

PRIP-TR-84

July 29, 2003¹

Physics-based Segmentation of Colour Images in Spherical Coordinates

Allan Hanbury

Abstract

We present image segmentation and highlight detection algorithms based on the dichromatic reflection model. For image segmentation, we use the model prediction that objects of a certain colour produce lines (the matte lines) radiating away from the origin of the RGB colour space. These lines therefore show up as peaks in a 2-dimensional histogram of the angular coordinates of a spherical polar coordinate representation of the RGB space. An algorithm for automatically locating these peaks is suggested. When the matte line locations are known, one can define cylindrical polar coordinate systems having their z -axes centred on the matte lines. We suggest a Hough-based algorithm for the detection of highlight lines in 2D-histograms of the ρ and z coordinates of the cylindrical polar coordinate system. Examples of the results of applying these algorithms are given.

¹Updated April 22, 2004.

Contents

1	Introduction	1
2	The Dichromatic Reflection Model	2
2.1	Physical model	2
2.1.1	Surface reflection	3
2.1.2	Body reflection	4
2.1.3	Complete model	4
2.2	Application to RGB histogram interpretation	5
3	Segmentation in Spherical Coordinates	6
3.1	Spherical coordinates	7
3.2	Segmentation based on the 2D histogram	9
3.3	Automatic segmentation of the 2D histogram	11
3.3.1	Histogram creation	12
3.3.2	Finding blob centres	12
3.3.3	Discussion	14
4	Detecting Highlights	15
4.1	Cylindrical coordinates	16
4.2	Hue-brightness histogram	17
4.3	Hough transform	18
4.4	Highlight marking	20
5	Further Examples	20
5.1	Guitar image	21
5.2	Sailboat image	21
5.3	Computer image	21
6	Conclusion	25

1 Introduction

Segmentation of colour images has recently become an extremely important topic of research, given the recent surge of interest in content-based image retrieval and automatic extraction of semantics from images and videos. For example, the Blobworld image retrieval system [4] relies on an initial segmentation of an image in order to have regions in which colour and texture characteristics can be calculated. In applications aiming at automatically extracting semantic information from images, the first step is often a segmentation of the image, followed by various post-processing steps aimed at grouping the resulting regions into larger regions representing physical objects and then describing them semantically [7]. The segmentation algorithms applied in the first step often place boundaries at any colour discontinuity in the image, irrespective of whether these are due to object boundaries or lighting artefacts, such as shadows. This “over-segmentation” then makes the grouping and characterisation steps more complex due to the lack of some basic information on which to base decisions. By including physics-based information into the initial segmentation, one should be able to distinguish at this early stage between at least object boundaries and shadow boundaries, and possibly even between boundaries due to different materials [23]. This would improve the initial segmentation by allowing it to better distinguish between actual physical objects in the scene, thereby simplifying the tasks of the subsequent processing steps.

There are three main approaches to segmenting colour images:

Spatial segmentation techniques These are the most commonly used. They take only local colour or intensity information into account, and so they place boundaries at any colour or brightness change in an image, irrespective of whether it is due to an object boundary or a shadow boundary. Examples of such techniques are the watershed algorithm and region growing.

Clustering techniques These use data clustering algorithms in a feature space containing features extracted from the image. These could be, for example, pixel colour coordinates or texture features [29].

Physics-based segmentation techniques These techniques use models based on the physical properties of light and of objects in an image. Such techniques can, for example, attempt to find actual material boundaries and ignore boundaries due to illumination changes in an image. Physics-based techniques are often variations of the spatial and clustering techniques containing additional constraints based on physical knowledge.

An extremely powerful spatial segmentation technique is the watershed algorithm [25] applied to the gradient of an image. Tests of this algorithm applied to the segmentation of images for video indexing have been presented in [6]. Further multimedia applications making use of the watershed algorithm are the compression of sequences for video telephony [12] and interactive segmentation of video sequences [36].

Physics-based methods, especially those based on the *dichromatic reflection model* for dielectric materials [21], have been successfully applied to the segmentation and analysis of colour images based on their RGB histograms, as well as the estimation of various scene properties such as surface roughness, illumination intensity and illuminant colour [26]. In general,

the use of physics-based techniques has been limited to images captured under carefully controlled conditions which contain only objects which have physical characteristics described by the model. It is thus known in advance that the colour histogram will have a form which is convenient to process.

To meet the aim of segmenting actual objects instead of colour discontinuities, it is necessary to make use of physics-based segmentation algorithms [23]. In this technical report, we investigate the applicability of the dichromatic reflection model to the analysis of general public images, for which it is not known that they contain only objects which satisfy the model. The main problem with this approach is that it is certain that situations will be encountered in which the objects in the image disobey the chosen physical model to such an extent that the resulting segmentation is useless. The algorithms suggested here are therefore not intended to be a final solution to all segmentation problems. One can envisage an approach to the segmentation of an image in which many segmentation algorithms, both spatial and physics-based, are applied to an image. The physics-based methods could be based on various physical models: dichromatic reflection model, unichromatic reflection model [19], etc. After one has obtained a number of possible image partitions, a partition merging scheme could be applied to obtain the most probable partition. Maxwell and Shafer [24] have suggested a hypothesis merging scheme applicable to images containing only dielectric and metallic objects. As we intend to deal with unconstrained images, a more flexible approach, which can also take incorrectly segmented images into account, will have to be investigated. A promising approach to this is the matching of hierarchies of partitions [13, 16, 18].

Our aim is to develop a segmentation algorithm based on characteristics of the dichromatic reflection model and to demonstrate that it works well on a number of images, thereby supporting its usefulness in a multiple-segmentation scheme. After a brief overview of the dichromatic reflection model (Section 2), we develop a segmentation algorithm which uses properties predicted by the model to reduce the segmentation problem to finding peaks in a 2D histogram (Section 3). We next investigate the detection of highlights (specular reflections) in images based on their properties predicted by the model, also using 2D histogram analysis (Section 4). Lastly, Section 5 presents the application of the suggested model-based segmentation and highlight detection algorithms to further images. A shorter version of this technical report is available [15].

2 The Dichromatic Reflection Model

In this section, we give an overview of the dichromatic reflection model, which we use extensively in this work. More detail can be found in [22, 30]. We first consider the physical model, and then how it can be used in the interpretation of RGB histograms.

2.1 Physical model

When a ray of light interacts with a material of some kind, it splits into two, as shown in Figure 1. One part is directly reflected by the surface and the other part enters the material. The light that enters the material then interacts with it, and some of it will eventually, through scattering and refraction by the material, exit by the same surface through which it entered.




Figure 1: Diagram showing the interaction of light with a dielectric material.

These two processes are generally referred to as *surface reflection* and *body reflection*. The ratio between the amount of surface and body reflection depends on the material. Metals tend to reflect over 90% of incident light at the surface, and therefore appear shiny but colourless. With non-conductors, such as dielectrics, a larger proportion of the light interacts with the material, causing the body-reflected light to have a specific colour. The surface-reflected light causes highlights on the material, with their position being dependent on factors such as scene geometry (position of the light source and observer) and surface roughness. Examples of dielectric materials include porcelain (ceramic), glass, plastics, and the oxides of various metals. More specifically, Tominaga [32] has done experiments demonstrating that the following materials are well described by the dichromatic reflection model: plastics, paints, ceramics, vinyls, tiles, fruits, leaves and woods; whereas metal, cloth and paper are not. The proportion of light transmitted and reflected at the surface of a dielectric material can be calculated using the Fresnel equations, which are functions of the angle of incidence of the light, and refractive indices and magnetic permeabilities of the two materials. As refractive index depends on the wavelength of the light, the amount of light reflected and transmitted at the surface is also a function of wavelength.

2.1.1 Surface reflection

The *angle of incidence* θ_i of a light ray incident on a material surface is the angle between the normal to the surface and the direction of the light ray (see Figure 1). Similarly, the *angle of reflection* θ_r is the angle between the reflected light ray and the surface normal. A basic law of geometric optics states that $\theta_i = \theta_r$. Hence, for a perfectly smooth surface illuminated by parallel light rays, all the light will be reflected into the same direction (perfect mirror reflection). For rough surfaces, where the direction of the surface normal is a function of position, the reflected light will be spread out to some extent around the angle of perfect mirror reflection.

