is to use the aesthetic qualities of some features in the images under study to recover the colour mapping functions such that the material colour is as close as possible to that selected by the user on the

image being edited while avoiding changes that would make the colour of the rest of the available images less aesthetically appealing.

In one example the work of Birkhoff is used to further improve the results. In a different example “colourfulness” is used, that is a measure that includes colour contrast. In another example the aesthetic score is a more general utility measure that allows the visualisation to go beyond aesthetics. For example, the utility measure may be about how hazardous a substance is (red for very dangerous, green for innocuous). In different examples the utility score is a score based on benefit, profit, added financial or monetary value, enhanced visualisation of multiple images or the nature of the rendering context. This would allow different colour profiles to be added so as to provide the best colour match possible.

In a different example, eight features introduced in [3] are used to relate the update of the colour mapping functions to the aesthetic qualities of the imagery. In [3], the authors classified images in terms of aesthetic quality by obtaining the top 15 out of 56 features which yield the highest classification accuracy. Following this approach, method 200 uses the eight features in their system which operate on HSV values (hue, saturation and value). Since these features are based on the HSV values for the image, they can be expressed in terms of the endmember library making use of a set of image-wise parameters which depend on the mixture coefficients  $\rho_{m,v}$ . These features are computed as given by the formulae in Table 1.

Most abundant material	$\xi_{1,m} = \frac{\sum_{v \in \mathcal{I}} \rho_{m,v} \delta_{v,1}}{\sum_{v \in \mathcal{I}} \delta_{v,1}}$
2 <sup>nd</sup> -most abundant material	$\xi_{2,m} = \frac{\sum_{v \in \mathcal{I}} \rho_{m,v} \delta_{v,2}}{\sum_{v \in \mathcal{I}} \delta_{v,2}}$
Rule of Thirds	$\xi_{3,m} = \frac{h}{9} \sum_{v \in \mathcal{I}_{thirds}} \rho_{m,v}$
Whole image average	$\xi_{4,m} = h \sum_{v \in \mathcal{I}} \rho_{m,v}$
Macro Feature	$\xi_{5,m} = \frac{\sum_{v \in \mathcal{I}_{inner}} \rho_{m,v}}{\sum_{v \in \mathcal{I}_{rest}} \rho_{m,v}}$

Table 1. Aesthetic features per material in each image.

It is noted that feature  $\xi_{i,m}$  may relate to a specific material or endmember  $m$ . The features are computed making use of the mixture coefficients  $\rho_{m,v}$  and the function  $\delta_{v,i}$  which is unity if and only if the  $i$ -th most abundant material in the image has a non-zero mixture coefficient in the composition of the pixel  $v$ . Making use of these features, in

- one example the method 200 determines the following feature vector:  
 $\varphi_{\mathcal{I},m} = \left[ [\xi_{1,m}, \xi_{2,m}, \xi_{3,m}, \xi_{4,m}, \xi_{5,m}] \times S_m, [\xi_{1,m}, \xi_{4,m}, \xi_{5,m}] \times V_m \right]^T$  where  $S_m$  and  $V_m$  are the  
 saturation and the brightness components of the HSV values equivalent to the RGB  
 responses yielded by the colour mapping function  $f_c(S_m(\cdot))$  for the endmember  $m$ . As a  
 5 result, there is a close relationship between the colour matching function  $f_c(S_m(\cdot))$  and  
 the parameters  $S_m$  and  $V_m$ .

This relationship is significant since the cost function below optimises the aesthetic  
 score. During this optimisation the aesthetic features  $\xi_1$  to  $\xi_5$  do not change but  $S_m$  and  
 10  $V_m$  are updated to achieve a lower cost. As a result of the close relationship, the colour  
 mapping function  $f_c(S_m(\cdot))$  is also updated and the RGB image generated by applying  
 $f_c(S_m(\cdot))$  also changes at each optimisation step. The optimisation procedure will be  
 explained in detail below.

- 15 It is noted that the features  $\xi_2$  and  $\xi_3$  are not present in the second element of the  
 aesthetic score. This makes the computation more efficient with negligible influence  
 on accuracy. In a different example the features  $\xi_2$  and  $\xi_3$  are also present in the second  
 element of the aesthetic score.

- 20 The omission of the hue component from the feature is a direct development from [3]  
 where hue does not appear in the top 15 features as a result of five-fold cross  
 validation. As a result, in this example the optimisation approach leaves the hue to be  
 edited by the user while allowing for variations of saturation and brightness.

