Scene Segmentation
Using Polarisation Information

Nitya Subramaniam

Submitted for the degree of Master of Science by Research

University of York
Computer Science

April 2010

Abstract

This thesis draws its hypothesis from the quantitative description of reflection and refraction of light at interfaces, given by the Fresnel theory. According to the theory, initially unpolarised light is partially polarised upon reflection from a smooth dielectric surface. Hence, variations in the polarisation properties of reflected and scattered light must indicate a change in properties of the reflecting surface. Assuming internal changes in refractive index can be neglected, any substantial change in the polarisation of light is thus indicative of a change in the material remitting the light. The contribution of this thesis is to develop a method for image segmentation based on surface material characteristics. The novel aspects of the method are the expression of intensity distribution as a function of the surface zenith and azimuth angles, and the expansion of this function using spherical harmonics to estimate surface characteristics. The method begins with estimating shape from polarisation using [1]. The surface normals obtained are used in combination with pixel intensities to form a three dimensional function that describes intensity changes in the image. This function is then expanded as a harmonic series and the constants of the expansion are used as features to characterize image regions and to segment the image based on the hypothesis. Experimental evidence is presented through analysis of polarisation images and segmentation of images of different objects.

Contents

List of Figures

Acknowledgements

I would like to express my thanks to my supervisor, Prof Edwin Hancock, for his continued support and advice during my research and writing up. His guidance has been invaluable not only in conducting my research but also in beginning to understand the elements of good research.

I would also like to thank Dr Richard Wilson and Dr Will Smith for their patience and support whenever I approached them. I owe special thanks to Ms Filomena Ottaway for her constant support as a friend and mentor during my time at York.

My thanks go to all the people in York and at the Department who helped me settle into a new country and helped me with little things that made life just a little easier and helped me not miss home as much as I would have without them.

Finally I thank my mother and brother, without whose support and constant encouragement I would never have made it this far.

Chapter 1

Introduction

Image segmentation is a problem that is critical to many applications in image processing and computer vision [2]. While most commonly encountered type of images are intensity images, other types of images include range (depth) images, magnetic resonance images and thermal images. There are may techniques for image segmentation in literature which rely on colour, gray scale intensity, shape in combination with some prior knowledge about the objects. Other approaches include methods based on Markov Random Fields and neural networks and fuzzy set theory [3]. However, in the absence of intensity or colour based cues, the task of segmenting an image into distinct objects poses a tough challenge. This problem is typical of most camouflaged objects and one that this thesis concentrates on solving.

Information in a gray scale image is encoded by variations in pixel intensity. The variations may be caused by changes in shape or reflectance of an object and also depend on the direction of illumination [4]. The relationship between incoming illumination (irradiance) and reflected light is commonly expressed using the bidirectional reflectance distribution function or BRDF, which is defined as the ratio of light in outgoing direction to incident irradiation [5]. For a point on a surface specified by position x with light coming in through solid angle $d\omega$, the BRDF is written as

$$\rho_{bd}(\theta_o, \theta_i, \phi_o, \phi_i) = \frac{L_o(x, \theta_o, \phi_o)}{L_i(x, \theta_i, \phi_i) \cos\theta_i d\omega} \tag{1.1}$$

Where L is the radiance, θ_i, ϕ_i are the incoming directions and θ_o, ϕ_o are the outgoing directions. A Lambertian surface is one whose BRDF is independent of incoming and outgoing directions and thus, constant for any given point on the surface. That is $L_o(x, \theta, \phi) = L_o(x)$. For such surfaces, the intensity of reflected light is given by Lambert's cosine law which states that the radiant intensity observed from the surface is directly proportional to the cosine of the angle between the observer's line of sight and the

surface normal. The radiance of reflected light for a Lambertian surface is expressed using the quantity radiosity $B(x)$, which is given by

$$B(x) = \int_0^{\pi/2} \int_0^{2\pi} L_o(x) \cos\theta \sin\theta \mathrm{d}\phi \mathrm{d}\theta \tag{1.2}$$

For a point light source that is far away compared to the size of the surface, assuming there is no ambient illumination [5] use a local shading model that defines the radiosity at a point x on the surface to be

$$B(x) = R(x)\mathbf{N}.\mathbf{S} = \cos\theta \tag{1.3}$$

where R is the surface albedo, N is the unit surface normal and S is a vector that represents the direction of the light source relative to the reflecting surface. Assuming the camera response is linear in the surface radiosity, the value of a pixel at location (x,y) is then given by

$$I(x, y) = kB(x, y) = kR(x)\mathbf{N}.\mathbf{S} = \cos\theta \tag{1.4}$$

The above equation succinctly expresses the main problems of physics based image analysis. Given an image, often the question in machine vision is to find an explanation such that the intensity changes $I(x, y)$ can be attributed to a change in albedo or shape, or a combination of both. This problem is heavily under-constrained as there may be many combinations of albedo, shape and illumination that result in a certain intensity pattern. Hence limiting assumptions become necessary to make the problem tractable. For instance, most simple shape recovery algorithms assume that the surface has a constant albedo.

Shape recovery is another well studied problem in computer vision. Common methods used for shape recovery include inference of shape from shading, multi-view techniques and others that fall under a broad category of methods called *Shape-from-X*. These include the use of texture, shading and stereo. A recent addition to this family of methods is the use of polarisation information for shape recovery. Polarisation vision has steadily gained use in the Computer Vision community over the last two decades. The areas in which polarisation finds applications range from medical imaging and industrial inspection to navigation. Recently a number of methods have been proposed that use cross-polarisation to remove specularities, or use specularities for shape recovery [6], [7], [8].

This thesis presents a method for image segmentation based on material surface characteristics. Assuming that camouflaged objects will be almost identical to the original in colour, the method presented uses only gray scale intensities to extract information about

the polarisation of light captured in an image. The features representing material characteristics are generated using the shape from diffuse polarisation method and analyzed using spherical harmonic decomposition.

The thesis begins with a review in Chapter 2 of recent work in polarisation vision and relevant literature in rough surface scattering theory, characterization of reflectance function and shape recovery. The chapter also draws attention to some interesting work in human shape perception and distinction of shape from albedo.

Chapter 3 presents a review of the theory and previous work that underpin this research. The chapter begins with a review of the basic physics behind scattering and polarisation in Section 3.2. Section 3.3 introduces the concepts of Fresnel theory and presents the important relations that are used in this work. Section 3.4 applies polarisation concepts to images and defines the key parameters used throughout the rest of the thesis and Section 3.5 presents an introduction to the work by Atkinson and Hancock [1] in recovering surface normals from polarisation information.

Chapter 4 draws on the theory in Chapter 4 to develop the results that are used in this research. To begin, Section 4.1 explains the motivation for the hypothesis explored in this thesis. The method used to estimate the polarisation parameters and its advantages are discussed in Section 4.2. The harmonic analysis used to compute feature descriptors is covered in Section 4.3 while Section 4.4 covers a brief overview of the graph theory relevant to this application before an overview of the normalized graph cuts method for segmentation.

Chapter 5 presents the experimental results of applying the method developed in Chapter 4 to various image sets. Sections 5.2 and 5.3 present experimental results for two chosen application areas to demonstrate the effectiveness of the developed method for surface segmentation based on material characteristics. Section 5.2 explores the first application in segmentation of image regions containing different materials. An obvious use of this method is in detection of camouflaged targets in images based on differences in polarisation characteristics of the camouflage material. Next, Section 5.3 presents results demonstrating a second application area in quality inspection of fruits and vegetables, for early detection of rots before they appear as skin lesions. Sections 5.4 and 5.5 discuss the accuracy of segmentation results obtained. The estimation error in the spherical harmonic expansion is analyzed by calculating the mean error in the image reconstructed by superposition of the spherical harmonic functions.

Finally Chapter 6 provides a brief summary of the work presented in this thesis. The chapter discusses the merits and drawbacks of this work and proposes directions for future research.

Chapter 2

Literature Review

This chapter presents a review of the recent work in the areas of texture, shape recovery and polarisation imaging and the research carried out n this thesis is placed in the context of current research in target detection.

2.1 Effect of Illumination on Texture

Although texture is a topic widely studied and applied to problems in computer vision, there is no agreement on a standard definition of what texture is. Julesz [9] introduced the idea of textons or texture elements to represent a description of texture. Julesz hypothesized that humans could pre-attentively distinguish between textures whose second order statistics were different. However there was no precise definition for how this may be applied to gray scale images. This drawback was removed by Malik et al in [10] which introduced an operational definition of textons using clustered filter responses for filters at various orientations and scales. This was then extended to a 3 dimensional definition by Leung and Malik in [11], by concatenating filter responses over different illumination angles. Dana et al made another effort to classify textures by building a database of bidirectional texture functions (BTF) [12], analogous to BRDF functions. They built a database of BTF and BRDF measurements for various common textures. The database was built from images of sample textures for a range of illumination angles. However the samples for each texture occupied a large amount of space.

The perception of a texture and its analysis in the spatial domain are strongly related to viewpoint. Chantler noted in [4] that the perceived texture of and object is also strongly dependent on illumination conditions. He published the first comprehensive study on the

effect of illumination on texture. Chantler uses empirical studies to show that the output of the Laws operators are significantly affected by illumination direction. Assuming a Lambertian reflectance and orthogonal projection, Chantler derives a BRDF function that included a dependence of illumination tilt angle. Chantler et al took this further in [13], developing a texture classifier that simultaneously estimates illumination direction. Recently Xu et al [14] have proposed a fractal based method for texture description that is invariant to changes in viewpoint and illumination.

2.2 Shape vs Shading

There have been several interesting experiments conducted to understand the nature shape and texture perception. This area of study is called psychophysics, the study of the relation between stimulus and response. Given an intensity image humans can intuitively assign it either a shape or an albedo interpretation. Freeman and Viola proposed a model [15] to distinguish between paint and shape.The model assigned prior probabilities based on data from a psychophysical computation and interpreted intensity changes as either shape or reflectance based on Bayesian probabilities. Bell and Freeman [16] train a classifier to separate steerable pyramid coefficients into shape or shading based on local image information. The advantage of this classifier over [15] was that it could separate components in images with both shape and reflectance changes. However a low frequency residual component in the algorithm causes limitations on the accuracy of recovery. Tappen et al [17] use colour information and gray scale patterns to classify image derivatives as being caused by shape or a change in surface reflectance. They then use belief propagation to resolve ambiguities and recover intrinsic shape and shading images from a single image. This method overcomes the residual component issue faced in [16].

In 2008, Padilla et al reported a new methodology for investigating the visually perceived properties of surface textures in [18]. The method uses synthetic textures rendered and rotated in real-time to study roughness perception and develops a method to estimate it. This method is used by Clarke et al in [19] to perform a psychophysical experiment on visual search patterns over synthetic textured backgrounds. The study shows that the saliency model of feature conspicuity [20] by Itti and Koch fails to adequately explain human performance in search tasks involving low contrast stimuli. In general however the saliency model by Itti and Koch is one of the most popular and a standard reference for all research on saliency. The model has also been extended further to include motion features to create a saliency model for videos.

2.3 Polarisation Vision

Although humans are insensitive to polarisation it contains a wealth of information about the reflecting surface. Sunlight traveling through the atmosphere is naturally polarised due to scattering by gas and dust particles. Light is also polarised when it passes underwater after refraction at an air-water interface, as in oceans and other water bodies. Reflection from a surface also causes light to be partially polarised. The information contained in polarisation states of light has been applied to develop Computer Vision techniques only very recently. Polarisation information is characterized by three basic parameters: intensity, polarisation degree and polarisation phase.

Wolff first introduced the idea of using polarisation for material discrimination in [21]. He noted that at points of specular reflection, the ratio of maximum to minimum intensity as viewed through a polarising filter can be used as a discriminator between dielectrics and metals. For metals this ratio is between 1.0 and 2.0 whereas for dielectrics the ratio is usually higher. The limiting assumption however is that the diffuse component of reflection be negligible compared to the specular component. To overcome this problem, Chen and Wolff developed a polarisation phase based method for discrimination between metals and dielectrics [22]. The method is based on the principle that linearly polarised light upon reflection from a conducting material is phase shifted to an elliptical polarised state. For non conducting dielectrics however, the reflected light remains linearly polarised. The high phase sensitivity of the technique makes it effective in both controlled and natural lighting. Polarisation has also been in a number of shape recovery methods. Rahmann and Canterakis use polarisation for reconstructing specular surfaces [7] and Miyazaki et al [8] use a polarisation-based method to model transparent objects. Atkinson and Hancock use the Fresnel theory to recover surface orientation [1] and estimate 2-dimensional BRDF [23]. The technique described in [1] is used to compute the surface normals that for the intensity function presented in this thesis.

2.4 Conclusions

Current research in the field of texture analysis and polarisation vision has established the ground work necessary for determining object shape from polarisation. Methods for highly accurate discrimination between metals and dielectrics based on polarisation phase are also available. However the use of polarisation for material characterization within dielectrics is an area of computer vision that is still largely unexplored. This thesis aims to

fill this gap by proposing a method for discriminating between different dielectric materials using a combination of intensity, polarisation phase and degree. Surface texture and illumination direction encoded in the polarisation state of diffusely reflected light are used to discriminate among dielectric materials based on their index of refraction.