In theory, the colour of the reflected light is different to that of the incident light, due to the refractive index being dependent on wavelength. In practice, however, the value of the refractive index varies very little in the visible range of the spectrum. One can therefore make the assumption that the reflected light is always the same colour, or the slightly stronger assumption that the colour of the reflected light is identical to that of the incident light.

2.1.2 Body reflection

The part of the light which enters the body moves through the material, interacting with pigment particles by refraction and scattering, as well as by passing through them. This alters the spectral composition of the light. Some of this light will eventually be directed back to the material surface, where some of it will be reflected back into the material and the rest of it will exit the material. This entire process is referred to as *body reflection*. The spectral composition of the light leaving the material depends on factors such as the shape, distribution, orientation and density of the pigments as well as their light scattering and transmittance properties. In order to simplify the model, we make the assumption that the pigments are randomly distributed in the material, and that none of them cross the surface of the material. This assumption implies that the light exits in random directions from the material. In the case where the exiting light is uniformly distributed, its distribution can be described by the reflection law for a Lambertian surface.

2.1.3 Complete model

The complete model includes both surface and body reflection. The assumption of randomly distributed pigments which are completely embedded in the material has already been described. Essentially, this assumption implies that the body reflection and surface reflection have no effect on each other, and hence give rise to two independent spectra. In addition, the assumption is made that there is only one light source, with no ambient light and no reflections between the objects (inter-reflections). The light L reflected from a point on the object is then simply described as an additive mixture of the light reflected at the material surface L_s and the light reflected by the body L_b , or

$$L(\lambda, \theta_i, \theta_e, \theta_g) = L_s(\lambda, \theta_i, \theta_e, \theta_g) + L_b(\lambda, \theta_i, \theta_e, \theta_g) \quad (1)$$

As shown in the equation, the light reflected is a function of its wavelength λ and of the geometry of the reflection, described by the three *photometric angles* illustrated in Figure 2. The *incidence angle* θ_i is between the illumination direction \mathbf{I} and the surface normal \mathbf{N} . The *excitance angle* θ_e is between \mathbf{N} and the viewing direction \mathbf{V} . Finally, the *phase angle* θ_g is between \mathbf{I} and \mathbf{V} .

Due to the assumption that the photometric properties of the reflection are the same over the whole surface of the material, the surface and body reflection terms can be written as the product of two terms, a spectral power distribution which depends only on the wavelength and a geometric scale factor which depends on the photometric angles:

$$L(\lambda, \theta_i, \theta_e, \theta_g) = m_s(\theta_i, \theta_e, \theta_g) c_s(\lambda) + m_b(\theta_i, \theta_e, \theta_g) c_b(\lambda) \quad (2)$$

where $c_s(\lambda)$ and $c_b(\lambda)$ are the spectral power distributions for the surface and body reflections respectively, and $m_s(\theta_i, \theta_e, \theta_g)$ and $m_b(\theta_i, \theta_e, \theta_g)$ are the geometric scale factors. According to the assumptions made, both spectral power distributions are constant over the entire material, and in addition $c_s(\lambda)$ is identical to the spectral power distribution of the illuminant.

The geometric scale factors are real numbers between 0 and 1. It is possible to replace the spectral power distributions by vectors \mathbf{c}_s and \mathbf{c}_b in an infinite dimensional vector space, with the components giving the amount of reflection at each wavelength. The reflected light




Figure 2: Photometric angles.

can then be described as a linear combination of these two vectors. It is then clear that the spectral coordinates of the reflected light all lie in a plane in the infinite dimensional vector space spanned by the vectors \mathbf{c}_s and \mathbf{c}_b , called the *dichromatic plane*. However, the plane is not evenly covered by reflection points. In general, the majority of points on a dielectric material exhibit very little surface reflection. These points, referred to as the *matte points*, therefore have $m_s(\theta_i, \theta_e, \theta_g) \approx 0$ and hence cluster along the vector \mathbf{c}_b . If the regions of strong surface reflection occupy a very small surface area of the material, forming *highlights* as they often do in actual situations, then it can be assumed that $m_b(\theta_i, \theta_e, \theta_g)$ is approximately constant over this area. The points due to a region of strong surface reflection therefore cluster on a line parallel to the vector \mathbf{c}_s and intersect the vector \mathbf{c}_b at the position given by the value of $m_b(\theta_i, \theta_e, \theta_g)$ over the region. Such a configuration is sketched in Figure 3, in which the *matte line* and *highlight line* clusters are shown. The matte and highlight lines together in general form a skewed-P shape, of which the skewed-T or skewed-L shape are special cases [2]. In addition, if there is more than one highlight region on an object, then more than one highlight line could be visible.

2.2 Application to RGB histogram interpretation

In image analysis, one does not in general measure the spectral power for each wavelength, but uses a finite set of numerical samples, obtained by integrating over a number of wavelengths, to describe the light. The use of three such samples, obtained by measuring light through red, green and blue filters is the standard, although the use of more samples is gaining popularity [28]. A sample C_j is obtained as follows:

$$C_j = \int L(\lambda, \theta_i, \theta_e, \theta_g) s_j(\lambda) d\lambda \quad (3)$$

where $s_j(\lambda)$ is the spectral sensitivity used for sample C_j . By this method, the infinite dimensional space of spectra is reduced to a three-dimensional space, with a colour represented by

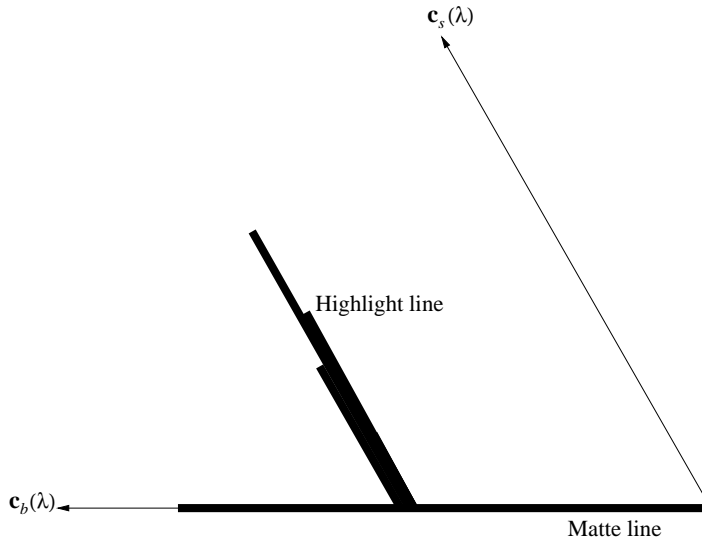


Figure 3: The dichromatic plane defined by body and surface reflection vectors.

the triplet $\mathbf{C} = (R, G, B)$. Spectral integration from the infinite-dimensional vector space of spectral power distributions to a three-dimensional vector space is a linear transformation [22]. For this reason, the relation stated in equation 2 is still valid after spectral integration, and can be written thus:

$$\mathbf{C} = m_s(\theta_i, \theta_e, \theta_g) \mathbf{C}_s + m_b(\theta_i, \theta_e, \theta_g) \mathbf{C}_b \quad (4)$$

The vectors \mathbf{C}_s and \mathbf{C}_b hence span a plane in the RGB histogram, with the vector \mathbf{C}_b passing through the origin (as long as there is no ambient illumination in the scene [26]). Theoretically, all points belonging to an object described by the dichromatic reflection model should lie in a parallelogram with two of the edges given by the vectors \mathbf{C}_s and \mathbf{C}_b . Due to the imprecision of cameras, these pixels do not lie exactly in a plane, but they do tend to cluster around the matte vector and highlight vectors. We intend to capitalise on this tendency for doing physics-based image segmentation.

3 Segmentation in Spherical Coordinates

While physics-based segmentation using the dichromatic reflection model has already been attempted, researchers have tended to restrict themselves to segmenting images which are known to be well-described by the model. Klinker et al. [21], for example, describe a rather complex pixel clustering algorithm which is demonstrated to work well on an image of a group of plastic objects. Healey’s algorithm [20] is demonstrated to work for metallic objects too. Ong and Matsuyama [27] have suggested a clustering algorithm in the RGB space, but their experimental results are limited to images containing only one or two materials. The Markov Random Field approach based on the dichromatic reflection model proposed in [8] is only tested on a single artificial image.