- 25 One possible interpretation of the above feature vector is as follows. The features in  
 Table 1, from top to bottom, correspond to the weighted average of the two most  
 abundant materials in the image, the rule of thirds (the centre square  $\mathcal{I}_{thirds}$  of the image  
 when divided into a lattice of nine regions), the mean image-wise mixture coefficient  
 and the macro weighted average per image. Here, following [3] the macro weighted  
 30 average is computed by dividing the image into a lattice of 16 regions so as to recover  
 the four centre squares denoted as  $\mathcal{I}_{inner}$  and the outer region denoted as  $\mathcal{I}_{rest}$  such that  
 $\mathcal{I} = \mathcal{I}_{inner} \cup \mathcal{I}_{rest}$ . In the table,  $h$  acts as a bandwidth parameter given by the  
 quantity  $\frac{1}{|\mathcal{I}|}$ .

Other exemplary features are abundance of materials in the image, most abundant component, second most abundant component, rule of thirds, whole image average, macro feature, domain knowledge, or rules and regulations pertaining the utility measure;

5

To express the aesthetics of images in terms of the saturation, the brightness and the mixture coefficients of endmembers, the score  $\gamma_{\mathcal{I},m}$  is defined for each image provided as input to the user. This score is computed making use of the set  $\Upsilon$  of images which will be affected in aesthetic terms by the colour mapping function update, although not  
10 being edited directly.

In one example the user is editing an image  $\mathcal{I}_1$ . If  $\mathcal{I}_1$  is part of a greater set  $\Upsilon$  of images stored, such as on a hard drive, the method will try to compute the profile that best matches the edits by the user while maximising the aesthetic score of the entire  
15 imagery  $\Upsilon$ . The aim is to allow the user to modify the image while preventing frivolous changes to the colour mapping functions. In another example the user edits several images at once. In yet another example hard constraints are imposed on some endmembers not to change or even weights assigned to some of them to control how much these endmembers are modified with respect to their initial values.

20

In one example the method 200 determines 208 the score of each endmember  $m$  in the image  $\mathcal{I}$  by using a support vector machine, a linear discriminant, or a logistic regressor.

25

In another example the method 200 determines 208 the score of each endmember  $m$  in the image  $\mathcal{I}$  by using a classifier, such as a quadratic discriminant as follows:

$$\gamma_{\mathcal{I},m} = K_m + \varphi_{\mathcal{I},m}^T \mathbf{W}_m + \varphi_{\mathcal{I},m}^T \mathbf{Q}_m \varphi_{\mathcal{I},m} \quad (13)$$

30

where the matrix  $\mathbf{Q}_m$ , the vector  $\mathbf{W}_m$  and the constant  $K_m$  correspond to the endmember  $m$  in the library. These parameters can be recovered making use of quadratic discriminant analysis, which relaxes the constraint that both high and low aesthetic quality image groups have a common covariance structure, as compared to linear discriminant analysis.

35

The quadratic discriminant analysis, that is the training of the quadratic discriminant model of equation (13) is performed based on the Photo.net database following [3]. Photo.net is an online photo sharing community where each photo has a score that is predetermined by its members rating photos in the database from one to seven, while higher score indicates better aesthetic quality. For the recovery of the matrix  $\mathbf{Q}_m$ , the vector  $\mathbf{W}_m$  and the scalar  $K_m$ , features for the photos in the database are extracted and the images are separated into high and low aesthetic quality using a cut-off rating of  $5.0 \pm \frac{t}{2}$  for the images of high and low aesthetic quality, where  $t$  is a parameter chosen by cross validation so as to maximise the accuracy of the quadratic estimator on the Photo.net database.

$\mathbf{Q}_m$ ,  $\mathbf{W}_m$  and  $K_m$  are found by optimising a discriminant function with respect to the top 15 aesthetic features used in [3]. The purpose of this optimisation is to discriminate good and bad looking pictures. The discriminant function relies on RGB features from [3] rather than an endmember database. Therefore, these features are different from the aesthetic score in Table 1. The features used for the classifier are based on the saturation and brightness for the RGB responses, which allows all the training to be done in RGB and the testing to be effected using the RGB values for the endmembers.