Chapter 3

Polarisation Information

3.1 Introduction

This chapter provides the foundation to the theory and experiments discussed further in this thesis. The chapter covers the fundamentals of light reflection and the theoretical explanation behind polarisation of light upon reflection from an interface. This background theory is placed in the context of surface reflectance analysis from polarisation information contained in scattered light. A brief overview of the Fresnel theory of light is provided before linking it with the work by Atkinson and Hancock in recovering surface normals [1] and multiple viewpoint reconstruction [24]. Finally the last section gives a demonstration of the motivation behind the study of polarisation information contained in light reflected off material surfaces, and its applicability for image segmentation based on surface material characteristics.

3.2 Scattering and Polarisation

When light is incident on a homogeneous material the time-varying electromagnetic field creates oscillating atomic dipoles, which in turn generate secondary wavelets. This process is called scattering [25], [26]. Within a homogeneous dielectric material the scattered waves reinforce each other in the forward direction resulting in transmission and propagation of light through a medium. Waves traveling in every other direction are cancel out because of out-of-phase waves from neighboring atoms. However at an interface between two materials of different refractive index, some light is always scattered in the backward direction because there an imbalance of atomic oscillators. This phenomenon is called reflection. The incident light that is not reflected continues to propagate across the interface

and into the second material, although with a phase shift caused by the change in refractive index. This phenomenon is called refraction.

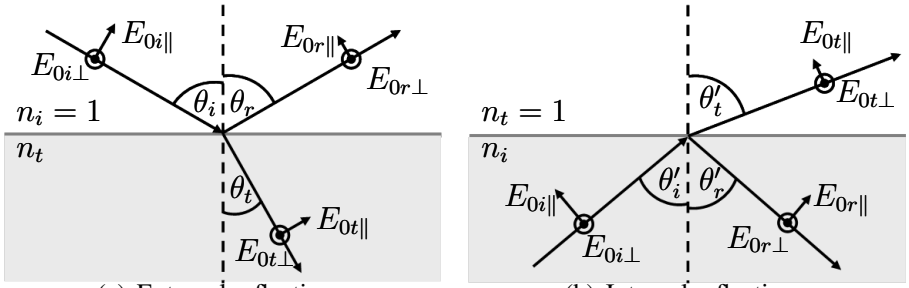

(a) External reflection (b) Internal reflection

Figure 3.1: Reflection of light at an interface between air and an optically denser material.

In the representation of Figure 3.1, the rays represent the direction of propagation of each wave. On extension of the concept to a two dimensional boundary, the plane that contains the incident, reflected and refracted rays (or propagation vectors of the respective waves) is called the plane of specular incidence. The electric field in each of these waves is resolved into two components, $E_\parallel$ in the plane of incidence and $E_\perp$ perpendicular to it. Consider Figure 3.1(a) for the case of external reflection at the boundary. We know from the Snell's law that the angles of incidence and refraction, θ_i and θ_t are related as follows:

$$n_i \sin\theta_i = n_t \sin\theta_t \tag{3.1}$$

where where n_i and n_t are the refractive indices of the first and second medium. If the values of n_i and n_t are known, for any given angle of incidence the above relation can be used to determine the direction of specular reflection.

In a simple model [25] of light reflection from a non-conducting and non-magnetic surface resolution of the electric field into $E_\parallel$ and $E_\perp$ can be used to explain the polarisation of light upon reflection. When the incident E-field is perpendicular to the plane of incidence, the bound electrons are driven into oscillation by the incident wave creating dipoles. These re-radiate creating a reflected and a refracted wave, both of which are polarised normal to the plane of incidence. However with an incident field in the plane of incidence, the atomic dipoles are driven into oscillation by the refracted wave and a component of this wave appears in the reflected direction. Since the direction of reflection is at an angle from the dipole axis, the electric field intensity in the reflected direction is significantly attenuated. Thus, combining the above two cases, incident unpolarised light upon reflected is partially polarised in a direction perpendicular to the plane of incidence. At a certain angle of incidence, when sum of the angles of incidence and refraction that is $\theta_i + \theta_t$ equals 90° the reflected light direction is such that the parallel component electric

field is completely extinguished. This is called the Brewster's angle, when light incident at the polarising angle $\theta_i = \theta_p$ results in completely polarised reflected light. Mathematically the polarising angle is expressed as $\tan\theta_p = n_t/n_i$.

3.3 Fresnel Theory of Light

The Fresnel theory of light can be used to understand the way in which polarised light interacts with surfaces. If E_{0i} and E_{0r} are the peak amplitudes of the electric field in the incident and reflected directions, the theory expresses the relation between the amplitudes of the reflected wave and the incident light wave in terms of the material constants and the angle of incidence. By starting with applying boundary conditions to the Maxwell's equations, the ratio of the amplitude of reflected field to that of the incident field for the component perpendicular to the plane of incidence is derived by the Fresnel theory as:

$$\left(\frac{E_{0r}}{E_{0i}}\right)_{\perp} = \frac{\frac{n_i}{\mu_i}\cos\theta_i - \frac{n_t}{\mu_t}\cos\theta_t}{\frac{n_i}{\mu_i}\cos\theta_i + \frac{n_t}{\mu_t}\cos\theta_t} \tag{3.2}$$

Here $\mu_i$and μ_t the magnetic permeabilities of the first and second medium.

For dielectrics, the permeabilities $\mu_i \approx \mu_t \approx \mu_0$, the permeability of free space. Under this assumption, (3.2) reduces to

$$r_{\perp} \equiv \left(\frac{E_{0r}}{E_{0i}}\right)_{\perp} = \frac{n_i\cos\theta_i - n_t\cos\theta_t}{n_i\cos\theta_i + n_t\cos\theta_t} \tag{3.3}$$

Similarly, the ratio of the components in the plane of incidence is then given by:

$$r_{\parallel} \equiv \left(\frac{E_{0r}}{E_{0i}}\right)_{\parallel} = \frac{n_t\cos\theta_i - n_i\cos\theta_t}{n_i\cos\theta_t + n_t\cos\theta_i} \tag{3.4}$$

The coefficients $r_{\perp}$ and $r_{\parallel}$ are called the Fresnel amplitude coefficients. The expressions (3.3) and (3.4) apply to any linear isotropic and homogeneous media. However the quantity measured by most light sensing devices is proportional not to the wave amplitude but the intensity or radiant flux density, which is in turn proportional to the square of the electric field [25]. Thus, the intensity coefficients can be derived as:

$$R_{\perp}(n,\theta) = (r_{\perp})^2 = \left(\frac{E_{0r}}{E_{0i}}\right)_{\perp},$$

$$R_{\parallel}(n,\theta) = (r_{\parallel})^2 = \left(\frac{E_{0r}}{E_{0i}}\right)_{\parallel} \tag{3.5}$$

The discussion above was focused on the polarisation phenomenon in dielectrics. For metals however the situation becomes more complicated due to their higher conductivity.

The presence of free electrons causes the time varying electromagnetic field of incident light to induce surface currents. There is also severe attenuation of transmitted light due to the conversion of electromagnetic energy to heat in conductors, which renders metals practically opaque in the visible wavelengths [26]. These phenomena make the electromagnetic analysis of metal surfaces less tractable. Hence the scope of analysis in this work is restricted to dielectrics.

3.4 Polarisation Image

A dielectric surface may polarise incident light in two different ways. In the case of specular reflection, initially polarised light is reflected from the air-dielectric boundary in the specular direction. In case of diffuse reflection, initially unpolarised light is refracted into the surface. The remitted light acquires a spontaneous polarisation due to multiple reflections and refraction at the dielectric-air surface. In both cases, the zenith angle of the reflected or remitted light is determined by the degree of polarisation while the azimuth angle is determined by the phase angle of light. However the relations between the surface normal angles and the polarisation is different in each case.

It has been established in the previous section that dielectric surfaces partially polarise light upon scattering. When the scattered light is passed through a polarising filter placed in front of the camera, the intensity of light observed though it fluctuates as a sinusoidal function of the orientation of the polariser. Thus the polarisation components of the remitted light can be resolved into components parallel and perpendicular to the plane of incidence by taking images of the object for different orientations of the filter. The parameters that represent the partial linear polarisation of incident light by reflection from the object surface are the minimum and maximum observed intensities I_{max} and I_{min}, and the phase ϕ of the transmitted sinusoid with respect to the polariser orientation α_p.

$$I(\alpha_p) = \frac{(I_{max} + I_{min})}{2} + \frac{(I_{max} - I_{min})}{2}\cos(2\alpha_p - 2\phi) \tag{3.6}$$

The radiant flux intensity or radiance of the light $I(\alpha_p)$ undergoing reflection at the surface is attenuated according to the Fresnel intensity coefficients defined in (3.5), which take on values between 0 and 1 inclusive. At an air-surface boundary, the relative refractive index n can be written as as $n \equiv n_t/n_i = n_t$ since the refractive index of air at room temperature is approximately 1. Assuming initially unpolarised light incident on a surface, the intensity bounds of the transmitted light can be expressed in terms of the specular component of reflectance I_s and the Fresnel intensity coefficients for reflection in the following

expression [27]:

$$I_{max} = \frac{R_\perp(n, \theta_i)}{R_\parallel(n, \theta_i) + R_\perp(n, \theta_i)} I_s,$$
$$I_{min} = \frac{R_\parallel(n, \theta_i)}{R_\parallel(n, \theta_i) + R_\perp(n, \theta_i)} I_s \tag{3.7}$$

For diffuse reflection the formula is similar although not as elegant. This is because the polarisation is the result of scattering within the material and refraction at the surface-air boundary. Hence the relative refractive index is now $1/n$. The Fresnel coefficients are now given by $R_\perp(1/n, \theta_i')$ and $R_\parallel(1/n, \theta_i')$, where θ_i' is the internal angle of incidence. The maximum and minimum intensities can hence be written in terms of the diffuse reflectance I_d as

$$I_{max} = \frac{1 - R_\parallel(1/n, \theta_i')}{2 - R_\parallel(1/n, \theta_i') - R_\perp(1/n, \theta_i')} I_d,$$
$$I_{min} = \frac{1 - R_\perp(1/n, \theta_i')}{2 - R_\parallel(1/n, \theta_i') - R_\perp(1/n, \theta_i')} I_d \tag{3.8}$$

Although there are many ways of defining the degree of polarisation, one way to define the quantity is as a ratio of the intensity of polarised light to the total light intensity. In terms of the minimum (I_{min}) and maximum (I_{max}) intensities observed as a polarising filter is rotated through a full cycle, this is given by:

$$\rho = \frac{(I_{max} - I_{min})}{(I_{max} + I_{min})} \tag{3.9}$$

which varies from 0 to 1 inclusive, indicating partial polarisation states ranging from completely unpolarised to completely polarised light.

3.5 Estimation of surface normals

For fixed light source direction and approximately planar samples, provided that the range of refractive indices for different materials in a scene is limited, the angular distribution of reflected or remitted light can be estimated from the polarisation image.

From the Fresnel theory it is straightforward to show that the azimuth angle for reflected polarised light or remitted diffusely polarised light is equal to the phase angle ϕ [1]. As established in Section 3.2 reflected light is attenuated to a greater extent in the parallel plane than in the normal plane relative to incidence. By extension, maximum transmission through a polariser occurs when the polariser is oriented at $90°$ to the azimuth angle ϕ of the surface normal.

Equations (3.6) through (3.9) can be applied with the Fresnel theory to analyze the distribution of reflectance from approximately planar samples of different material. This section presents the review of the relations required to analyze the distribution of reflectance from approximately planar samples of different materials. From the above results, substituting (3.7) into (3.9) gives an expression for the specular degree of polarisation in terms of the Fresnel reflection coefficients:

$$\rho_s = \frac{R_\perp(n, \theta_i) - R_\parallel(n, \theta_i)}{R_\perp(n, \theta_i) + R_\parallel(n, \theta_i)} \tag{3.10}$$

From the geometry of Figure 3.1(a), the zenith angle of the surface normal with respect to incident light direction is given by $\theta = \theta_i$. If (3.10) is expanded further by substituting for the Fresnel coefficients, the result is the following expression for ρ_s in terms of the zenith angle θ and the refractive index n

$$\rho_s = \frac{2\sin^2\theta\cos\theta\sqrt{n^2 - \sin^2\theta}}{n^2 - \sin^2\theta - n^2\sin^2\theta + 2\sin^4\theta} \tag{3.11}$$

Similarly, by subsitituting (3.8) in to the definition of degree of polarisation in (3.9) the expression for diffuse degree of polarisation ρ_d is derived as

$$\rho_d = \frac{R_\perp(1/n, \theta_i') - R_\parallel(1/n, \theta_i')}{2 - R_\perp(1/n, \theta_i') - R_\parallel(1/n, \theta_i')} \tag{3.12}$$

Using Snells's law for the refraction of light from the material boundary into air, the internal angle of incidence can be replaced with the exitance angle, which is the observed angle at which the light is scattered at the material surface. This angle θ_t' is the zenith angle θ of the surface normal with respect to the remitted light direction. As in the previous case, expanding the Fresnel coefficients and simplifying the resultant expression gives the relation between the diffuse degree of polarisation, the zenith angle and the refractive index.

$$\rho_d = \frac{(n - 1/n)^2\sin^2\theta}{2 + 2n^2 - (1 + 1/n)^2\sin^2\theta + 4\cos\theta\sqrt{n^2 - \sin^2\theta}} \tag{3.13}$$

Provided it is known whether the observations measure the specular polarisation of reflected polarised light, or the diffuse polarisation of remitted initially unpolarised light, the zenith angle of light with respect to the surface normal can then be calculated from (3.13) and (3.11). This is valid under the assumption that the range of refractive index is small, and can be treated as a constant. These results have been used in further work by Atkinson and Hancock in shape and BRDF recovery [23, 24, 28]. Further to this, if data is collected from multiple light source directions, a set of simultaneous non-linear equations can be solved to infer both refractive index and zenith angle of the surface normal. Figure 3.2 shows the degrees of diffuse and specular polarisation plotted against

the zenith angle θ, as given by (3.13) and (3.11). An inspection of the plots reveals that the range of values taken by diffuse polarisation degree is substantially lower than that taken by the specular polarisation, for the same zenith angles. This seems to match with what one would intuitively expect. Specifically, this is because diffuse polarisation arise from a weak polarisation by refraction at the material surface according to Fresnel theory. Specular polarisation on the other hand arises typically from specular reflection of initially polarised light.