As our goal is the segmentation of arbitrary images captured by the general public, we attempt to use the dichromatic reflection model to as great an extent as possible, while keeping the segmentation algorithms as simple as possible. We present a segmentation algorithm which



Figure 4: Colour test image.

takes advantage of the fact that the object colours in the RGB space tend to form linear clusters radiating out from the origin, in other words, clustered around the body reflection vector. We begin by demonstrating the use of a spherical coordinate system for highlighting the linear clusters, and then describe an algorithm for automatically segmenting the images based on their 2-dimensional histograms in spherical coordinates. The colour image shown in Figure 4 is used to demonstrate the algorithms developed. Examples on other images are presented in Section 5.

3.1 Spherical coordinates

As a demonstration of the clustering of colours into linear structures, we begin with the image in Figure 4, for which no special effort has been made to include only objects satisfying the dichromatic reflection model, and plot its RGB histogram in Figure 5. The points in this histogram have colours given by their corresponding RGB coordinates. It is quite obvious from the histogram that there is a linear cluster corresponding to the blue slide, a linear cluster corresponding to the skin and a (less-linear) cluster corresponding to the grass. Two artefacts due to the acquisition procedure are also visible:

- The blue cluster changes direction and lies along one of the edges of the RGB cube in the high brightness region of the RGB space. This is referred to as *colour clipping* [22] and is due to the limited dynamic range of the camera. Clipped colour pixels are not taken into account by the dichromatic reflection model.
- The clusters are not perfectly linear, but curve slightly. This is due to the gamma-correction of the camera [22]. It could potentially be corrected by one of the gamma-correction compensation algorithms available [9].

As we assume that the linear clusters radiate out from the origin, a representation of the RGB values in spherical coordinates should result in pixels belonging to the same linear cluster having similar angular coordinates. Bajcsy et al. [2] have used a similar approach in their S space. The disadvantage of the S space in the present context is that it requires the scene illumination colour to be known. While algorithms to estimate this exist [10], they introduce

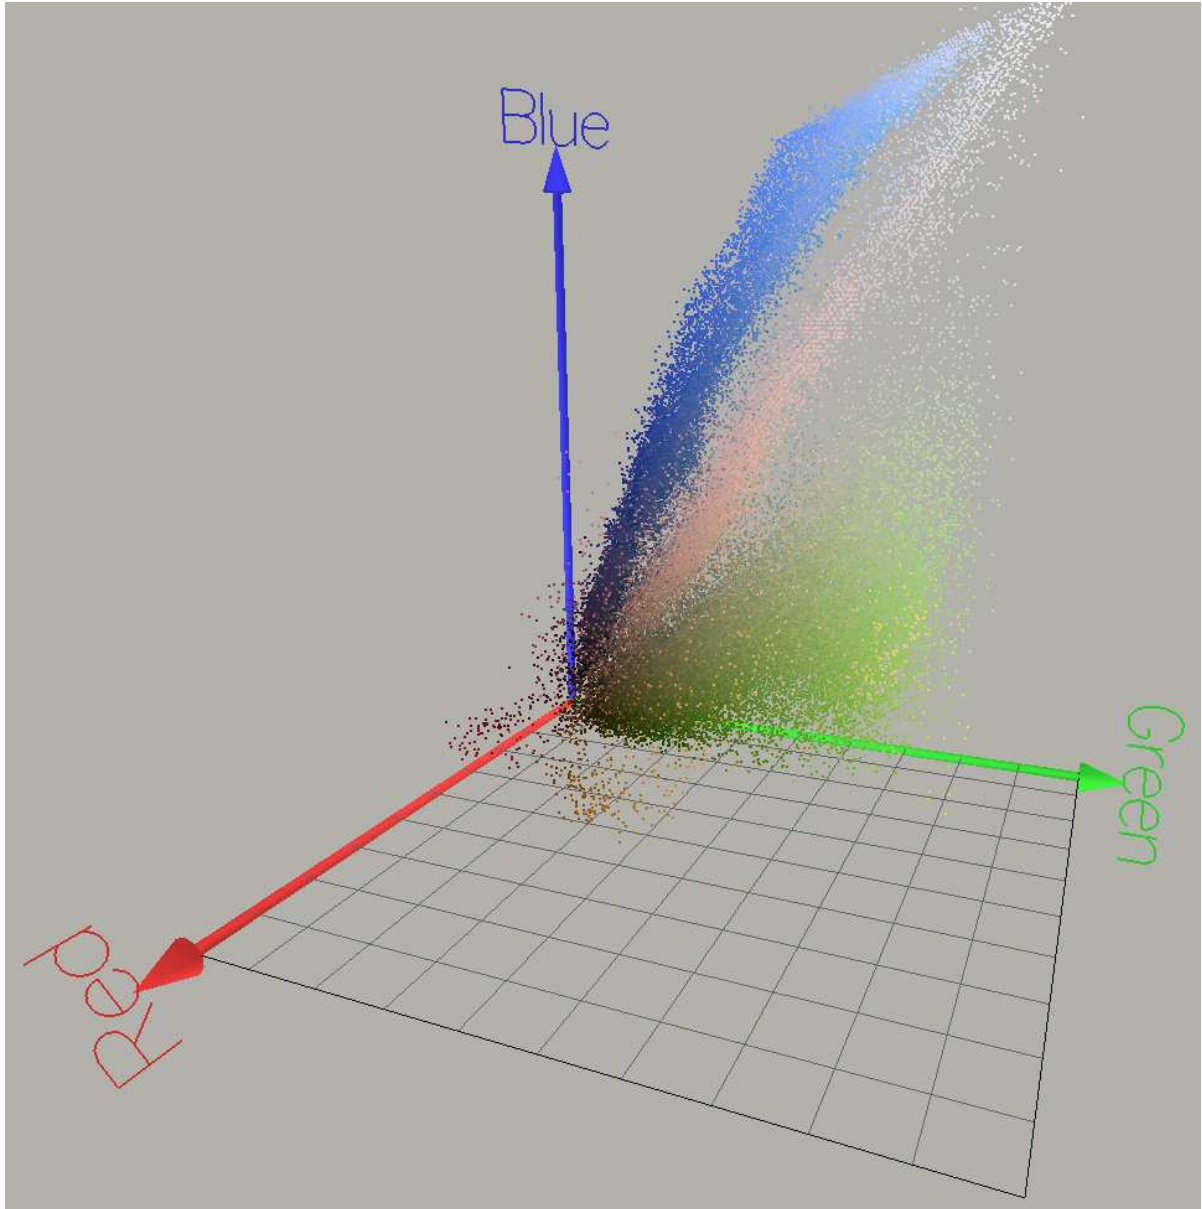


Figure 5: RGB histogram of Figure 4. (Created using the Colorspace software [5]).




Figure 6: The spherical coordinate system in terms of the R , G and B axes.

extra uncertainty into the process. We demonstrate that good results can be achieved directly in the RGB space, although this does not exclude using the \mathbf{S} space with an illumination colour estimation algorithm in the proposed multi-segmentation scheme. The relation of the spherical coordinates M , θ and ϕ to the R , G and B axes is shown in Figure 6, and the conversion is done as follows:

$$M = \sqrt{R^2 + G^2 + B^2} \quad (5)$$

$$\theta = \tan^{-1} \left(\frac{G}{R} \right) \quad (6)$$

$$\phi = \cos^{-1} \left(\frac{B}{M} \right) \quad (7)$$

For completeness, we also present the inverse transformation from spherical to RGB coordinates:

$$R = M \cos \theta \sin \phi \quad (8)$$

$$G = M \sin \theta \sin \phi \quad (9)$$

$$B = M \cos \phi \quad (10)$$

The images corresponding to the spherical coordinate representation of Figure 4 are shown in Figure 7. Notice that for the θ and ϕ images, the coordinates corresponding to a specific object tend to be uniform. This can be better seen by looking at the two-dimensional histogram of the θ and ϕ coordinates, shown in Figure 8. On this histogram, the clusters corresponding to various objects in the initial image have been manually annotated.

3.2 Segmentation based on the 2D histogram

Based on the information visible in the 2-dimensional histogram, it is simple to produce a segmentation of the initial image. One simply places straight lines in the RGB space passing through the origin in the direction of each detected cluster. For each pixel, the line closest to it in the RGB space is found, and it is assigned to the cluster described by that line.