Given the aesthetic score of the imagery, method 200 proceeds with the recovery of the colour mapping function for the endmembers in an optimisation setting. The method determines the value of the following cost function:

$$\epsilon_{fc} = \sum_{\mathcal{I} \in \Upsilon} \sum_{v \in \mathcal{I}} \left( I_c(v) - \sum_{m \in \mathcal{S}} \rho_{m,vc} (S_m(\cdot)) \right)^2 + \tau \sum_{\mathcal{I} \in \Upsilon} \sum_{m \in \mathcal{S}} \exp(-\gamma_{\mathcal{I},m}) \quad (14)$$

The aim is to minimise this cost function. In the above equation,  $I_c(v)$  is the value of the  $c$ -component at the pixel  $v$  as provided by the user while editing the image  $\mathcal{I}$ . In the example described with reference to Fig. 4a the colour of all pixels is initialised using a default mapping function, such as a profile defined by the CIE standards, and the second sum in equation (14) iterates over all pixels of the image.

In a different example, the colour values of only a subset  $\mathcal{J}$  of pixels is available, that is the subset for which the user has specified a colour. In this example the second sum iterates over the subset  $\mathcal{J}$  instead of the entire image  $\mathcal{I}$ .

5 In yet another example, equation (14) is not expressed in terms of endmembers but in terms of materials. In that example, the contributions  $p$  are replaced by  $P(\omega|v)$  and the endmember spectra  $S_m(\cdot)$  are replaced by the material spectra  $S_{\omega}(\cdot)$ . Further, the score  $\gamma_m$  is determined based on the individual scores of the materials rather than the endmembers.

10 Since the score  $\gamma_m$  is a quadratic function of the material HSV values  $S_m$  and  $V_m$ , the user-provided pixel colour  $I_c(v)$  and the mapping  $f_c(S_m(\cdot))$  are converted to the HSV colour space.

15 The user edits the RGB value itself, which is much more natural than editing in the HSV space. The method 200 uses the HSV colour space to compute the aesthetic score used for the optimisation. In equation (14), the method 200 finds the closest RGB colour mapping  $f_c(S_m(\cdot))$  for the end members  $S_m(\cdot)$  in the image such that the aesthetic score is maximised. The score itself is computed using the features in HSV, but this is  
20 not a problem since equation (14) uses the score yielded by the classifier in equation (13), not the features themselves.

Furthermore,  $\tau$  is a parameter that controls the contribution of the aesthetic score  $\gamma_{I,m}$  in the optimisation process. In one example the value of  $\tau$  is 5. It is noted that, in the  
25 equation above, there is used the exponential as a monotonically decreasing function whose value tends to zero as the score of the imagery increases. This has the effect of penalising choices of  $f_c(S_m(\cdot))$  which yield poor aesthetic results. In one example a quasi-Newton method [2] is employed in order to minimise the cost function in equation (14). The quasi-Newton method is an iterative method embedded into method  
30 200. In each step the method 200 determines 210 mapping functions  $f_c(S_m(\cdot))$  with lower cost than the current solution.

Method 200 then checks 212 whether a termination condition is met, such as exceeding a maximum number of iterations or staying below a minimum cost difference  
35 compared to the previous step. If the termination condition is not met, the method 200 goes back to the step of determining 208 the utility score using the updated mapping



functions. If the termination condition is met, the method applies 216 the mapping functions to the image data to generate RGB images from the endmember decompositions applying equation (2) or (4).

- 5 The endmember output colour of an image can be used in pre and post processing of the image data.

Processing of the image includes, but is not limited to:

Image editing

- 10 re-illumination, such as colour manipulation to change an image from warm to cold, or change illumination spectrum of an image to change the illumination from daylight to sunlight, or sunlight to tungsten

re-shading

- 15 re-colouring, for example to change a black and white image to colour based on the properties of a known colour set or colouring the image such that materials for each of the identified hyperplanes are represented by distinctive colours

material modification

highlight removal

surface rendering

- 20 material modification

Material recognition or classification

material, pattern or object recognition and classification

- 25 Hardware calibration

improve photometric calibration, such as that of sensor 112 that captured the image

- 30 The method for determining an output colour value for a component of a decomposed multispectral or hyperspectral image may be directly implemented in hardware, such as an application specific integrated circuit or a field programmable gate array. This hardware may even be integrated with the sensor 112 into a single integrated circuit.