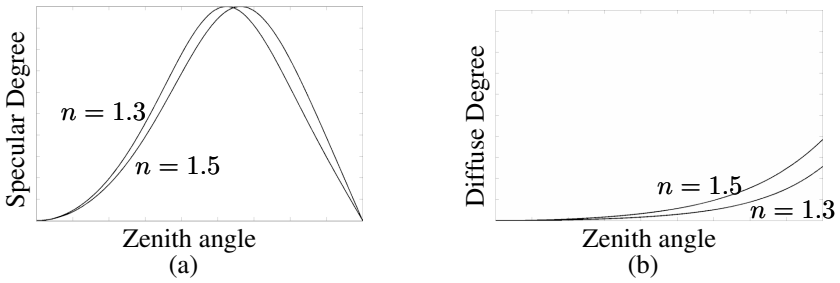

Figure 3.2: Polarisation degree versus zenith angle of surface normal for (b) diffuse and (a) specular reflection.

Now it is more convenient to rewrite (3.6) in terms of the mean-intensity $\hat{I} = 1/2(I_{max} + I_{min})$ and the degree of polarisation ρ as defined in (3.9), giving:

$$I(\alpha_p) = \hat{I}(1 + \rho\cos(2\alpha_p - 2\phi)) \tag{3.14}$$

The three components of the polarisation information as expressed in (3.14) are the mean intensity $\hat{I}$, the degree of polarisation ρ and the polarisation phase ϕ. These shall henceforth be referred to collectively as the polarisation image. It is worth nothing that there is an ambiguity in the determination of phase. Polarising filters cannot distinguish between two states of polarisation that are 180° apart. This is because the intensity of light remitted through a filter is a function of twice the polariser orientation angle, that is 2ϕ. This leaves an ambiguity of 180° in determining the phase of polarisation. This ambiguity can be resolved using simple techniques in computer vision [1]. However there is no attempt to resolve this ambiguity in this work as it does not affect the analysis.

3.6 Conclusions

The Fresnel theory of light provides a quantitative analysis of the electromagnetic interaction of light at an interface. From the theory reviewed, it can be seen that the polarisation degree depends on the maximum and minimum intensities observed through a polarising

filter. The azimuth angle of the surface normal with respect to light source direction equals the polarisation phase. Also, for the the same zenith angle, the degree of diffuse polarisation is higher than specular polarisation. These results are used in the following chapter to derive the relations used in the analysis of images for material-based segmentation.

Chapter 4

Surface Segmentation using Polarisation

The results of the Fresnel theory have been use to develop a number of machine vision techniques including methods for surface quality inspection [29], [30] and surface shape recovery [1], [24], [8], [7]. Polarisation can also be used to infer information concerning the reflectance properties of surfaces. For instance, Atkinson and Hancock have shown in [23] how diffuse polarisation can be used to estimate the bidirectional reflectance function. However, their method is computationally demanding, using simulated annealing to estimate the BRDF. The work presented in this thesis takes a simpler approach to reflectance characterization.

The motivation behind analyzing polarisation information for material segmentation was established in Chapter 3. Observation of intensity distributions of different materials revealed distinct patterns depending on material type. This chapter takes the above observation forward and develops the theoretical results used in this thesis for image segmentation based on differences in material surface properties.

4.1 Motivation

The observation underpinning this thesis is that polarisation information in light remitted from an object's surface allows the measurement of its reflectance properties. Under the restrictions of sample planarity and slowly varying refractive index, the polarisation state of remitted light represents the distribution of mean intensity $\hat{I}$ with respect to the zenith and azimuth angles of remitted light. To provide some illustrative motivation, Figure 4.1 shows a scatter plot of the intensity versus the degree of polarisation and versus surface azimuth angle for real and plastic leaves.

There are a number of features to note from the plot. First, the distributions are quite

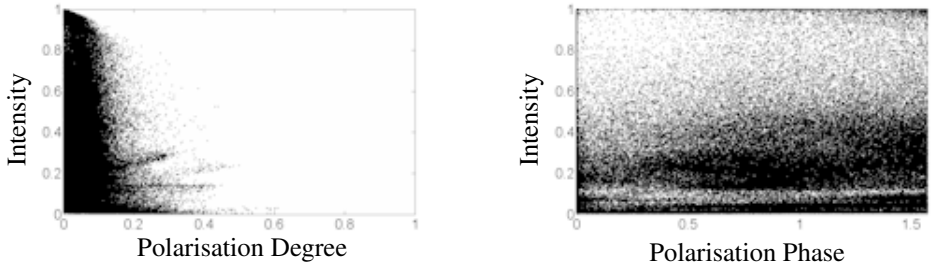

(a) Plastic leaf in unpolarised light

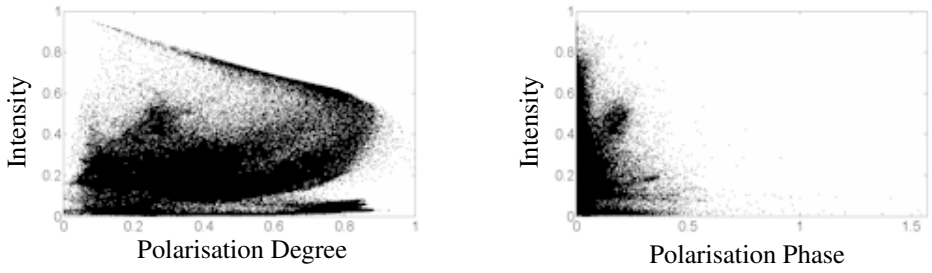

(b) Natural leaf in unpolarised light

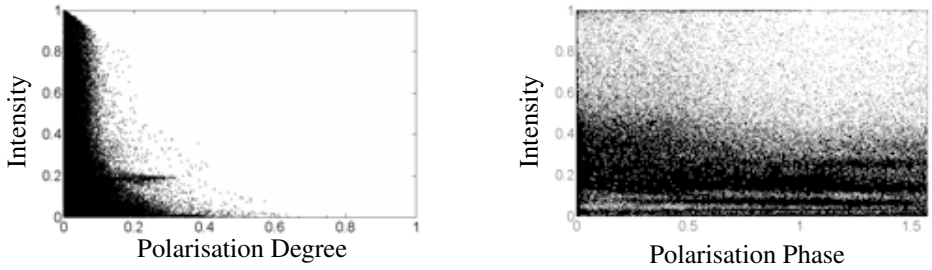

(c) Plastic leaf in polarised light

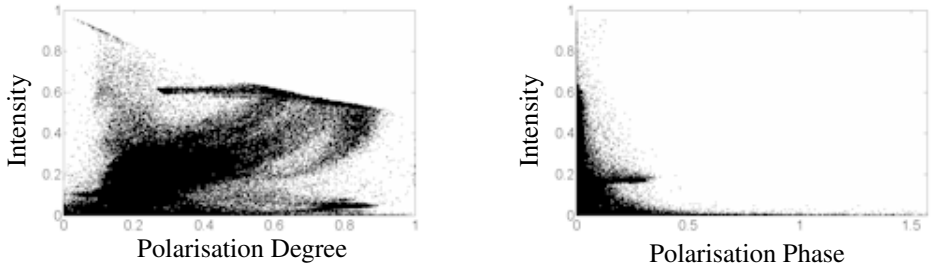

(d) Natural leaf in polarised light

Figure 4.1: Scatter plot: Distribution of intensity function

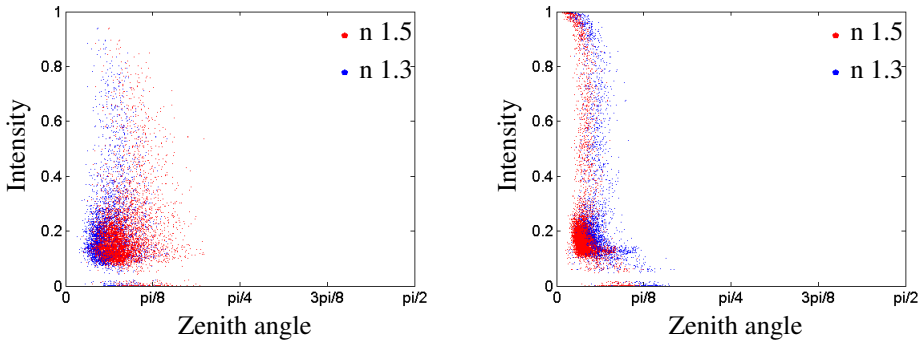

(a) Real leaves (b) Plastic leaves

Figure 4.2: Shift in scatter plot of zenith angle against intensity, for a change in assumed refractive index

different for the two materials. This is attributed to the fact that natural leaves have a layered sub-surface structure, which affects distribution of remitted light through subsurface refraction according to Snell's law. Artificial leaves are frequently made from synthetic polymers which have a more jumbled molecular structure and hence do not exhibit the same reflectance properties.

Second, when the refractive index is changed within the known range for dielectrics, there is a small shift in the plots at all zenith angles, demonstrated in Figure 4.2. This is because as established in Section 3.4 equations (3.12) and (3.10), the diffuse degree of polarisation is a function of the zenith angle and the refractive index. However since the shift is uniform across all angles, the effect of approximating refractive index in calculations can be neglected.

Figure 4.1 demonstrates a difference in the patterns exhibited by light returned from plastic and natural leaves. There is a difference in intensity plots with respect to both polarisation degree and phase. As explained earlier, the difference in plots arises from the structural difference between different materials in the studied scene.

This work attempts to exploit the difference in material structure as presented in the properties of captured light to segment a scene into areas of different materials. One immediately relevant application for this study will be in the detection of military camouflage from amongst natural foliage. This may be used in military surveillance and target detection systems. Another application is in the detection of the quality of natural fruits. Studies carried out as part of this research showed that degradation in fruits and vegetables due to aging and natural enzyme action has an effect on the polarisation properties of remitted light. This property can be exploited to build automatic quality inspection systems for food products. There are several other possible applications for the method for example,

in airport security screening, automatic defects inspection and automobile lane guidance systems.

Figure 4.3 shows the values of polarisation triplet (average intensity, polarisation degree and phase) at various points on the surface, plotted as an image, for plastic leaves in unpolarised light in (a) and polarised light in (b). The values range from 0 to 1 for normalized average intensity and degree of polarisation, and from 0 to $\pi/2$ for polarisation phase.

It was noted in Section 3.5 that the diffuse polarisation degree takes lower values than specular for the same zenith angle. The polarisation image sets in Figure 4.3(a) and (b) for example also reveal that polarisation degree is higher in initially polarised light, as expected. This is in agreement with the expectation that spontaneous polarisation of initially unpolarised light is much weaker than the polarisation degree of remitted light that is polarised before incidence.

It is also observed that as the incidence tends to normal (i.e. the more the angle between the surface normal and incident light tends to 90 degrees), the spontaneous depolarisation of incident light increases. The depolarisation is also greater at the veins when compared to the leaf surface. This may be attributed to a difference in structure of veins. The veins in plastic leaves also contain a thicker layer of plastic as compared to the remaining leaf surface. This changes the nature of subsurface scattering in the veins.

The polarisation phase closely represents the shape of the object surface. The values of polarisation phase vary more widely in initially unpolarised incident light than in polarised incident light. This may be attributed to the random nature of electric field orientations in unpolarised light. The small phase shifts in specular polarisation are caused by spontaneous depolarisation of the incident light by the object surface, which is much weaker than the polarisation of incident light. The polarisation images demonstrate that the phase accurately represents the azimuth angle of the surface normal with respect to incident light direction, as predicted by equation (3.14).