Figure 7: The (a) M , (b) θ and (c) ϕ coordinates of Figure 4 converted to spherical coordinates.

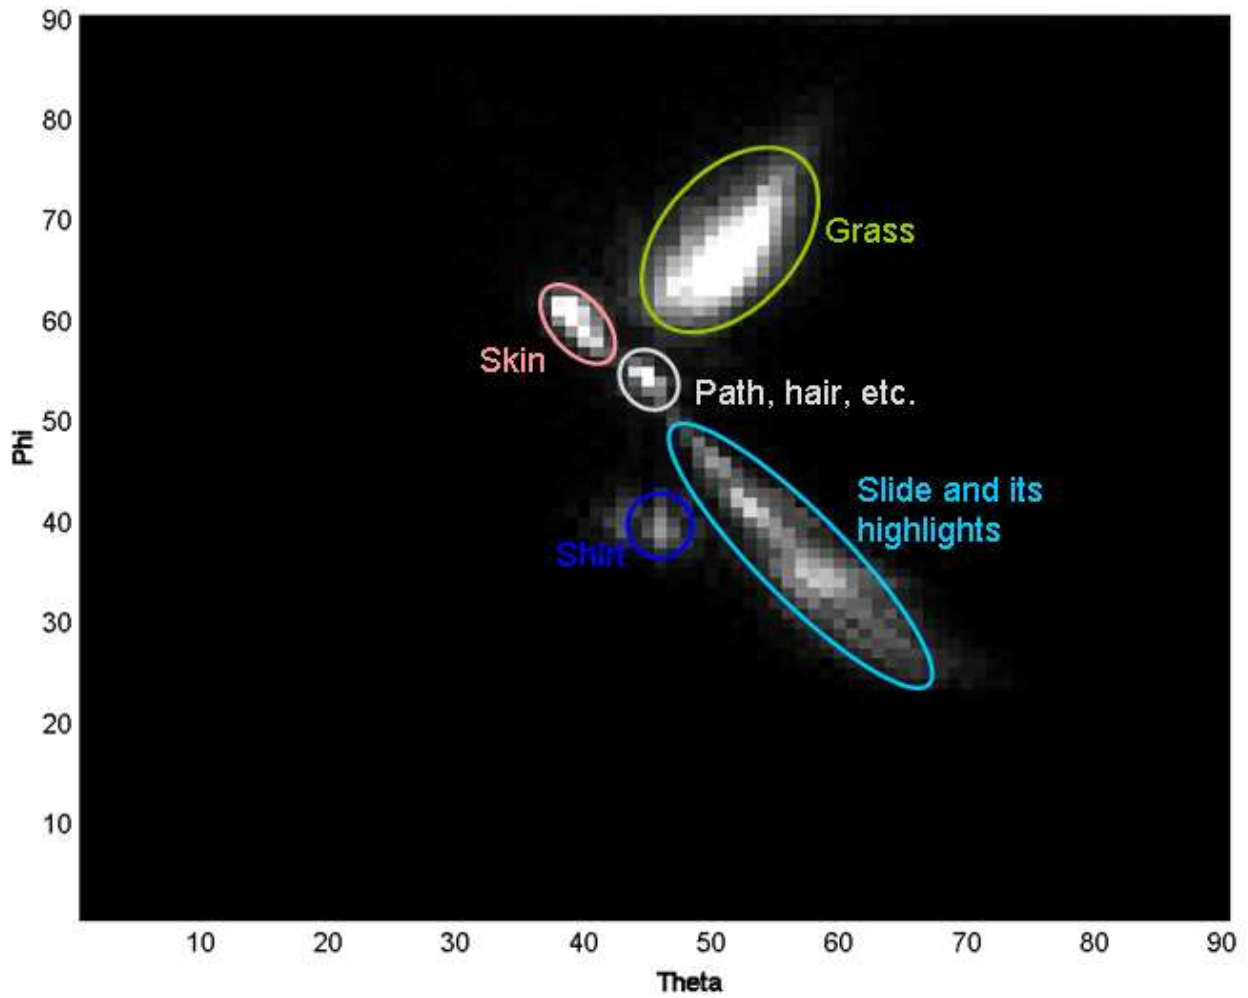


Figure 8: The 2-dimensional θ - ϕ histogram. The regions to which the clusters correspond have been entered by hand.

The most difficult step to automate in this process is the automatic detection of the clusters in the 2D histogram. We therefore ignore this step in this section and present the results of a segmentation based on the manually found clusters indicated in Figure 8. An algorithm for automatically detecting clusters is presented in Section 3.3.

In 3D space, the shortest distance d between a line passing through points $\mathbf{x}_1 = (x_1, y_1, z_1)$ and $\mathbf{x}_2 = (x_2, y_2, z_2)$, and a point $\mathbf{x}_0 = (x_0, y_0, z_0)$ is given by [34]

$$d = \frac{|(\mathbf{x}_2 - \mathbf{x}_1) \times (\mathbf{x}_1 - \mathbf{x}_0)|}{|\mathbf{x}_2 - \mathbf{x}_1|} \quad (11)$$

where \times is the cross product between two vectors.

We begin with a colour image described by the function $f(\mathbf{x}) = [R(\mathbf{x}), G(\mathbf{x}), B(\mathbf{x})]$, which gives the value of R , G and B at each position \mathbf{x} in the image. The algorithm for assigning each pixel of $f(\mathbf{x})$ to a region works as follows:

1. The 2D θ - ϕ histogram is analysed, either manually or automatically, to produce N pairs of coordinates $\mathbf{L}_i = (\theta_i, \phi_i)$, $i = 1, 2, \dots, N$ corresponding to the cluster centres in the histogram.
2. The coordinates \mathbf{L}_i are converted to RGB coordinates $\mathbf{L}_i^C = (R_i, G_i, B_i)$ using Equations 8 to 10. The value of M in this conversion is arbitrary, and is taken to be 1.
3. For every pixel \mathbf{x} in the input colour image:
 - (a) Equation 11 is evaluated for each \mathbf{L}_i^C , taking $\mathbf{x}_1 = (0, 0, 0)$, $\mathbf{x}_2 = \mathbf{L}_i^C$ and $\mathbf{x}_0 = f(\mathbf{x})$, producing N distances d_i .
 - (b) The pixel is assigned to the cluster j corresponding to the smallest distance, i.e. $d_j = \min_i (d_i)$.

As a demonstration of this algorithm, we have manually estimated the centres of the five clusters indicated in Figure 8. Their θ and ϕ coordinates are shown in Table 1. Applying the above algorithm to the image in Figure 4 using the indicated cluster centres produces the segmentation shown in Figure 9b. In this image, each cluster is indicated in a different randomly-chosen colour. In general, the segmentation corresponds well to the categories identified in the histogram. Almost all regions which are misclassified correspond to highlights, such as those on the slide and on the boy's forehead. This is understandable, as the highlights give rise to side branches of the matte clusters, and hence do not satisfy the assumptions made. The identification of highlights based on this initial segmentation is discussed in Section 4.

3.3 Automatic segmentation of the 2D histogram

In the previous section, we manually located the centres of the clusters in the 2D histogram. In this section we propose an algorithm which should find the centres automatically. The algorithm is divided into two steps:

1. Histogram creation.
2. Finding blob centres.

These steps are described in more detail in the following subsections.

Cluster No.	Description	Coordinates
1	Grass	(51, 65)
2	Skin	(39, 60)
3	Path, hair, etc.	(45, 55)
4	Shirt	(46, 39)
5	Slide	(58, 34)

Table 1: The (θ, ϕ) coordinates of the manually chosen cluster centres from the 2D-histogram in Figure 8.