Applications of the methods described here include to the fields of:

- 35 digital photography, such as image editing  
manufacturing, such as quality and production control.

product analysis, such as determining whether a vehicle had been in an accident, and

surveillance, such as face identification and tracking.

- 5 It will be appreciated by persons skilled in the art that numerous variations and/or modifications may be made to the specific embodiments without departing from the scope as defined in the claims.

10 It should be understood that the techniques of the present disclosure might be implemented using a variety of technologies. For example, the methods described herein may be implemented by a series of computer executable instructions residing on a suitable computer readable medium. Suitable computer readable media may include volatile (e.g. RAM) and/or non-volatile (e.g. ROM, disk) memory, carrier waves and transmission media. Exemplary carrier waves may take the form of electrical, 15 electromagnetic or optical signals conveying digital data streams along a local network or a publically accessible network such as the internet.

It should also be understood that, unless specifically stated otherwise as apparent from the following discussion, it is appreciated that throughout the description, discussions 20 utilizing terms such as "estimating" or "processing" or "computing" or "calculating", "optimizing" or "determining" or "displaying" or "maximising" or the like, refer to the action and processes of a computer system, or similar electronic computing device, that processes and transforms data represented as physical (electronic) quantities within the computer system's registers and memories into other data similarly represented as 25 physical quantities within the computer system memories or registers or other such information storage, transmission or display devices.

## References

- [1] E. Angelopoulou, R. Molana, and K. Daniilidis. Multispectral skin color modeling. In *Computer Vision and Pattern Recognition*, pages 635–642, 2001.
- [2] J. Bonnans, J. C. Gilbert, C. Lemarechal, and C. Sagastizabal. Numerical  
5 Optimization-Theoretical and Practical Aspects. Universitext. Springer Verlag, Berlin, 2006.
- [3] R. Datta, D. Joshi, J. Li, and J. Z. Wang. Studying aesthetics in photographic images using a computational approach. In *European Conference on Computer Vision*, pages III:288–301, 2006.
- 10 [4] C. P. Huynh and A. Robles-Kelly. A nurbs-based spectral reflectance descriptor with applications in computer vision and pattern recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2008.
- [5] E. T. Jaynes. Information Theory and Statistical Mechanics. *Phys. Rev.*, 106(4):620–630, 1957.
- 15 [6] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, 1983.
- [7] S. A. Shafer. Using color to separate reflection components. *Color Research and Applications*, 10(4):210–218, 1985.
- [8] W. S. Stiles and J. M. Burch. N.P.L. colour-matching investigation: Final report  
20 1958. *Optica Acta*, 6:1–26, 1959.

## CLAIMS:

1. A computer implemented method for determining an output colour value for a component spectrum of a decomposed multispectral or hyperspectral image, the method comprising:
  - 5 determining or accessing the component spectrum and a first colour value for the component spectrum;  
receiving an input colour value for a pixel location in the image, wherein the component spectrum contributes to the pixel;  
determining the output colour value for the component spectrum based on the
  - 10 first colour value for the component spectrum, the input colour value and a utility score; and  
storing the output colour value for the component spectrum on a datastore.
2. The computer implemented method of claim 1, wherein determining the output
- 15 colour value comprises optimising a cost function.
3. The computer implemented method of claim 1 or 2, wherein the input colour value is received from a user.
- 20 4. The computer implemented method of any one of the preceding claims, wherein the image is decomposed into material spectra and the component spectrum is one of the material spectra.
5. The computer implemented method of claim 4, wherein the method comprises
- 25 determining or accessing a decomposition of the material spectra into endmember spectra; and  
determining an output colour value for each endmember spectrum based on the output colour value for the material spectrum.
- 30 6. The computer implemented method of any one of the claims 1 to 3, wherein the image is decomposed into endmember spectra and the component spectrum is one of the endmember spectra.
7. The computer implemented method of any one of the preceding claims, wherein
- 35 the method is performed, such that an output colour value for each component spectrum represented in the image is determined.