4.2 Statistical Moments

The fundamentals of polarisation image analysis were covered in Section 3.4. Light scattered from surfaces acquires an instantaneous polarisation because of refraction. For images captured through a polarising filter at different orientations, the components of the

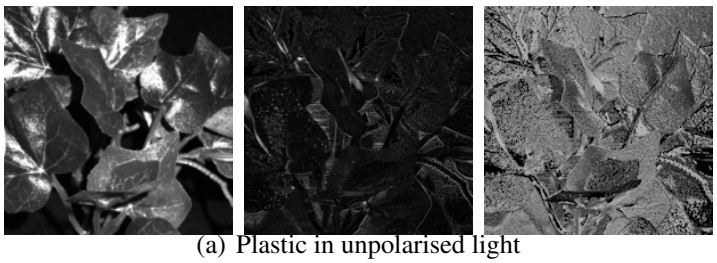

(a) Plastic in unpolarised light

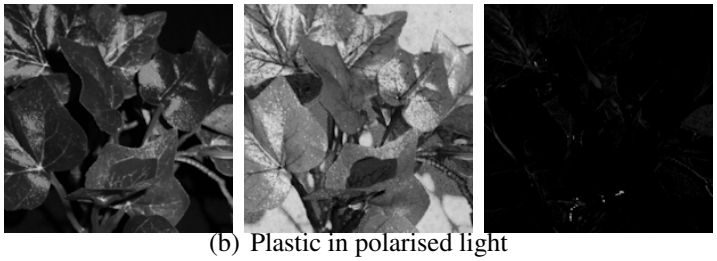

(b) Plastic in polarised light

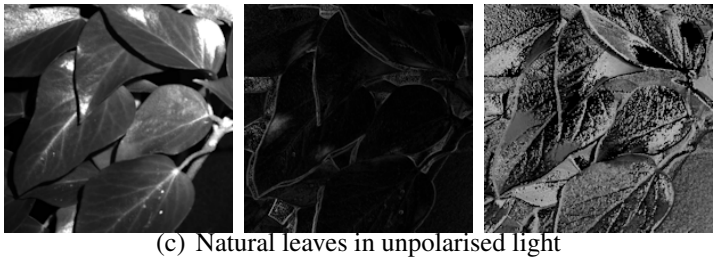

(c) Natural leaves in unpolarised light

(d) Natural leaves in polarised light

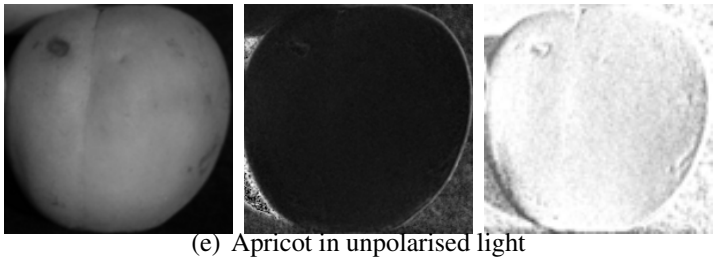

(e) Apricot in unpolarised light

Figure 4.3: Polarisation image components (L-R) $\hat{I}$, ρ, ϕ for different materials

polarisation image are related by (3.14). The three parameter polarisation image $(\hat{I}, \theta, \phi)$ recorded for a given scene can be treated as a distribution of data points containing information about the scene. Since the moment descriptors of a variable describe the nature of its distribution, moment descriptors of the polarisation image parameters can be used to represent this scene information.

Consider a variable x, described by a probability density function $f(x)$. By definition [31], the kth moment of the variable is given by

$$\mu_k = \mathbb{E}(x^k) = \int_{-\infty}^{\infty} x^k f(x)\ \mathrm{d}x \tag{4.1}$$

In the case of x being a discrete variable the moment is expressed by a summation. The centralized moments are then the expected values of the variable about its mean $\hat{x}$. Since the mean is also the first raw moment, the first centralized moment is zero and all centralized moments can be expressed as

$$\mathbb{E}\left((x - \hat{x})^k\right) = \frac{1}{N}\sum_{i=1}^{N}(x_i - \hat{x})^k \tag{4.2}$$

In other words, the kth central moment of x is the expected value of x^k about its mean, $\hat{x}$.

4.2.1 Moment Estimators of Polarisation

Suppose that we take N equally spaced images of an object through a polarising filter, with the polariser angle indexed as $p = 1, 2, ..., N$. Recall that the transmitted radiance sinusoid is a function of the recorded image intensity in terms of the polarisation image parameters given by (3.14), repeated here for ease of reference:

$$I(\alpha_p) = \hat{I}(1 + \rho\cos(2\alpha_p - 2\phi))$$

This may be rearranged as follows :

$$I(\alpha_p)/\hat{I} = 1 + \rho\cos(2\alpha_p - 2\phi) \tag{4.3}$$

$$I(\alpha_p)/\hat{I} - 1 = \rho\cos(2\alpha_p - 2\phi) \tag{4.4}$$

$$\frac{I(\alpha_p) - \hat{I}}{\hat{I}} = \rho\cos(2\alpha_p - 2\phi) \tag{4.5}$$

Now let $x_p = (I(\alpha_p) - \hat{I})/\hat{I}$. Substituting this in (4.5) and squaring both sides of the equation gives

$$x_p^2 = \rho^2\cos^2(2\alpha_p - 2\phi) \tag{4.6}$$

Since the variable x has zero mean, the second central moment of x is $1/N \sum_{p=1}^{N} /x_p^2$. The right hand side of the equation (4.6) now becomes an integral of the expression with respect to α. The limits of the integral are then derived from the total angle that polarising filter is rotated through to record the observations. Since the transmitted radiance sinusoid is a function of $2\alpha_p$, the function completes one cycle in π radians and further readings would only present repeated information. Therefore the limits of the integral are set from 0 to π giving

$$\sigma^2 = 1/N \sum_{i=1}^{n} x_p^2 = \rho^2 \int_0^{\pi} \cos^2(2\alpha - 2\phi)\ \mathrm{d}\alpha \tag{4.7}$$

Upon integration in (4.7), the right hand side reduces to $\pi/2\,\rho^2$. This gives us the moment estimates for image intensity and degree of polarisation as

$$\hat{I} = 1/N \sum_{p=1}^{N} I(\alpha_p) \tag{4.8}$$

$$\rho = \sqrt{2/\pi}\,\sigma \tag{4.9}$$

To derive an expression for the polarisation phase, both sides of (4.5) are multiplied by $\cos 2\alpha$ and the expected value is calculated on both sides. Let $y = (I(\alpha_p) - \hat{I})/\hat{I} \cos 2\alpha_p = x_p \cos 2\alpha_p$.

$$\hat{y} = \rho \int_0^{\pi} \cos(2\alpha - 2\phi) \cos 2\alpha\ \mathrm{d}\alpha \tag{4.10}$$

Integrating the left hand side with respect to θ eliminates the variable from the expression and reduces it to $\pi\rho\cos 2\phi$. Thus the moment estimate for the phase of polarisation is given by

$$\hat{y} = \pi\rho\cos 2\phi \tag{4.11}$$

$$\phi = \frac{1}{2}\cos^{-1}\left(\frac{2\hat{y}}{\pi\rho}\right) \tag{4.12}$$

Equations (4.8), (4.9) and (4.12) define the moment descriptors of the polarisation image used in this work. These represent the characteristics of the intensity distribution as a function of material characteristics The initial estimates of polarisation parameters obtained from moments are then fitted to the data using the Levenberg-Marquardt method to give a robust estimate of the degree and phase of polarisation in the scattered light.

4.3 Feature Descriptors

A number of object recognition problems involve looking for image windows that match an expected shape, texture or colour expectation. This is fundamentally a problem of se-

lecting groups of pixels in the image which match the required description and discarding those which do not. However it may not be possible to make the decision simply based on pixel gray scale or colour information. In such cases, it helps to have a compact representation of pixels that brings out the properties of interest. Obtaining this representation of the image is called segmentation or grouping [5]. The process of analyzing information contained in the image and producing a description that makes the segmentation task simpler is called feature generation. The important properties of a good feature representation are that there should be relatively few components in the feature so as to make the segmentation task computationally efficient. Secondly the feature should be a good representation of the characteristics that are of interest in the image.

The choice of feature descriptors is largely subjective and depends on the type of segmentation task at hand. The aim of this work is to enable segmentation of a given scene into objects of different material composition. Hence a good feature descriptor would need to adequately represent the reflectance properties of the objects as observed in the polarisation image.

4.3.1 Laplace Spherical Harmonics

The relations in Section 3.5 presented a representation of the surface reflectance properties in terms of the observed image intensity. Looking at (3.6), it is easy to see that the image intensity can be perceived as a 3-D function $I(\rho, \theta, \phi)$ dependent on the polarisation degree, zenith and azimuth angles of the surface normal. The choice of spherical harmonics to represent the function $I(\theta, \phi)$ takes advantage of the spherical symmetry that can be observed in the representation of image intensity as a function of spherical coordinate variables. At that note, a small diversion is taken to discuss some fundamental theory behind the chosen feature representation.

In applied mathematics, the spherical harmonics $Y_l^m(\theta, \phi)$ are the angular part of a set of solutions to the type of partial differential equations known and Laplace's equation. Laplace's equations find application in a number of practical problems that arise in the study of potentials. If a function f can be written in a form with separated angular variables θ and ϕ, that is $f = \Phi^m(\phi)\Theta_l^m(\theta)$, then the solution to the Laplace differential equation on f gives the angular functions

$$\Phi^m(\phi) = Ae^{-im\phi} + Be^{im\phi} \tag{4.13}$$

$$\Theta_l^m(\theta) = P_l^m(\cos\theta), \tag{4.14}$$

where l is the polynomial degree, m takes integer values from $-l$ to l and $P_l^m(z)$ are the associated Legendre polynomials [32]. The spherical harmonics are then arrived at by combining $\Phi(\phi)$ and $\Theta(\theta)$

$$Y_l^m(\theta,\phi) = \sqrt{\frac{2l+1}{4\pi}\frac{(l-m)!}{(l+m)!}} P_l^m(\cos\theta) e^{im\phi}. \tag{4.15}$$

with the normalization coefficient chosen such that the integral of the magnitude of the functions over the sphere equals the Kronecker delta function:

$$\begin{aligned}\int_0^{2\pi}\int_0^{\pi} Y_l^m(\theta,\phi)\bar{Y}_{l'}^{m'}(\theta,\phi)\sin\theta \mathrm{d}\theta\mathrm{d}\phi &= \int_0^{2\pi}\int_0^{\pi} Y_l^m(\theta,\phi)\bar{Y}_{l'}^{m'}(\theta,\phi)\mathrm{d}(\cos\theta)\mathrm{d}\phi \\ &= \delta_{mm'}\delta_{ll'}\end{aligned}$$

From the expansion theorem [33] it follows that the superposition of the solutions gives the boundary solution to the Laplace's differential equation on the sphere. That is

$$f(\theta,\phi) = \sum_{l=0}^{\infty}\sum_{m=-l}^{l} a_{lm} Y_l^m(\theta,\phi) \tag{4.16}$$

Furthermore, the spherical harmonics $Y_l^m(\theta,\phi)$ form a compete set of linearly independent orthonormal functions on the sphere. Using this property, we can derive the spherical harmonic coefficients for the function $f(\theta,\phi)$ as

$$a_{l,m} = \int_0^{2\pi}\int_0^{\pi} f(\theta,\phi) Y_l^m(\theta,\phi)\sin\theta\,\mathrm{d}\theta\,\mathrm{d}\phi \tag{4.17}$$

4.3.2 Harmonic Analysis

The information in the polarisation image is comprised of a set of triples consisting of the mean image intensity $\hat{I}$, the degree of polarisation ρ and the polarisation phase ϕ. For an image consisting of M pixels indexed by i = 1, 2, ..., M, the polarisation image may be written as

$$P = \{(\hat{I}_i, \rho_i, \phi_i), i = 1, ..., M\} \tag{4.18}$$

Recall from (3.13) in Ch. 3, Section 3.5 that the degree of polarisation is a function of a constant refractive index n and the zenith angle θ. From this the above triple can be written as

$$D = \{(\hat{I}_i, \theta_i, \phi_i), i = 1, ..., M\} \tag{4.19}$$

To exploit the spherical symmetry in the triplets we suggest that the mean image intensity $\hat{I}$ be expressed in terms of θ and ϕ giving a discrete 2-D function of the zenith and

azimuth angles. This function $\hat{I}(\theta, \phi)$ can then be expressed as the superposition of the weighted orthonormal basis functions $Y_l^m(\theta, \phi)$ as follows

$$\hat{I}(\theta, \phi) = \sum_{l=1}^{\infty} \sum_{m=-l}^{l} a_{l,m} Y_l^m(\theta, \phi), \quad a \in \mathbb{R} \tag{4.20}$$

where $Y_l^m(\theta, \phi)$ are the spherical harmonic functions as defined in (4.15).

The expression for the spherical harmonic coefficients for a continuous function can be calculated using (4.17). Extending this to the discrete domain using a moments estimation, the spherical harmonic coefficients for the function in (4.20) can then be obtained as

$$a_{l,m} = \frac{1}{M} \sum_{i=1}^{M} \hat{I}_i \, Y_l^m(\theta_i, \phi_i) \tag{4.21}$$

Estimation of harmonic functions in previous literature includes residual fitting approaches by [34], [35] and spherical FFT by [36]. This work uses a MATLAB function to compute the Legendre polynomials and a moments based approach to estimate the coefficients $a_{l,m}$. The image is divided into windows and the average coefficients are calculated over each window. The window size is chosen to ensure that the intensity function is a reasonable representation of shape while taking care to not over-smooth the features.