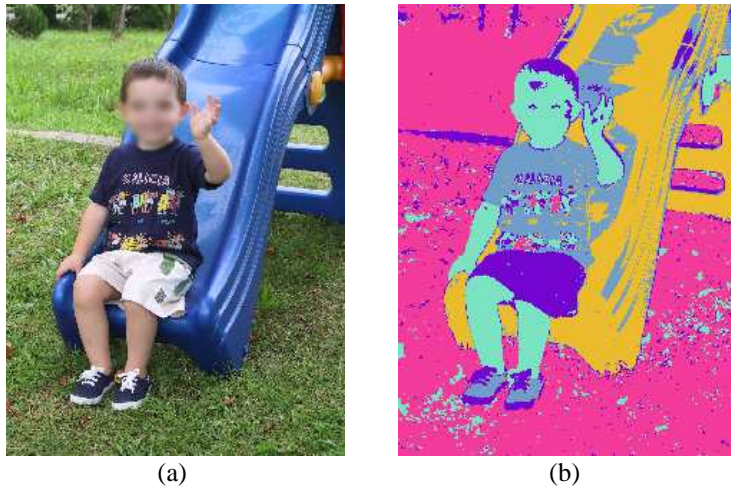


Figure 9: (a) The initial image (shown for comparison purposes) and (b) the segmented image based on the manually chosen cluster centres in the 2D histogram.

3.3.1 Histogram creation

One of the histogram parameters which can have an effect on the automatic extraction of maxima is the quantisation. It was experimentally determined that a histogram quantisation step of 2° worked well with the proposed algorithm. This larger quantisation step has the effect of reducing the number of local maxima in the histogram.

Colour clipping in the histogram can cause some of the blobs in the histogram to become larger, and hence can have an effect on the finding of blob centres. We therefore implemented a heuristic solution to this problem proposed by Klinker et al. [22]. Colour clipped pixels can be identified by placing a high threshold on the values of the three colour channels. Pixels which have a value in one of the R , G or B channels greater than the threshold are excluded from the histogram. A threshold of 240 was chosen for all examples. The histogram of Figure 4 calculated using these two modifications is shown in Figure 10a.

3.3.2 Finding blob centres

The blob centres are initially found by simply using the h -maxima operator [31]. The parameter h that has to be chosen for this operator indicates how high a peak should be with respect to the surrounding region to qualify as a valid maximum. We found experimentally that a value




Figure 10: (a) The θ - ϕ histogram for Figure 4 with bins of size 2° . The values on the axes should be multiplied by 2 to get the actual values in degrees. (b) The h -maxima of histogram (a). (c) The watershed segmentation of the inversion of (a) on which (b) have been imposed as minima.

of 5×10^{-4} multiplied by the number of pixels in the image under consideration works well. The results are not sensitive to small changes in the value of this parameter, as the peaks are generally well-defined. The positions of the five maxima found by applying the h -maxima operator to the histogram of Figure 10a are shown in Figure 10b. Note that the two points at the lower right actually indicate a single maximum, as 8-connectivity was used.

The next step is to refine the positions of the maxima. The aim is to calculate a greylevel-weighted centre of mass of each blob. Before this can be done, we need to find the “zones of influence” of each maximum, which we do by using the watershed operator [31]. This is applied to the inverted 2D-histogram (the flooding starts from the minima) on which the maxima found by the h -maxima operator have been imposed as the only minima from which the flooding starts [31]. Figure 10c shows the watershed lines and regions found by the watershed operator used on the example histogram. The greylevel-weighted centre of mass is now calculated within each region in the standard way. In other words, if there are N regions labelled by $i = 1, 2, \dots, N$,



(a)

θ	ϕ
44.9	40.3
44.9	53.3
39.1	61.1
51.7	68.4
59.7	34.4

(b)

Figure 11: (a) Automatically segmented image. (b) Coordinates of the blobs found in the histogram by the suggested algorithm.

the x and y coordinates of the centre of mass of region i , denoted by x_{CM}^i and y_{CM}^i are calculated as

$$x_{CM}^i = \frac{\sum_{\mathbf{x} \in i} x f(\mathbf{x})}{\sum_{\mathbf{x} \in i} f(\mathbf{x})} \quad \text{and} \quad y_{CM}^i = \frac{\sum_{\mathbf{x} \in i} y f(\mathbf{x})}{\sum_{\mathbf{x} \in i} f(\mathbf{x})}, \quad (12)$$

where the sums are over all the pixels which are part of region i . The coordinates of the centres of mass of the five regions located automatically by the algorithm are shown in Figure 11b. The segmentation obtained when using these values in the process described in Section 3.2 is shown in Figure 11a. The regions obtained agree well with those in the segmentation based on manually determined blob centres shown in Figure 9b. Results of applying this algorithm to further images are given in Section 5.

3.3.3 Discussion

This algorithm has only two parameters to be set:

1. The first is common to many histogram analysis algorithms, namely how high a peak should be to differentiate it from a local maximum. We have adopted the solution of setting this value to a fraction of the total number of pixels in an image. For the images to which this algorithm has already been applied, the real peaks have been significantly more pronounced than the local maxima, so a rough setting of this parameter is sufficient to distinguish between them.
2. The second parameter is the threshold on the R , G and B channel to remove the potentially colour clipped pixels. It is used to make the blobs slighter better defined by removing any low-value “tails” from them. The threshold step could also be omitted in the majority of cases without affecting the blob centre values significantly.

A simple measure of the prominence of the peak in region i can be obtained by summing the greyscale values in region i . This could also be turned into a relative prominence by dividing each

prominence value by the maximum prominence value over all regions. The usefulness and use of this prominence value remains to be investigated.

Many other algorithms exist which could potentially be used to locate the blob centres in the 2D-histogram, such as graph-theoretical clustering [11].

4 Detecting Highlights

So far we have only used one aspect of the dichromatic reflection model in the image segmentation step, namely the expected form of the matte line in the histogram. We now use the knowledge of the position of this line extracted during the segmentation step to attempt to locate any highlight lines radiating out from it. The crucial step in our approach is the conversion to a cylindrical polar coordinate system with its z -axis centred on the matte line to be analysed. Image analysis methods developed for use in cylindrical polar coordinate colour spaces (HSV, HLS, etc.) are also applicable in this representation [14].

Recently there has been some work done on the automated detection of highlights in images through the use of 2D histograms of the brightness and saturation coordinates [1, 33] in a cylindrical polar coordinate colour representation. In these histograms, the colours of pixels in highlight regions tend to form straight lines in the high brightness part of the histogram. In effect, the highlight lines in the RGB space are all projected onto a plane in the histogram. What is not mentioned is that this method is particularly good for detecting *colour clipped* pixels. These are the regions of a scene whose brightness exceeds the dynamic range of the camera. The pixels corresponding to these regions tend to lie along the faces of the RGB cube [21]. This implies that in the brightness-saturation histogram, they will lie along the line corresponding to the maximum saturation for a given brightness, the position of which can easily be determined. Torres et al. [33] take the rather brutal approach of performing a histogram equalisation of the brightness channel before calculating the brightness-saturation histogram. While the results demonstrated look good, this step should guarantee the detection of some highlights in every image, even if none are present.

The algorithm for finding these highlight lines suggested in this section consists of the following steps, which are described in more detail in the indicated sub-sections:

1. Conversion to cylindrical polar coordinates (Section 4.1).
2. Hue-brightness histogram analysis (Section 4.2), from which the expected angular position of the highlight line with respect to the corresponding matte line is determined.
3. Hough transform of the luminance and saturation (Section 4.3), from which the exact coordinates of the highlight line are determined.
4. Marking the highlights in the image (Section 4.4).

The general idea is to place the z -axis of the cylindrical coordinate system along a matte line (the first step of the above algorithm). Any lines radiating away from this matte line are then more easily detectable in this coordinate system. Steps 2–4 of the algorithm are only one way of making use of this simplified access to highlight lines. As an example, we extract the highlights on the slide in the example image (Figure 4).




Figure 12: The cylindrical polar coordinate system in terms of the R , G and B axes.

4.1 Cylindrical coordinates

We wish to represent the image pixel colour coordinates in a cylindrical polar coordinate system having its z -axis lined up with a matte line having coordinates (θ, ϕ) . This has the effect of making highlights associated with this matte line more visible and hence easier to detect. The relation of the cylindrical polar coordinates ρ , θ_c and z to the R , G and B axes is shown in Figure 12, and the conversion is done as follows:

$$\rho = \sqrt{R^2 + G^2} \quad (13)$$

$$\theta_c = \tan^{-1} \left(\frac{G}{R} \right) \quad (14)$$

$$z = B \quad (15)$$

Traditionally, both the cylindrical and spherical coordinate systems have an angular coordinate θ . We use θ_c for the cylindrical angular coordinate to distinguish between the two.