8. The computer implemented method of any one of the preceding claims, wherein the method comprises determining or accessing a component contribution for the pixel location and the cost function is based on the component contribution.
- 5
9. The computer implemented method of any one of the preceding claims, wherein the utility score is any one or more of
- an aesthetic score;
  - a score based on a predetermined scale of danger, benefit, profit or added
  - 10 financial or monetary value;
  - enhanced visualisation of multiple images; or
  - the nature of the rendering context.
10. The computer implemented method of any one of the preceding claims, wherein
- 15 the utility score is based on a classifier.
11. The computer implemented method of claim 10, wherein the classifier is any one or more of
- a quadratic discriminant;
  - 20 a support vector machine;
  - a linear discriminant; or
  - a logistic regressor.
12. The computer implemented method of claim 10 or 11, wherein the classifier is
- 25 based on one or more of:
- abundance of material spectra in the image;
  - most abundant component spectrum;
  - second most abundant component spectrum;
  - rule of thirds;
  - 30 whole image average;
  - macro feature;
  - domain knowledge; or
  - rules and regulations pertaining the utility measure.
- 35 13. The computer implemented method of any one of the preceding claims, wherein the utility score is based on the contrast of the image.

14. The computer implemented method of any one of the preceding claims, wherein the utility score is based on the hazardousness of the component spectrum.
- 5 15. The computer implemented method of any one of the claims 2 to 14, wherein the cost function is based on a utility score for each component spectrum.
16. The computer implemented method of any one of the preceding claims, wherein the method comprises accessing a library of component spectra.
- 10 17. The computer implemented method of any one of the preceding claims, wherein the method comprises determining or accessing illuminant spectra and illuminant contributions.
- 15 18. The computer implemented method of claim 17, wherein the method comprises determining a colour value for each illuminant spectrum.
19. The computer implemented method of claim 17 or 18, wherein the method comprises determining a colour value for each combination of component spectra and illuminant spectra.
- 20 20. The computer implemented method of any one of the claims 8 to 19, wherein the method comprises determining for each pixel location a pixel colour value based on the output colour value for each of the component spectra and the component contributions.
- 25 21. The computer implemented method of claim 20, wherein the method comprises displaying to a user an image based on the pixel colour values.
- 30 22. The computer implemented method of any one of the preceding claims, wherein the method comprises initialising the first colour value for the component spectrum with a default colour.
- 35 23. Software, that when installed on a computer causes the computer to perform the method of any one or more of the claims 1 to 22.

24. A computer system for determining an output colour value for a component spectrum of a decomposed multispectral or hyperspectral image, the system comprising:

5 a processor to determine or access the component spectrum and a first colour value for the component spectrum; to receive an input colour value for a pixel location in the image wherein the component spectrum contributes to the pixel; to determine the output colour value for the component spectrum based on the first colour value for the component spectrum, the input colour value and a utility score; and to store the output colour value for the component spectrum on a datastore.

10


25. The computer system of claim 24, wherein the computer system comprises a port to receive wavelength indexed multispectral image data and a storage to store the image data, the component spectrum and the output colour value.

15

26. A computer implemented method for determining an output colour value for a component spectrum of a decomposed multispectral or hyperspectral image substantially as hereinbefore described with reference to the accompanying drawings.

20

27. A computer system for determining an output colour value for a component spectrum of a decomposed multispectral or hyperspectral image substantially as hereinbefore described with reference to the accompanying drawings.



$$I_c(v) = \left( \sum_{\ell=1}^M \alpha_{\ell} g_{\ell}(\mathbf{L}_{\ell}(\cdot)) \right) \left( g(v) \sum_{m=1}^N \rho_{m,v} f_c(\mathbf{S}_m(\cdot)) + k(v) \right)$$