4.3.3 Feature Selection

The object of feature selection is to find a smaller subset of the available features that accurately represent the the full feature set. The advantages of reducing dimensionality include a smaller feature vector which translates to lesser computation time. It also serves to consolidate information that may be repeated in different dimensions of the feature. Linear transformations of the feature set can be used to identify the dimensions that contain maximum information and also realign the feature space so that a clearer partition of the data points becomes visible. Principal Component Analysis (also known as the KarhunenLoève transform) is one such linear transform in which the resultant features are linear combinations of the original features. The objective of principal component analysis is using the original data points to construct a lower dimensional linear subspace that best explains the deviation of the features from the mean. This is a classical technique from statistical pattern recognition ([5], from [37], [38], [39]).

4.3.4 Affinity Measure

The distances between feature vectors in the multi-dimensional vector space are a good indication of how they may be clustered into relevant groups. There exist many metrics for calculating the distance between two vectors $\mathbf{x}$ and $\mathbf{y}$. For example the Euclidean distance between vectors is defined as the the L_2-norm:

$$d(\mathbf{x}, \mathbf{y}) = \sqrt{(\mathbf{x} - \mathbf{y})^T (\mathbf{x} - \mathbf{y})} = \sqrt{\sum_i (x_i - y_i)^2} \tag{4.22}$$

The city block or Manhattan distance is similarly defined to be the L_1-norm:

$$d(\mathbf{x}, \mathbf{y}) = \sqrt{\sum_i |x_i - y_i|} \tag{4.23}$$

However in certain cases, for example if the features are correlated, the Euclidean distance does not provide an accurate representation of the distribution. In such cases, the Mahalanobis distance provides a better measure of similarity because it takes into account the variance of data while calculating distance. The Mahalanobis distance uses the feature covariance matrix to compensate any inconsistencies in scaling of the features relative to each other and account for correlation between features. The distance metric is given as

$$d(\mathbf{x}, \mathbf{y}) = \sqrt{(\mathbf{x} - \mathbf{y})^T \Sigma^{-1} (\mathbf{x} - \mathbf{y})} \tag{4.24}$$

where Σ is the covariance matrix for all the feature vectors available.

To calculate the Mahalanobis distance between PCA-mapped spherical harmonic coefficient vectors, the algorithm commences by computing the variance matrix over the image. Since the image is divided into blocks in the feature vector calculation, the variance is also calculated over image blocks. Let the the image blocks be indexed by $k = 1, 2, ..., L$ and the k-th block have the feature vector A_k. The mean coefficient vector is then

$$\hat{A} = \frac{1}{L} \sum_{k=1}^{L} A_k \tag{4.25}$$

and the covariance matrix is

$$\Sigma_A = \frac{1}{L} \sum_{k=1}^{L} (A_k - \hat{A})(A_k - \hat{A})^T \tag{4.26}$$

The Mahalanobis distance between the coefficient vectors for the blocks indexed k_1 and k_2 is

$$D_{k_1,k_2} = (A_{k_1} - A_{k_2})^T \Sigma_A (A_{k_1} - A_{k_2}) \tag{4.27}$$

4.4 Graph Theoretic Clustering

Image segmentation or clustering can be approached as a graph partitioning problem. A good segmentation is then a problem of cutting a graph into good pieces. A graph is represented by its nodes and its edges, thus a graph G = (V,E) has vertices V and edges E. The edges may also be assigned weights, in which case the graph is called a weighted graph. The approach in graph theoretic clustering is to treat the weighted graph as a sparse matrix, with a row and column for each vertex. The elements in the matrix represent the weights between the edges represented by the row and column, e.g. $w(i, j)$ is the weight associated with the edge connecting vertices i and j. Image features can be associated with vertices in a weighted graph where the weights on the edges represent a measure of similarity, in other words the feature distance. The segmentation algorithm would then attempt to cut the graph in a way that edges with high weights are grouped together and edges with relatively low weights are removed, thus grouping the graph into sub-graphs with similar features.

4.4.1 Normalized Graph Cuts

In a graph G = (V,E) comprised of vertices V and edges E with $w(i, j)$ being the weight associated with the edge connecting vertices i and j, the degree of dissimilarity between two sections A and B can be computed as a sum of the weights on the edges that have been removed to separate the sections. This is called a cut and can be expressed mathematically as follows:

$$cut(A, B) = \sum_{u \in A, v \in B} w(u, v). \tag{4.28}$$

The minimum cut criterion proposed by Wu and Leahy [40] attempts to minimize the sum of all edge weights connecting any two partitions. This criterion however has a tendency to partition out small sets of points in the graph. Shi and Malik introduced a graph cut criterion [41] that computes the cost as a fraction of the total edge connections to all the nodes in the graph. The expression for the normalized cut is then

$$Ncut(A, B) = \frac{cut(A, B)}{assoc(A, V)} + \frac{cut(A, B)}{assoc(B, V)} \tag{4.29}$$

where the total connection of all nodes in, for example, A to all nodes in the graph is denoted as

$$assoc(A) = \sum_{u \in A, t \in V} w(u, t). \tag{4.30}$$

With this definition of the disassociation between the groups, the minimum cut is that which partitions groups that have a combination of the largest number of high weight edges with within each group and smallest number of edges between groups. Therefore cuts that partition out small groups of points do not have a small normalized cut. A graph with $N = |V|$ vertices is represented as an $N \times N$ symmetrical matrix as explained above, consisting of the edge weights $w(i, j)$. This is called an affinity matrix. The definition of affinity measures used to assign weights in a segmentation problem depends on the application. Affinity measures may be defined by distance, intensity, colour or texture. The degree matrix is then a diagonal matrix defined from the affinity matrix, such that each diagonal element is the sum of the weights on edges to that vertex from all others in the graph, that is

$$A_{ij} = w(i, j),\ D_{ii} = \sum_j A_{ij} \tag{4.31}$$

If $\mathbf{y}$ is a vector of length N, that holds the segmentation results for the graph such that an element of $\mathbf{y}$ is 1 if the node belongs to a segment and $-b$ if it does not, then the N-cut criterion can be rewritten as the minimization of

$$\frac{\mathbf{y}^T(D - A)\mathbf{y}}{\mathbf{y}^T D\mathbf{y}} \tag{4.32}$$

The solution to this problem is then found by solving the generalized eigenvalue system $(D - A)\mathbf{y} = \lambda D\mathbf{y}$. The second smallest eigenvector of this system is the real values solution for bi-partitioning the graph. The segments thus obtained are then recursively partitioned until the desired segmentation is achieved.

4.5 Process Overview

The following block diagram gives an overview of the segmentation process. The main stages of the segmentation process can be divided into the following parts:

- Image acquisition and registration
- Polarisation image generation
- Calculation of spherical harmonic coefficients
- Principal component analysis
- Affinity matrix computation

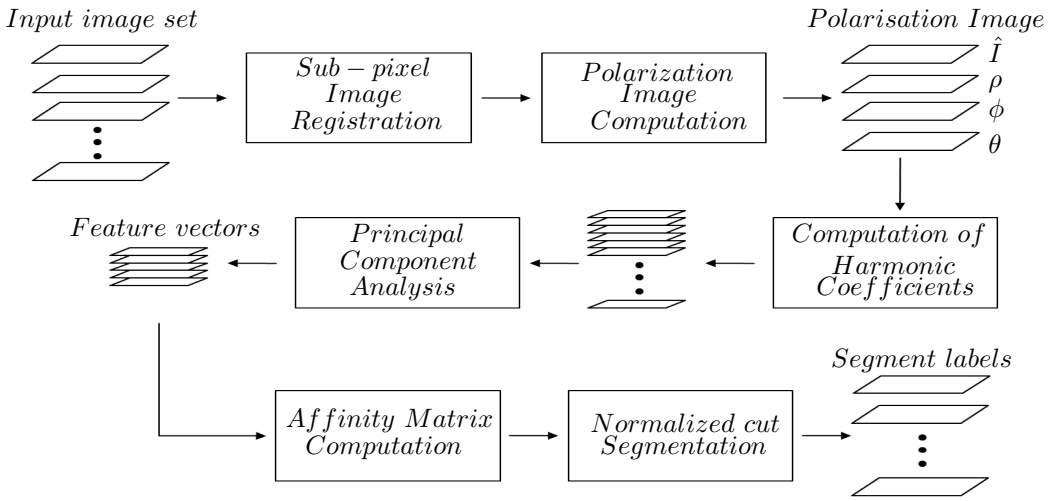

Figure 4.4: Process diagram for segmentation using the developed method.

- N-cut segmentation using Mahalanobis distance between feature vectors

Given the polarisation image, the proposed algorithm attempts to simultaneously segment and classify using the normalized cut method. Spherical harmonic decomposition of the intensity over image blocks yields the harmonic coefficient vectors $a_{l,m}$ as described in (4.21). These vectors are then transformed into low-dimension PCA-mapped feature vectors. This work uses the Mahalanobis distance between feature vectors as an affinity measure. The more different the feature vectors, the more likely that the windows having the feature belong to different segments. For an image consisting of L blocks, the Mahalanobis distances between the feature vectors for each block are arranged in an $L \times L$ block affinity matrix S with elements $S(k_1, k_2) = \exp[-D_{k_1,k_2}]$. The polarisation image is then segmented into regions by recursively applying Shi and Malik's [41] algorithm to the affinity matrix.

In practice the set of coefficients are estimated over non-overlapping pixel blocks of size $w \times w$, and the spherical harmonic expansion is truncated at a chosen value of l, with m varying from $-l$ to l. As a result the mean intensity distribution is each pixel block is parameterized by an $(l+1)^2$ element vector of spherical harmonic coefficients, $A = [a_{0,0}, a_{1,-1}, a_{1,0}, a_{1,1}, ..., a_{l,0}, a_{l,-1}, ..., a_{l,l}]^T$. Since application memory is a limited resource, there is a trade-off between window size w and degree of expansion l. The size of feature vectors calculable is also restricted by image size and resolution. Reliable segmentation results are obtained for values of x varying from 8 to 12 pixels and values of degree from 8 to 16.

4.6 Conclusions

The motivation behind this thesis is demonstrated in the scatter plots of intensity (Fig. 4.1). Since the degree of polarisation is a function of the zenith angle and the polarisation phase is equal to the azimuth angle, components of the polarisation image represent the spatial distribution of average image intensities. Figure 4.1 shows that the spatial distribution of intensity for foliage is different from that for plastic leaves. Since the distribution of normalized intensity spans a unit sphere, it can be represented using spherical harmonic functions. Hence image intensity is characterized using spherical harmonic coefficients. The coefficients calculated over blocks of image pixels are used to form a feature vector that brings out the differences in material surface properties. The distance between these feature vectors is then used to calculate an affinity matrix and normalized cuts are used to recursively segment the image. The next chapter demonstrates the segmentation results obtained using the method developed above. The technique is tested on images captured in different lighting conditions and for different camouflaged objects and the results are presented in the chapters that follow.

Chapter 5

Experimental Results

This chapter demonstrates the results obtained by applying the method developed in Chapter 4. The results show image segmentation in two main application areas. The first is material separation in scenes containing different objects, based on the polarisation properties of each material. The second is in quality inspection by segmenting fruits and vegetables to identify rotting before it manifests as skin lesions. Results shown include segmentation of images taken in polarised and upolarised light from a collimated helium light source as well as images taken outdoors in direct and diffuse sunlight.

5.1 Experimental set up

All indoor images have been captured in a darkened room with matte black walls and working surfaces to minimize noise caused by light scattered from objects other than those of interest. The test objects and the camera are positioned on the same axis, and a halogen source with light in the visible spectrum is positioned at approximately 15 degrees from the viewing axis. This arrangement is chosen in order to reduce specular reflection. Images are captured using a Nikon D200 camera, with fixed aperture size and exposure time. Linear polarising filters are placed in front of the camera and the halogen source in case of polarised incident light. The camera polaroid is rotated through increments of 10 degrees to give 19 images per scene, or through 30 degrees to give 7 images. The objects studied include natural and artificial leaves, man-made surfaces, military camouflage net and fruit.

5.2 Application I : Material-based Surface Segmentation

It has been established theoretically that polarisation information can be used to characterize material surfaces. This section provides the test results of the segmentation method applied to different scenes. Indoor images are captured in a dark room in both polarised and unpolarised light from a helium light source, with the remitted light being diffusely polarised in the first case and specularly in the second. Thus the effectiveness of the method is tested for both diffuse and specular polarisation. From the Fresnel theory reviewed in Section 3.3, the phase angle of remitted light is equal to the azimuth angle of the remitted light direction, and for materials with restricted ranges of refractive index the specular or diffuse polarisation determines the zenith angle. Based on this observation, the angular distribution of the mean intensity for remitted light is parameterized using spherical harmonics. The vectors of spherical harmonics are used to characterize varying surface material reflectance distributions, which in hypothesized to indicate a change in refractive index of the scattering material. The distances between the feature vectors is then used to segment the scene into different material patches.

The first segmentation demonstrated is for a scene of natural and plastic leaves. Figure 5.1(a) shows the test scene containing a bunch of artificial leaves in roughly the right half of the image and natural leaves in the left. The sub-figures (b) and (c) show components of polarisation image for scene images captured in unpolarised and polarised light from a collimated halogen lamp in the dark room.