In order to align the z -axis of the cylindrical coordinate system with the matte line having a direction given by the spherical coordinates θ and ϕ , we rotate the RGB axes so that the B axis is aligned with the matte line. This is done in two steps [35]. The R and G axes are initially rotated anti-clockwise around the B axis through an angle θ , described by the following rotation matrix:

$$R_B = \begin{bmatrix} \cos(-\theta) & -\sin(-\theta) & 0 \\ \sin(-\theta) & \cos(-\theta) & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad (16)$$

The axes resulting from this rotation are labelled R' , G' and B . The R' and B axes are then rotated anti-clockwise around the G' axis through an angle ϕ , described by the following rotation matrix:

$$R_{G'} = \begin{bmatrix} \cos(\phi) & 0 & \sin(\phi) \\ 0 & 1 & 0 \\ -\sin(\phi) & 0 & \cos(\phi) \end{bmatrix} \quad (17)$$

Given colour coordinates $\mathbf{c} = (R, G, B)^T$, the coordinates \mathbf{c}'' after the above transformation are given by $\mathbf{c}'' = R_{G'} R_B \mathbf{c}$. These coordinates are then simply substituted into Equations 13 to




Figure 13: The (a) ρ , (b) θ_c and (c) z coordinates of the cylindrical coordinate representation with the z -axis aligned with the matte line corresponding to the slide.

15 to obtain a system of cylindrical polar coordinates around the rotated B -axis. Note that if one substitutes spherical coordinates $\theta = 45^\circ$ and $\phi = 54.7356^\circ$ into the above equations, the B -axis will line up with the achromatic axis (i.e. the axis passing through all the grey colours), and the resulting coordinate system will be the 3D-polar coordinate colour system described in [17]. To save computation time, the two rotation matrices can be multiplied together, giving the following matrix:

$$R_{G'} R_B = \begin{bmatrix} \cos(\phi) \cos(-\theta) & -\cos(\phi) \sin(-\theta) & \sin(\phi) \\ \sin(-\theta) & \cos(-\theta) & 0 \\ -\sin(\phi) \cos(-\theta) & \sin(\phi) \sin(-\theta) & \cos(\phi) \end{bmatrix} \quad (18)$$

In the example image, the spherical coordinates of the blob corresponding to the slide are $\theta = 59.7^\circ$ and $\phi = 34.4^\circ$. The images showing the ρ , θ_c and z coordinates when the z -axis is aligned with the matte line corresponding to these coordinates are shown in Figure 13. In this representation, the pixels having a colour closest to the selected matte line have the smallest ρ values, as can be seen for the slide region in this example. Highlight lines extending away from the matte line will be characterised, in this coordinate representation, by increasing ρ values and relatively constant θ_c values. As the highlights are bright, they should also have high z values. It can be seen that the reflections on the slide have these characteristics.

Two possible approaches are conceivable: an analysis of the images to find, for example, areas adjacent to the slide which have the required characteristics to be highlights, as done in [2], or the analysis of histograms to find potential highlight pixels. We pursue the latter approach.

4.2 Hue-brightness histogram

In this step we find the direction in which a potential highlight line extends from a matte line through the use of a 1D z -weighted θ_c histogram. This histogram is created by first quantising




Figure 14: The z -weighted θ_c histogram for the cylindrical coordinate system centred on the slide.

the θ_c image into 360 integer levels ($0^\circ, 1^\circ, \dots, 359^\circ$). The total count W_ξ in bin ξ of the histogram is

$$W_\xi = \sum_x z_x \delta_{\xi \theta_{cx}} \quad (19)$$

where the sum is over all pixels in the image, θ_{cx} and z_x are the θ_c and z image values at position x , and δ_{ij} is the Kronecker delta. As the θ_c values of points on a highlight line are relatively constant, the peaks in a histogram of the θ_c image show the positions of possible highlight lines. The z -weighting is used as it is expected that the highlights are bright and therefore have high z values. The highest peak in the histogram, at position θ_{\max} , is the expected position of the highlight line.

The z -weighted θ_c histogram calculated using the cylindrical coordinate images centred on the slide in the example image is shown in figure 14. The highest peak in this histogram is at 327° . At present, we only take the position of the highest peak into account in further steps. The investigation of the other peaks remains to be done.

4.3 Hough transform

We now make use of a Hough-transform inspired method to locate the highlight line precisely. Figure 15a shows the 2D z - ρ histogram for the cylindrical coordinate system centred on the slide in the example image (the z values are on the horizontal axis). This histogram includes only those pixels which have θ_c values in the range of 20° on either side of the maximum angular value found from the 1D hue-brightness histogram, i.e. $327^\circ \pm 20^\circ$. Due to the large number of colours in the image, removing this limit results in a histogram containing too much extraneous information which could mask the highlight line. The 20° size was chosen by experiment. It should be large enough to include all the pixels in the highlight line, but to exclude as many




Figure 15: (a) The z - ρ 2D histogram for pixels in the angular range $\theta_c = 327 \pm 20^\circ$ for the images in Figure 13. (b) The Hough-type transform of histogram (a) described in the text.

extraneous pixels as possible. On this histogram, the highlight line is clearly visible on the right, originating on the horizontal axis. The cluster further up is another matte line.

Based on the underlying reflection model theory, we expect the highlight lines to:

1. lie in the high brightness area of the histogram.
2. start close to the matte line (which in this histogram is the horizontal axis).
3. slope towards the right in the histogram.

These lines are therefore most usefully characterised by the following two parameters: z_l , the starting position of the line on the horizontal axis and θ_l , the angle of the line with respect to the horizontal axis (measured between the positive direction of the horizontal axis and the line).

The Hough-based algorithm proceeds as follows:

1. Quantise the z and ρ images to be integers in the range 0 to 100.
2. Initialise a matrix H to have 91 rows (representing the θ_l values) and 101 columns (representing the z_l values).
3. For every pixel $(\rho_x, \theta_{cx}, z_x)$ having angular value θ_{cx} in the range $\theta_{\max} \pm \theta_w$ and $\rho_x \leq \rho_{\max}$:
 - (a) Let z_t go from z_{\min} to z_{\max} in steps of one histogram bin, where z_{\min} is a user parameter and $z_{\max} = z_x$:
 - i. Calculate the angle $\theta_t = \arctan\left(\frac{\rho_x}{z_x - z_t}\right)$ rounded to the nearest integer.
 - ii. Increment the value at row θ_t and column z_t of H by 1.
4. Find the maximum value in H , the row and column of which give the coordinates of the highlight line.

The only parameters to be set by the user are z_{\min} , θ_w and ρ_{\max} . All these parameters deal with the minimisation of the effect of pixels not belonging to the cluster of interest. The parameter θ_w gives the range around the selected angle θ_{\max} as for the construction of the 2D-histogram. The parameter z_{\min} implements the assumption that the highlight line is in the high brightness part of the histogram. Upon examining Figure 15a, one can see the large cluster to the left due to another matte and highlight cluster in the RGB space, which gives rise to a very strong peak in the Hough transform if the minimum brightness is not limited. Finally, the parameter ρ_{\max} implements the assumption that the highlight lines start close to the matte line and therefore that a significant part of them will lie close to this line. An additional heuristic which was found to be useful in differentiating between true and false peaks in the Hough transform is to ignore peaks which lie exactly on the brightness cut-off line. We take $z_{\min} = 50$, $\theta_w = 20$ and $\rho_{\max} = 25$ in all the experiments reported in this document. The user is of course free to choose any quantisation step in step 1 of the algorithm, although this should have little influence on the algorithm if sensible values are chosen.

The Hough matrix H corresponding to the 2D-histogram of Figure 15a is shown in Figure 15b. The maximum value in this histogram corresponds to line coordinates $z_l = 77$ and $\theta_l = 69^\circ$, which corresponds well to the highlight line visible on the corresponding 2D-histogram.

Another way of limiting the extraneous clutter due to the other matte lines would be to limit the construction of the histogram to pixels in a spatial region including and surrounding the objects corresponding to the matte line under consideration. This remains to be investigated.