Fig. 1b



2/4

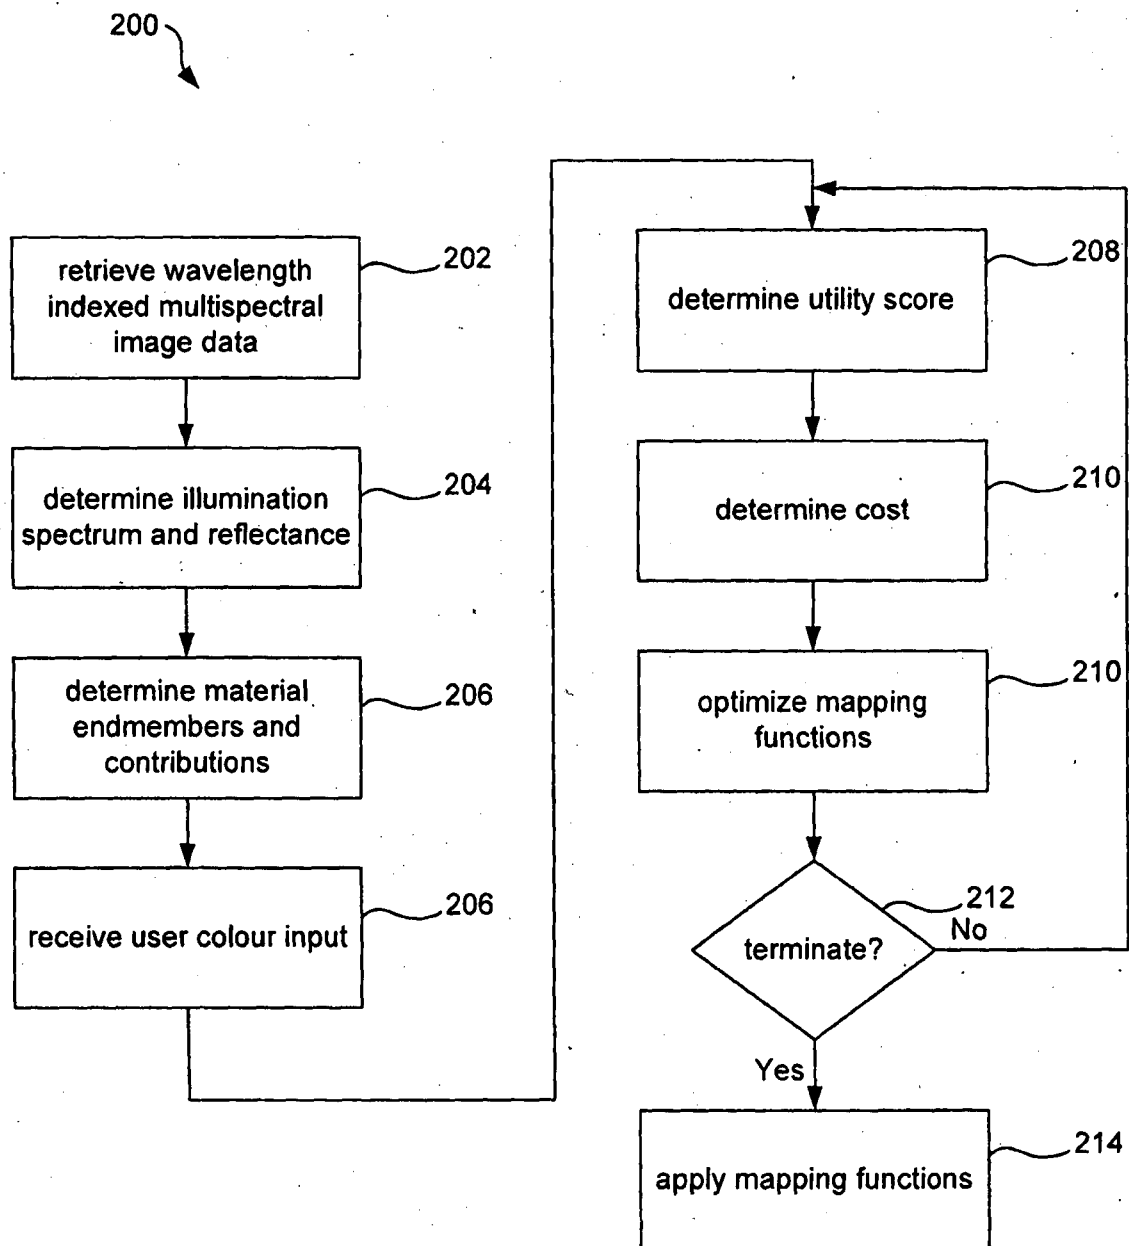


Fig. 2




Fig. 3

4/4




Fig. 4a




Fig. 4b

## INTERNATIONAL SEARCH REPORT

International application No.

PCT/AU2012/001352

## A. CLASSIFICATION OF SUBJECT MATTER

G06K 9/00 (2006.01) G06T 5/00 (2006.01) H04N 1/62 (2006.01)

According to International Patent Classification (IPC) or to both national classification and IPC

## B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

WPI, EPODOC, GOOGLE PATENTS, GOOGLE SCHOLAR, PATENT LENS; Keywords: +spectral, imag+, colour value, component, change, input, user, metamer+, and similar terms.