It is evident from the figure that polarisation image captures information about the scattering surface. Because the remitted light is polarised spontaneously upon refraction from the surface, the polarisation degree captures edges and fine surface texture in unpolarised light. The degree of polarisation is a ratio of the intensities of polarised and unpolarised light. Hence when incident light is polarised, the contribution from diffuse polarisation is smaller and the polarisation degree captures coarser surface features. It is worth noting that the degree of polarisation is higher in polarised light and near occluding boundaries in unpolarised light. It is also higher for specular reflection in polarised light. The polarisation phase also captures more fine surface details in unpolarised than in polarised light.

The polarisation image is used to calculate the spherical harmonic coefficients using the relations derived in Section 4.3.2. Figures 5.2 and 5.3 show the first few spherical harmonic coefficients for plastic and natural leaves respectively. However the information

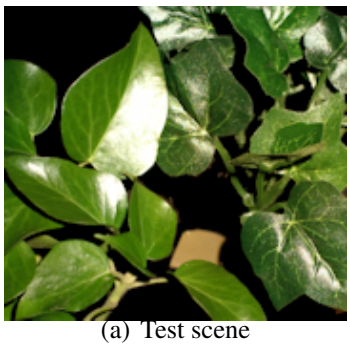

(a) Test scene

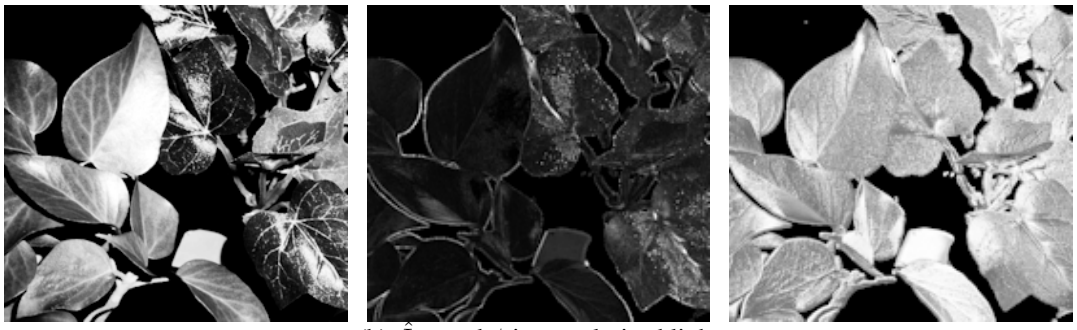

(b) $\hat{I}$, ρ and ϕ in unpolarised light

(c) $\hat{I}$, ρ and ϕ in polarised light

Figure 5.1: Components of the polarisation image computed for the scene in (a) are shown in (b) and (c).

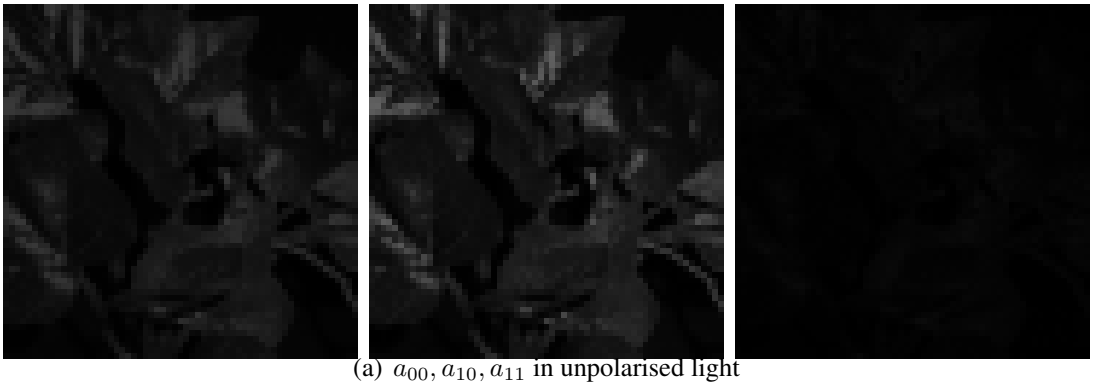

(a) a_{00}, a_{10}, a_{11} in unpolarised light

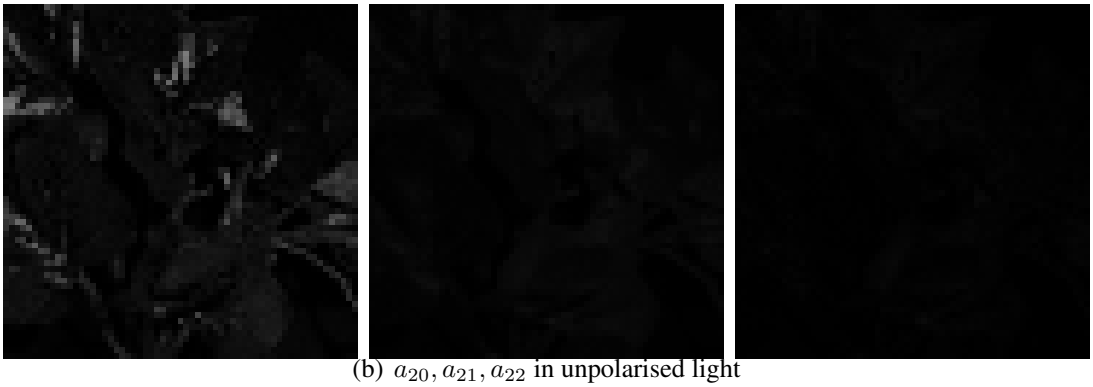

(b) a_{20}, a_{21}, a_{22} in unpolarised light

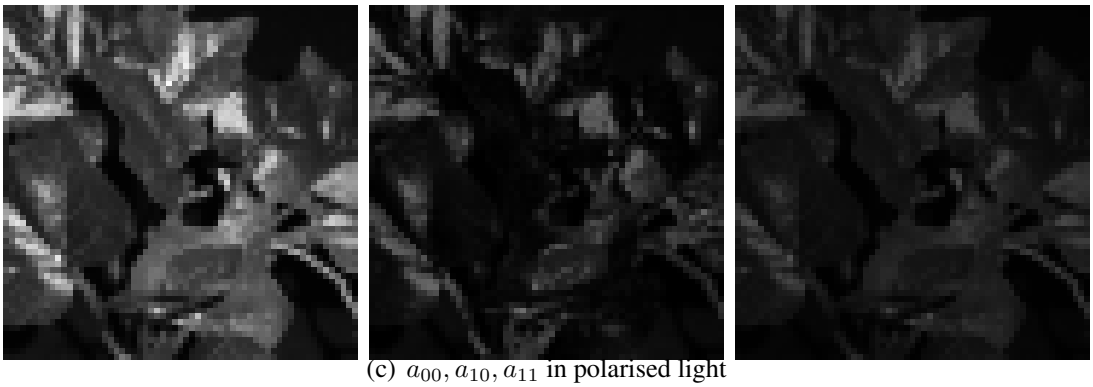

(c) a_{00}, a_{10}, a_{11} in polarised light

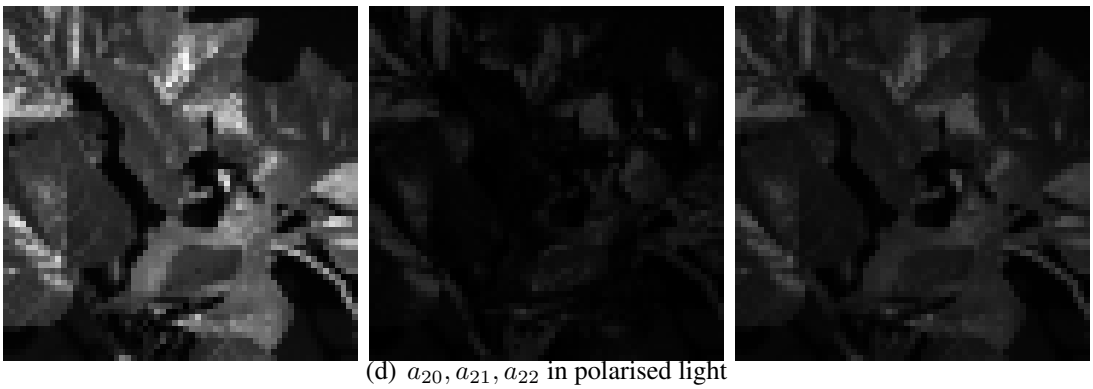

(d) a_{20}, a_{21}, a_{22} in polarised light

Figure 5.2: Spherical harmonic coefficients for plastic leaves, corresponding to positive values of m for l from 0 to 2

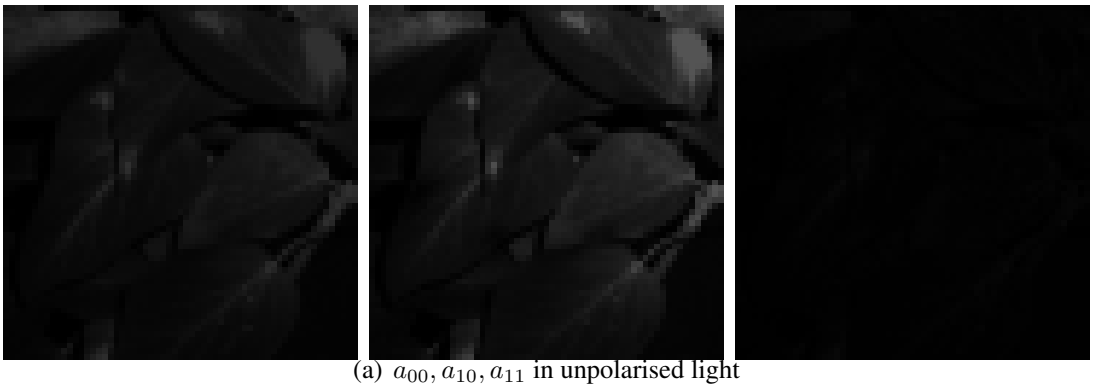

(a) a_{00}, a_{10}, a_{11} in unpolarised light

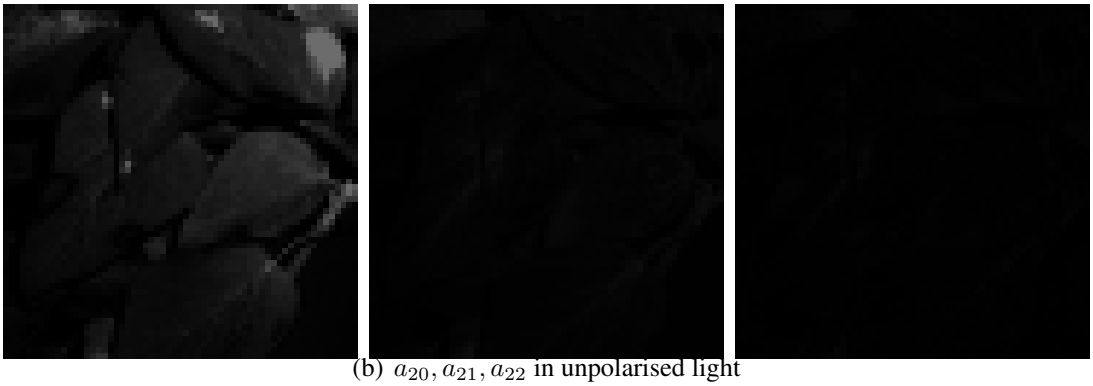

(b) a_{20}, a_{21}, a_{22} in unpolarised light

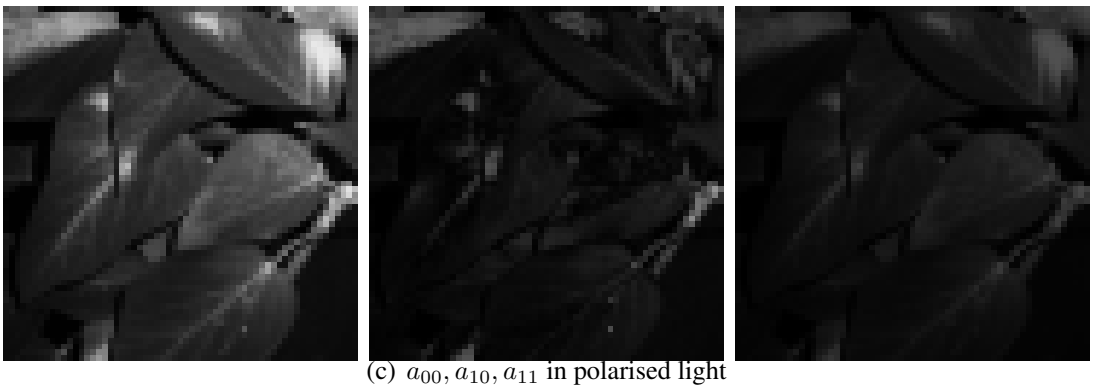

(c) a_{00}, a_{10}, a_{11} in polarised light

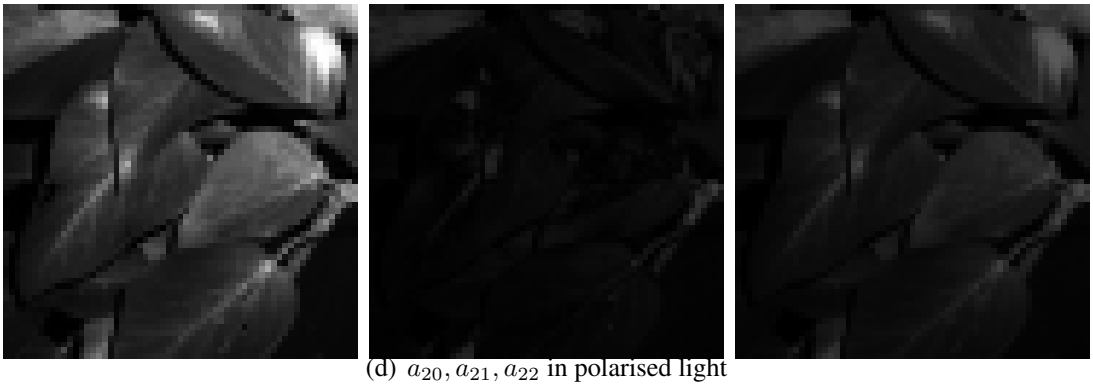

(d) a_{20}, a_{21}, a_{22} in polarised light

Figure 5.3: Spherical harmonic coefficients for plastic leaves, corresponding to positive values of m for l from 0 to 2

contained in the spherical harmonic coefficients is distributed among a large number of data dimensions. One way to compress the representation is by transforming the data in terms of a set of basis functions that express the variation better. The KLT (or PCA) is a powerful statistical tool for such data analysis and compression. Figure 5.4 shows the feature vectors derived by applying the PCA transform to Σ_A , the covariance matrix of spherical harmonic coefficients. The first three principal components for the test image account for 95% of the total variance. For a spherical harmonic expansion truncated at $l = 20$ for example the coefficient vector is typically reduced from 441 to 4 or 5 dimensions.