4.4 Highlight marking

Finally the pixels in the image corresponding to the highlight line detected in the previous steps are marked. We do this by finding all pixels whose colours lie in a cylinder surrounding the highlight line detected in the previous step. The central axis of this cylinder lies on the highlight line, it has a radius of R_c , starts at a distance of ρ_{\min} from the matte line and ends at a distance of ρ_{\max} from the matte line. In all our experiments, we have used experimentally determined values of these parameters: $R_c = 0.15$, $\rho_{\min} = 0.05$ and $\rho_{\max} = 0.5$. The highlight regions on the slide determined for the example image are shown in white on the segmented image in Figure 16b. To aid comparison, the test image is reshowed in Figure 16a. It is clear that the highlights on the slide are extremely well detected, with those on the steps behind the slide having been detected too.

This part of the algorithm has much scope for improvement. At the moment, the highlight pixels are marked without taking the other matte lines into account for example. The marking should also take into account that the highlight line is usually thicker closer to the matte line than further away.

5 Further Examples

We present some further examples of the algorithms described applied to general public images. In each figure, we show (a) the initial image, (b) the segmented image and (c) the 2D θ - ϕ histogram. The locations of the histogram maxima found by the automated algorithm described

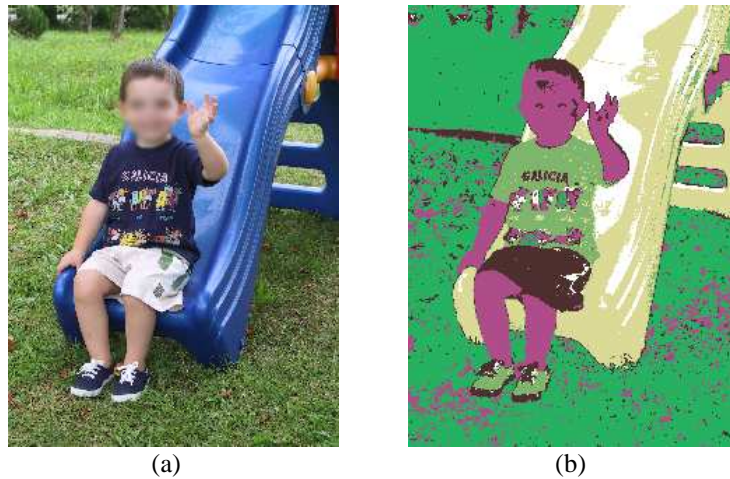


Figure 16: (a) The test image (identical to Figure 4). (b) The test image segmented into regions with the highlights marked.

in section 3.3 are also given. For some of the examples, the results of highlight detection are also shown.

5.1 Guitar image

The results are shown in Figure 17. The image has been divided into five regions by the automatic segmentation process, with the guitar body very well segmented. There is some confusion between the hand and the inside of the guitar, and the black regions of the pullover and the background. For highlight detection, we used the coordinates of the hand colour ($\theta = 37.0, \phi = 64.1$), producing the results shown in Figure 17e. The highlights on the hand are extremely well detected, although there are some false detections inside the guitar.

5.2 Sailboat image

The results are shown in Figure 18. Two groups are visible in the histogram. The upper left one corresponds to the greenery, and the lower right one to the water, sky and clouds. The automatic cluster detection algorithm finds two maxima in the lower blob, thus separating the clouds from the sky. In addition, the maximum in the upper left corner is detected, which leads to the (possibly unwanted) detection of the black parts of the trees.

5.3 Computer image


The results are shown in Figure 19. The image has been divided into regions corresponding to the skin, hair, white regions, back wall and side wall/curtains.

We demonstrate a more detailed highlight analysis on the skin regions of this image to demonstrate a potential problem which arises. The cylindrical coordinate system is constructed around the line corresponding to the skin, in spherical coordinates $\theta = 33.6^\circ$ and $\phi = 67.5^\circ$. The maximum on the brightness-weighted hue histogram is at 146° , and the 2D $\rho - z$ histogram



(a)

(b)



(c)



θ	ϕ
59.0	30.6
37.0	64.1
31.5	89.3
41.8	78.5
49.5	53.5

(d)



(e)

Figure 17: Guitar example. (a) Initial image. (b) Segmented image. (c) θ - ϕ histogram. (d) Coordinates of the blobs found in the histogram. (e) Image with highlights on the hand marked.



θ	ϕ
28.6	70.1
2.1	89.9
48.6	55.5
55.3	51.4

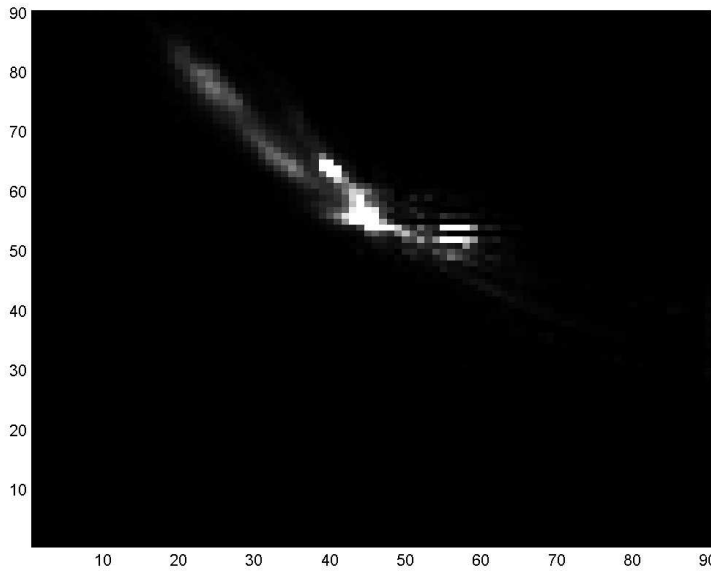
(d)

Figure 18: Sailboat example. (a) Initial image. (b) Segmented image. (c) θ - ϕ histogram. (d) Coordinates of the blobs found in the histogram. The image is from the USC-SIPI image database (<http://sipi.usc.edu/services/database/>).



(a)

(b)



(c)

θ	ϕ
45.0	56.4
33.6	67.5
24.5	78.8
58.2	51.7
40.4	64.2

(d)

Figure 19: Computer example. (a) Initial image. (b) Segmented image. (c) θ - ϕ histogram. (d) Coordinates of the blobs found in the histogram. The image is from NASA.

corresponding to this direction is shown in Figure 20a. The Hough-transform version of the histogram is shown in Figure 20b. In the 2D histogram, the actual highlight line is visible faintly towards the bottom right of the histogram, although it is overshadowed by the strong matte line also visible in the histogram (the greyvalue scale has been adjusted so that the highlight line is visible). In the Hough-transform version, some curves corresponding to the correct line are visible, but they are also overshadowed by the curves generated by the matte line. The maximum of the Hough-transform version, at coordinates $\theta_t = 86^\circ$ and $z_t = 61$ does not therefore indicate the actual position of the highlight line. Interestingly, the highlights on the skin regions, shown in white in Figure 20c, are still rather well detected. For comparison, the highlight regions found when the correct (manually-determined) coordinates of the highlight line are used, $\theta_t = 60^\circ$ and $z_t = 75$, are shown in Figure 20d.




Figure 20: Computer image highlight analysis for the skin region. (a) 2D $\rho - z$ histogram for the direction $\theta_c = 146^\circ$. (b) Hough-transform version of histogram (a). (c) Highlight regions (in white) corresponding to the automatically determined highlight coordinates. (d) Highlight regions corresponding to the manually determined highlight coordinates.

6 Conclusion

The use of the dichromatic reflection model in the segmentation of arbitrary images has been discussed. We initially make use of the prediction by the model that all matte lines radiate out from the origin in the RGB space. This suggests that a representation of RGB coordinates in terms of spherical polar coordinates is useful for segmentation, which is demonstrated through image segmentation based on finding clusters in 2D histograms of the angular coordinates θ and ϕ in the spherical coordinate system. At present, we optimise the coordinates of the matte lines found by finding the centres of the clusters in the 2D histogram as accurately as possible. An improvement in accuracy is likely if the matte lines are fitted to the clusters in the RGB space once the approximate centres have been found on the histogram. A point to be kept in mind is that because all the matte lines intersect at the origin, classification of dark-coloured pixels

into regions is most likely arbitrary. The M coordinate, however, could be used as an indicator of the certainty of the classification. Furthermore, if a detected cluster does not contain any pixels near the origin and does not lie along a line radiating away from the origin, this could be evidence of a non-dielectric object in the image.