## C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
	Documents are listed in the continuation of Box C	

☒ Further documents are listed in the continuation of Box C☒ See patent family annex

* Special categories of cited documents:	
"A" document defining the general state of the art which is not considered to be of particular relevance	"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"E" earlier application or patent but published on or after the international filing date	"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"O" document referring to an oral disclosure, use, exhibition or other means	"&" document member of the same patent family
"P" document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search  
29 January 2013Date of mailing of the international search report  
29 January 2013

Name and mailing address of the ISA/AU

AUSTRALIAN PATENT OFFICE  
PO BOX 200, WODEN ACT 2606, AUSTRALIA  
Email address: pct@ipaustalia.gov.au  
Facsimile No.: +61 2 6283 7999

Authorised officer

I G BIGGS  
AUSTRALIAN PATENT OFFICE  
(ISO 9001 Quality Certified Service)  
Telephone No. 0262256118

**Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)**

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos.:  
because they relate to subject matter not required to be searched by this Authority, namely:
2. ☒ Claims Nos.: 26, 27  
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:  
**See Supplemental Box**
3. ☐ Claims Nos.:  
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a)

**Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)**

This International Searching Authority found multiple inventions in this international application, as follows:

1. ☐ As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2. ☐ As all searchable claims could be searched without effort justifying additional fees, this Authority did not invite payment of additional fees.
3. ☐ As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:
4. ☐ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

**Remark on Protest**

- ☐ The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.
- ☐ The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.
- ☐ No protest accompanied the payment of additional search fees.

INTERNATIONAL SEARCH REPORT		International application No.
C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		PCT/AU2012/001352
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2004/0017586 A1 (NICHOGI) 29 January 2004 Abstract; figure 1-2; paragraph [0015] – [0017], [0032]-[0047].	
A	US 2002/0012461 A1 (MACKINNON et al.) 31 January 2002 Abstract; figure 6; paragraph [0073]-[0079]	
A	US 2009/0128649 A1 (OSORIO et al.) 21 May 2009 Abstract, figure 3-5; paragraph [0010]-[0049].	
T	GU L. et al., “Material-Specific User Colour Profiles from Imaging Spectroscopy Data”, 2011 IEEE Conference on Computer Vision ICCV, 6-13 Nov 2011. [Retrieved on 16 March 2012 ]. Retrieved from the Internet <URL: <a href="http://users.cecs.anu.edu.au/~huynh/publications/gu_huynh_robles-kelly_zhou_iccv2011_colour.pdf">http://users.cecs.anu.edu.au/~huynh/publications/gu_huynh_robles-kelly_zhou_iccv2011_colour.pdf</a> > abstract, figure 1, section 2, 3.1-3.3	1-25

**INTERNATIONAL SEARCH REPORT**

International application No.

**PCT/AU2012/001352****Supplemental Box**Continuation of **Box II**

Claims 26 and 27 do not comply with rule 6.2(a) because they rely on references to the description and/or drawings.

**INTERNATIONAL SEARCH REPORT**

Information on patent family members

International application No.

**PCT/AU2012/001352**

This Annex lists known patent family members relating to the patent documents cited in the above-mentioned international search report. The Australian Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

Patent Document/s Cited in Search Report		Patent Family Member/s	
Publication Number	Publication Date	Publication Number	Publication Date
US 2004/0017586 A1	29 Jan 2004	CN 1454011 A	05 Nov 2003
		EP 1357736 A2	29 Oct 2003
		JP 2004005566 A	08 Jan 2004
		US 2004017586 A1	29 Jan 2004
		US 7312899 B2	25 Dec 2007
US 2002/0012461 A1	31 Jan 2002	AU 5999101 A	26 Nov 2001
		AU 2001259991 B2	06 Jul 2006
		EP 1287681 A2	05 Mar 2003
		JP 2003533831 A	11 Nov 2003
		US 2002012461 A1	31 Jan 2002
		US 6961461 B2	01 Nov 2005
		US 2006170994 A1	03 Aug 2006
		WO 0189201 A2	22 Nov 2001
US 2009/0128649 A1	21 May 2009	EP 1854277 A1	14 Nov 2007
		EP 1854277 B1	27 Jan 2010
		JP 2008532401 A	14 Aug 2008
		US 2009128649 A1	21 May 2009
		US 8044969 B2	25 Oct 2011
		WO 2006092559 A1	08 Sep 2006
End of Annex			

Due to data integration issues this family listing may not include 10 digit Australian applications filed since May 2001.

Form PCT/ISA/210 (Family Annex)(July 2009)