Spherical harmonic coefficients are calculated up to degree $l = 12$ for results in Figure 5.5, resulting in a coefficient vector with 169 dimensions. The first 6 of these dimensions for positive m are shown in Figures 5.2 and 5.3. However, after applying the PCA transform to coefficient vectors, 95 % of the variance is accounted for by the first 5 dimensions, 3 of which are shown in Figure 5.4(a) - (c). Thus almost all information in the feature vector is compressed from the original 169 dimensions into 5 dimensions.

The results of segmenting the scene in Figure 5.1(a) using normalized graph cuts algorithm from [41] are shown in Figure 5.5. These results were obtained by truncating the spherical harmonic expansion at order 8 on the 660×720 image. The affinity matrix was computed using the Mahalanobis distances between the feature vectors in blocks of 10×10 pixels. A small minority of the image blocks are classified incorrectly. Segmentation is better in unpolarised light even in the presence of specularities due to stronger spontaneous polarisation degree and therefore stronger discrimination in features. Specularities in the image cause some difficulty in correct segmentation when using polarised light.

The results of segmentation for outdoor scenes are shown in Figures 5.6, 5.7 and 5.8. The images to be segmented are shown in (a), and the components of polarisation image for each are shown in (b) to (d). Subfigure (e) shows the segments generated by the normalized cut algorithm. Segmentation results were tested for different values of window size and orders of expansion. The window size that produces optimal segmentation is found to vary from scene to scene. It is noted that for the window size that produces optimal segmentation, spherical harmonic expansion truncated at order 12 gives the desired segments and the results are stable over multiple runs. Adding higher order terms to the feature vector does not improve performance noticeably. This implies that most of the key feature information is contained in low order harmonics, up to order 12. For other window sizes however the normalized cuts seem to produce unpredictable results. This is

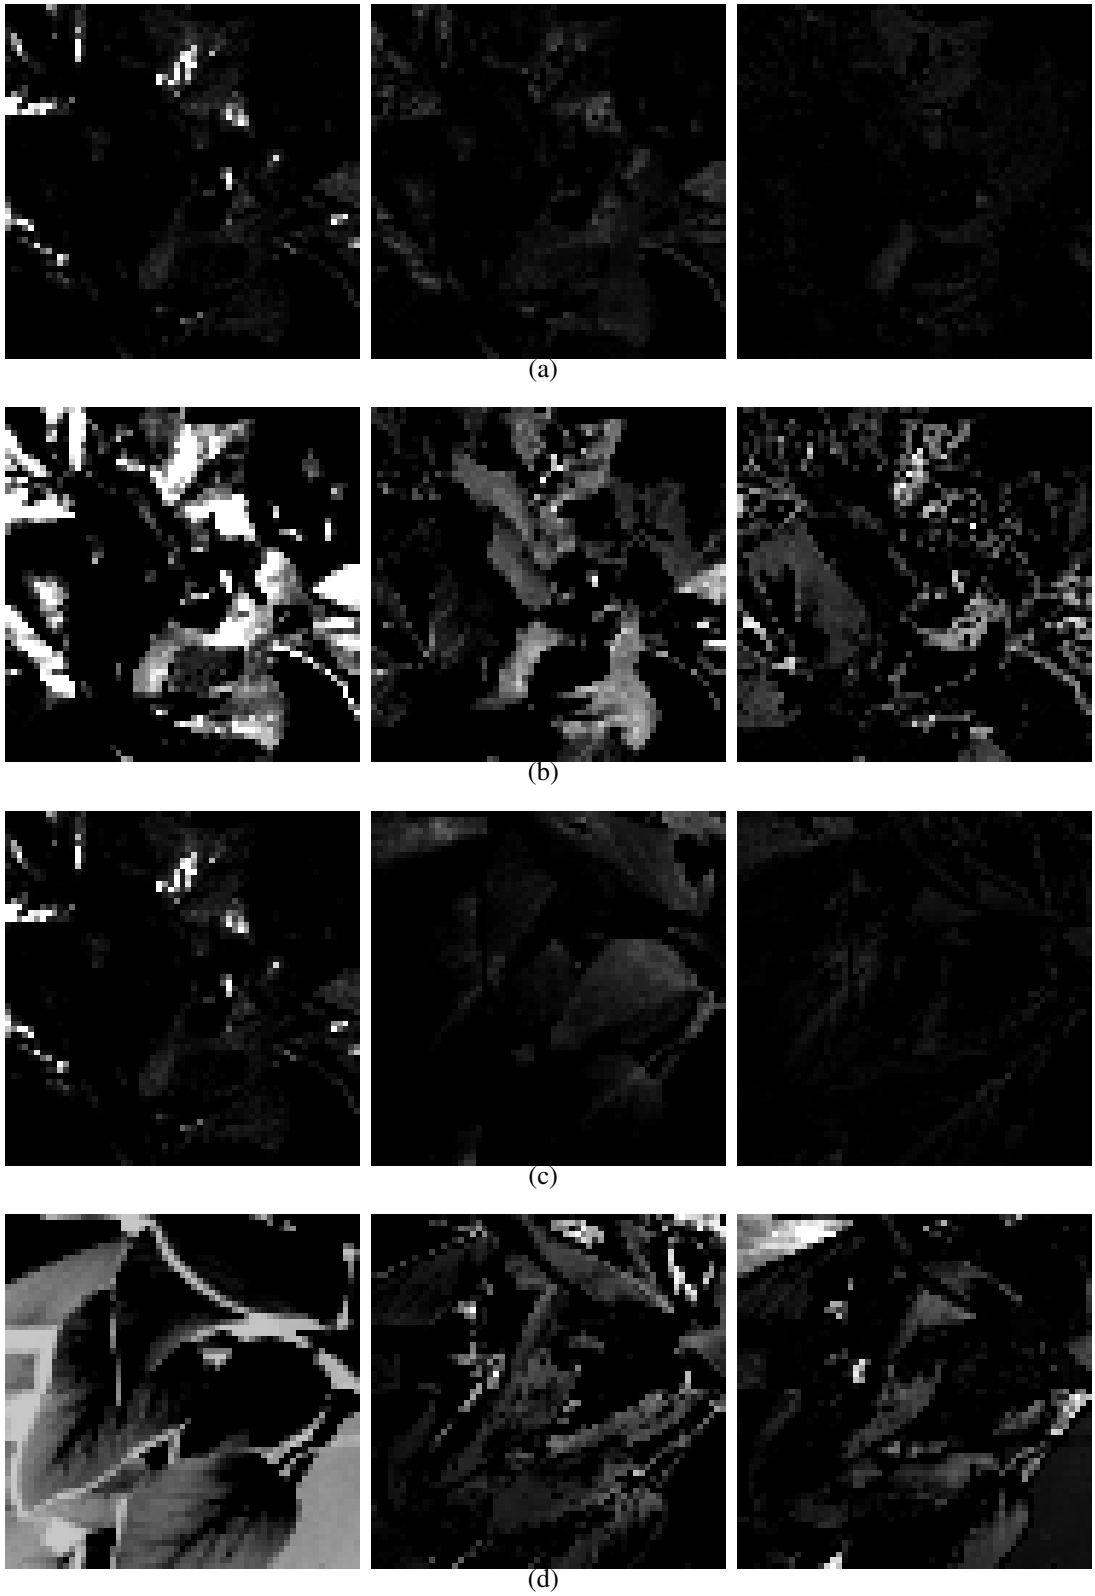

Figure 5.4: Dimensions of the feature vector for plastic leaves in (a) unpolarised and (b) polarised light, and natural leaves in (c) unpolarised and (d) polarised light.

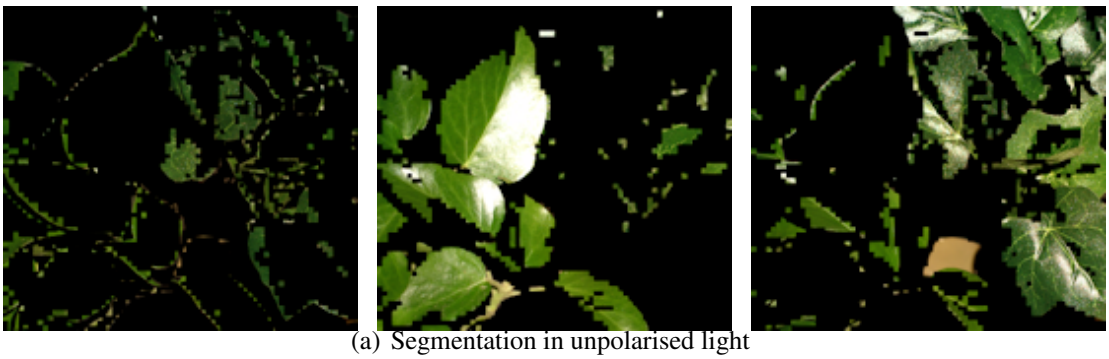

(a) Segmentation in unpolarised light

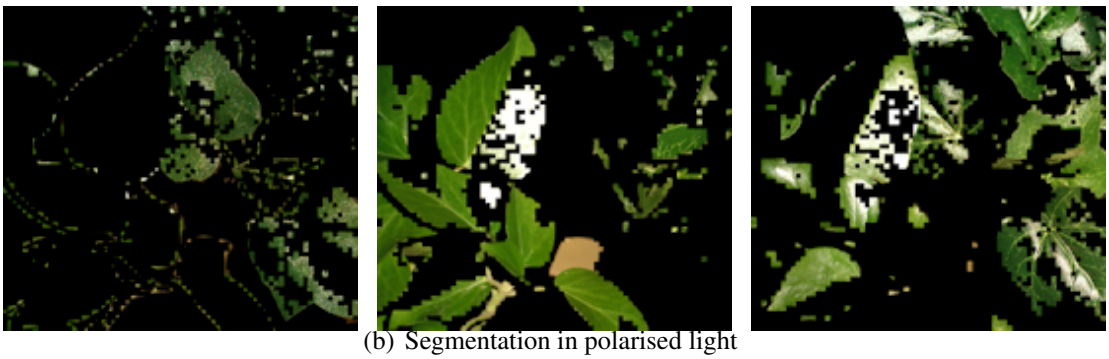

(b) Segmentation in polarised light

Figure 5.5: Segmentation of (a) into (L-R): background, natural leaves and plastic leaves.

an area that requires further investigation to analyze the factors that affect the accuracy of segmentation.

5.3 Application II: Surface Quality Inspection

Layered surfaces are common in nature, and are found in leaves, fruits and skin. This presents another application for the segmentation technique developed, in inspection of the quality of layered surfaces. The angular distribution of light is estimated from the polarisation image and parameterized using spherical harmonics. The feature vectors thus obtained are transformed along the principal axes and used to characterize the reflectance distribution on a pixel-by-pixel basis. This section presents the analysis of images of bruised fruit. The segmentation task is then to separate the fresh and bruised regions in a fruit before such damage appears as a lesion on the skin. The fruits used in the experiment are bruised with a soft hammer and allowed to rot. The hammered areas on the fruit are marked in ink to allow easy identification. Images of the fruit are captured three and four days after the being bruised. The fruits used in the experiment are an apple, pear, orange, two plums and two apricots.