Once the matte lines have been found, we suggest a further analysis in a series of cylindrical polar coordinate systems having their z -axes centred on the matte lines. This coordinate system transform has the effect of making the highlights on the objects corresponding to the selected matte line more visible and hence easier to detect. The method for the automatic detection of highlight lines in the cylindrical polar coordinate system is still in need of improvement. As has been suggested, the background clutter in the $\rho - z$ histograms could be reduced by excluding all pixels which are close to matte lines not under consideration. This could be done by limiting the analysis to pixels labeled as corresponding to the selected matte line and their surrounding spatial region. The surrounding region should be taken into account as highlights are often mislabeled in the segmentation step (this naturally poses the problem of how much of the surrounding region should be used). Further information that could be taken into account, for example, is that all the highlight lines are parallel and in the direction of the illumination colour [22]. Methods for estimating the illumination colour [3, 10, 30] could potentially be used to limit the search directions for the highlight line, thereby eliminating the need to use the z -weighted θ_c histogram. A comparison to clustering methods, such as the one described in [27], as well as the potential complementarity of these approaches remains to be investigated.

Having demonstrated that this simple physical model can be useful in the analysis of arbitrary images, it remains to test it out more thoroughly. Methods of comparing and combining segmentations obtained based on different assumptions or different physical models will also be developed.

References

- [1] Jesús Angulo and Jean Serra. Colour feature extraction from luminance/saturation histogram in L_1 representation. Technical Report N-01/03/MM, Centre de Morphologie Mathématique, Ecole des Mines de Paris, March 2003.
- [2] Ruzena Bajcsy, Sang Wook Lee, and Aleš Leonardis. Detection of diffuse and specular reflections and inter-reflections by color image segmentation. *International Journal of Computer Vision*, 17(3):241–272, 1996.
- [3] Vlad C. Cardei, Brian Funt, and Kobus Barnard. White point estimation for uncalibrated images. In *Proceedings of the Seventh Color Imaging Conference (CIC'99)*. Imaging Science and Technology Society, November 1999.
- [4] Chad Carson, Serge Belongie, Hayit Greenspan, and Jitendra Malik. Blobworld: Image segmentation using expectation-maximization and its application to image querying. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(8):1026–1038, August 2002.
- [5] Philippe Colantoni. Colorspace software. <http://www.couleur.org>, 2003.

- [6] Claire-Hélène Demarty. *Segmentation et Structuration d'un Document Vidéo pour la Caractérisation et l'Indexation de son Contenu Sémantique*. PhD thesis, CMM, Ecole des Mines de Paris, 2000.
- [7] Jianping Fan, Xingquan Zhu, and Lide Wu. Automatic model-based semantic object extraction algorithm. *IEEE Transactions on Circuits and Systems for Video Technology*, 11(10):1073–1084, October 2001.
- [8] Paul Fieguth and Slawo Wesolkowski. Highlight and shading invariant color image segmentation using simulated annealing. In *Proceedings of the 3rd International Workshop on Energy Minimization Methods in Computer Vision and Pattern Recognition*, pages 314–327. Springer-Verlag, 2001.
- [9] Graham Finlayson and Ruixia Xu. Illuminant and gamma comprehensive normalisation in log *rgb* space. *Pattern Recognition Letters*, 24:1679–1690, 2003.
- [10] Graham D. Finlayson and Gerald Schaefer. Solving for colour constancy using a constrained dichromatic reflection model. *International Journal of Computer Vision*, 42(3):127–144, 2001.
- [11] K. Fukunaga. *Introduction to Statistical Pattern Recognition*. Academic Press, 1990.
- [12] Cristina Gomila. *Mise en Correspondance de Partitions en vue du Suivi d'Objets*. PhD thesis, CMM, Ecole des Mines de Paris, 2001.
- [13] Cristina Gomila and Fernand Meyer. Tracking objects by graph matching of image partition sequences. In *Proceedings of the 3rd IAPR-TC15 Workshop on Graph-Based Representations in Pattern Recognition*, pages 1–11, 2001.
- [14] Allan Hanbury. A 3D-polar coordinate colour representation well adapted to image analysis. In *Proceedings of the Scandinavian Conference on Image Analysis (SCIA)*, pages 804–811, 2003.
- [15] Allan Hanbury. Highlight detection by 2D-histogram analysis. In *Proceedings of the CGIV 2004 conference*, pages 167–172, 2004.
- [16] Allan Hanbury, Jocelyn Marchadier, and Walter G. Kropatsch. The redundancy pyramid and its application to image segmentation. In *Proceedings of the AAPR Conference 2004*, 2004.
- [17] Allan Hanbury and Jean Serra. Colour image analysis in 3D-polar coordinates. In *Proceedings of the DAGM'03 conference*, pages 124–131, 2003.
- [18] Yll Haxhimusa and Walter G. Kropatsch. Hierarchy of partitions with dual graph contraction. In *Proceedings of the DAGM conference*, 2003.
- [19] Glenn Healey. Using color for geometry-insensitive segmentation. *Journal of the Optical Society of America A*, 6(6):920–937, 1989.

- [20] Glenn Healey. Segmenting images using normalized color. *IEEE Transactions on Systems, Man, and Cybernetics*, 22(1):64–73, 1992.
- [21] Gudrun J. Klinker, Steven A. Shafer, and Takeo Kanade. Image segmentation and reflection analysis through color. In *Applications of Artificial Intelligence VI*, volume 937, pages 229–244. SPIE, 1988.
- [22] Gudrun J. Klinker, Steven A. Shafer, and Takeo Kanade. The measurement of highlights in color images. *International Journal of Computer Vision*, 2:7–32, 1988.
- [23] Bruce A. Maxwell and Steven A. Shafer. Physics-based segmentation: looking beyond color. In *Proceedings of the Image Understanding Workshop*, 1996.
- [24] Bruce A. Maxwell and Steven A. Shafer. Physics-based segmentation of complex objects using multiple hypotheses of image formation. *Computer Vision and Image Understanding*, 65(2):269–295, 1997.
- [25] Fernand Meyer and Serge Beucher. Morphological segmentation. *Journal of Visual Communication and Image Representation*, 1(1):21–46, 1990.
- [26] Carol L. Novak and Steven A. Shafer. Estimating scene properties from color histograms. Technical Report CMU-CS-92-212, School of Computer Science, Carnegie Mellon University, November 1992.
- [27] Chun-Kiat Ong and Takashi Matsuyama. Robust color segmentation using the dichromatic reflection model. In *Proceedings of the 14th International Conference on Pattern Recognition*, volume 1, pages 780–784, 1998.
- [28] Jussi Parkkinen. Spectral color imaging. In *Proceedings of the Scandinavian Conference on Image Analysis (SCIA)*, pages 800–803, 2003.
- [29] E. J. Pauwels and G. Frederix. Finding salient regions in images: Nonparametric clustering for image segmentation and grouping. *Computer Vision and Image Understanding*, 75(1/2):73–85, 1999.
- [30] Valéry Risson. *Application de la morphologie mathématique à l’analyse des conditions d’éclairage des images couleur*. PhD thesis, CMM, Ecole des Mines de Paris, 2001.
- [31] Pierre Soille. *Morphological Image Analysis: Principles and Applications*. Springer-Verlag, 1999.
- [32] Shoji Tominaga. Surface identification using the dichromatic reflection model. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(7):658–670, July 1991.
- [33] Fernando Torres, Jesús Angulo, and Francisco Ortiz. Automatic detection of specular reflectance in colour images using the MS diagram. In *Proceedings of the CAIP*, pages 132–139, 2003.
- [34] Eric W. Weisstein. Point-line distance–3-dimensional, Eric Weisstein’s World of Mathematics. <http://mathworld.wolfram.com/Point-LineDistance3-Dimensional.html>.

- [35] Eric W. Weisstein. Rotation matrix, Eric Weisstein's World of Mathematics. <http://mathworld.wolfram.com/RotationMatrix.html>.
- [36] Francisca Zanoguera. *Segmentation interactive d'images fixes et de séquences vidéo basée sur des hiérarchies de partitions*. PhD thesis, CMM, Ecole des Mines de Paris, 2001.