Figure 5.9 shows the arrangement of fruit. There are four fruit subject to decay, namely the pear (first from left), the large apple (second from left), the plum (third from left) and

(a) Input image

(b) Polarisation image: $\hat{I}$

(c) Polarisation image: ρ

(d) Polarisation image: ϕ

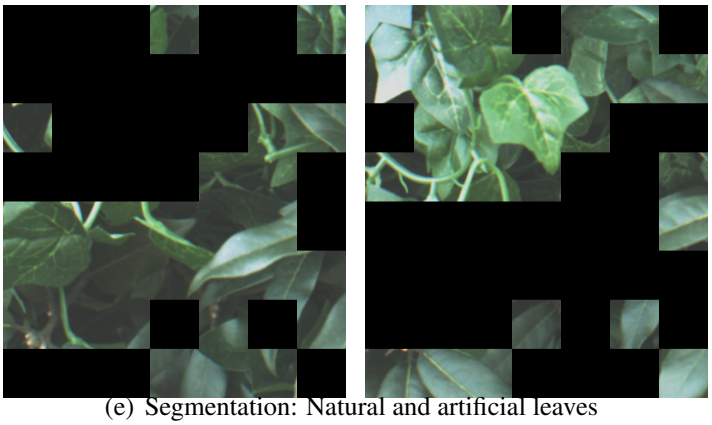

(e) Segmentation: Natural and artificial leaves

Figure 5.6: Segmentation for artificial leaves and foliage in diffuse sunlight with $w = 24$, $l = 30$

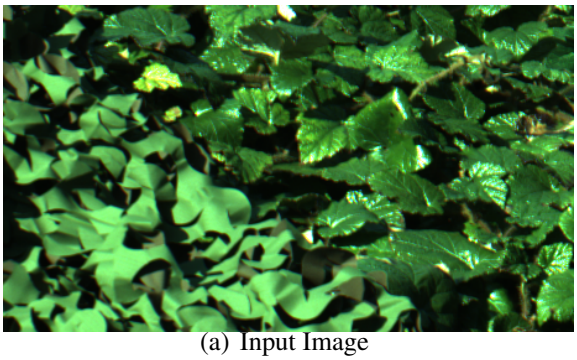

(a) Input Image

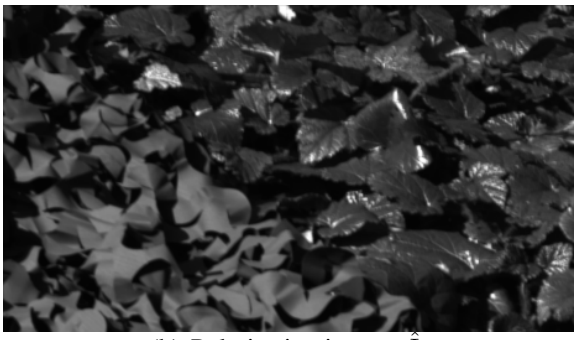

(b) Polarisation image: $\hat{I}$

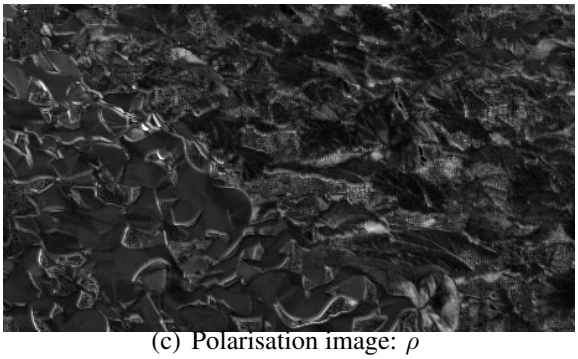

(c) Polarisation image: ρ

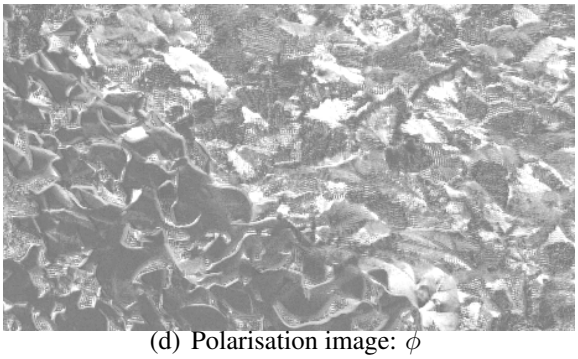

(d) Polarisation image: ϕ

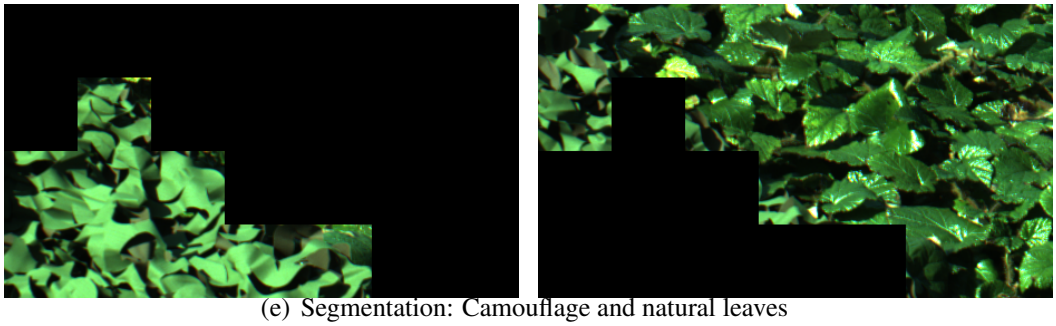

(e) Segmentation: Camouflage and natural leaves

Figure 5.7: Segmentation for camouflage and foliage in direct sunlight with $w = 66$, $l = 12$

(a) Input Image

(b) Polarisation image: $\hat{I}$

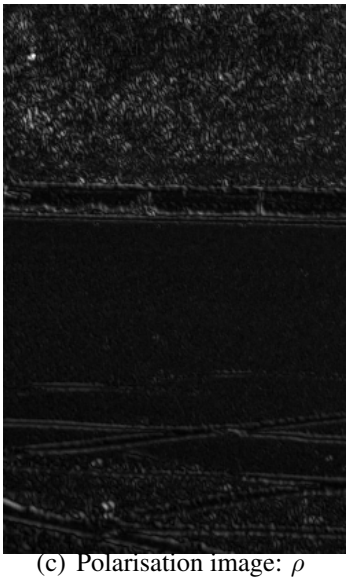

(c) Polarisation image: ρ

(d) Polarisation image: ϕ

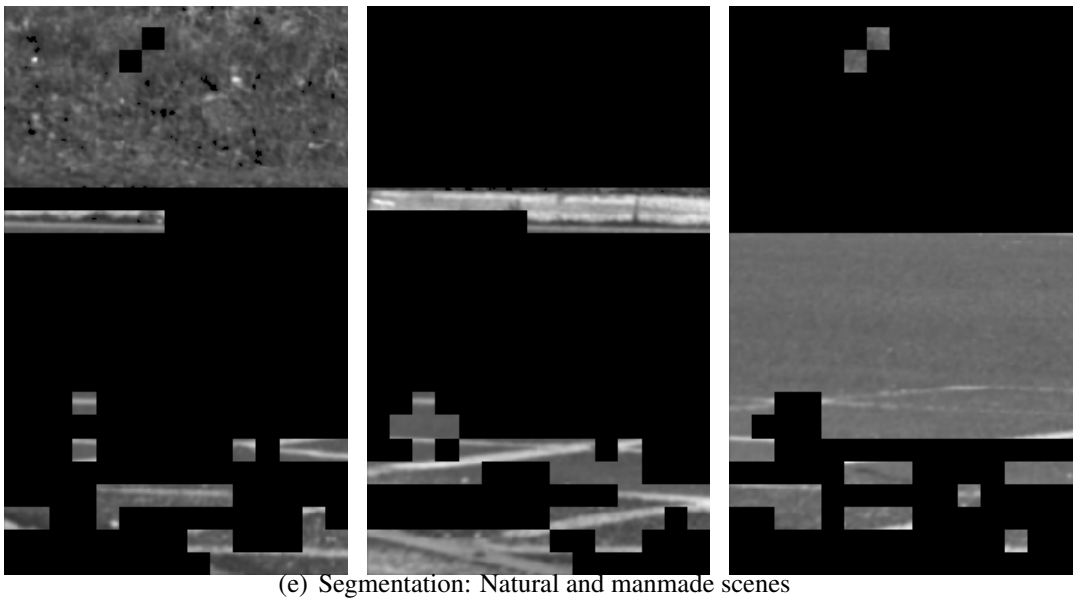

(e) Segmentation: Natural and manmade scenes

Figure 5.8: Segments for man made and natural objects in direct sunlight with $w = 16$, $l = 30$

apricot (fifth from left). The regions bruised on the pear, apple and apricot have been circled with a marker pen. Results of a study on two image sections are presented here, each of size 240×240 pixels. The first sub-image shows the affected regions of the apple and pear, and the second shows the affected area of the plum. The degree and phase of

(a) Polarisation image: $\hat{I}$

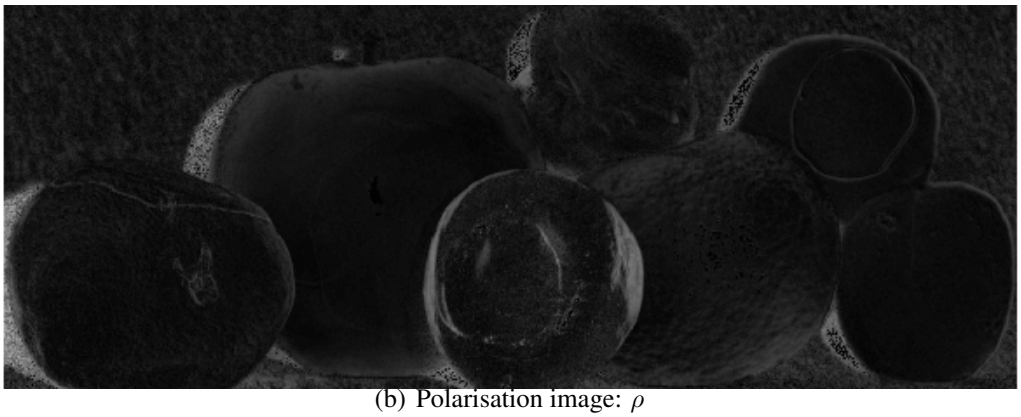

(b) Polarisation image: ρ

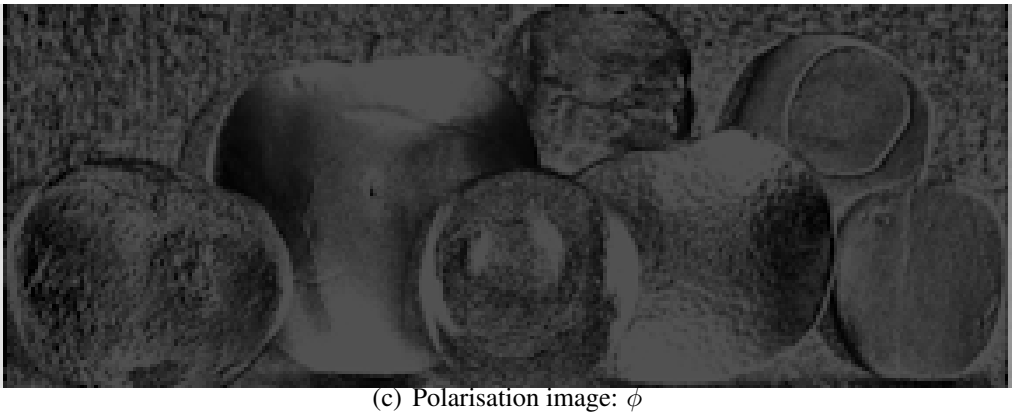

(c) Polarisation image: ϕ

Figure 5.9: Polarisation image components for scene with fruit

diffuse polarisation for the full image are shown in Figure 5.9 (b) and (c) for the scene, and the polarisation phase is shown in Figure 5.10 for different stages of rotting. There are a number of features to note from the polarisation data. First, the degree and phase of polarisation reveal the boundaries between the undamaged and bruised surface regions in the scene. Second, fresh fruit exhibits a higher shift in polarisation phase than rotting

fruit.

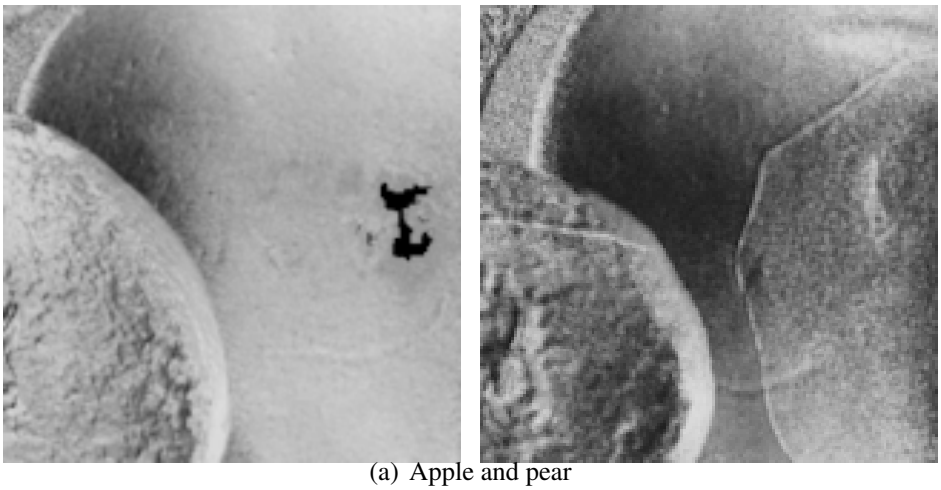

(a) Apple and pear

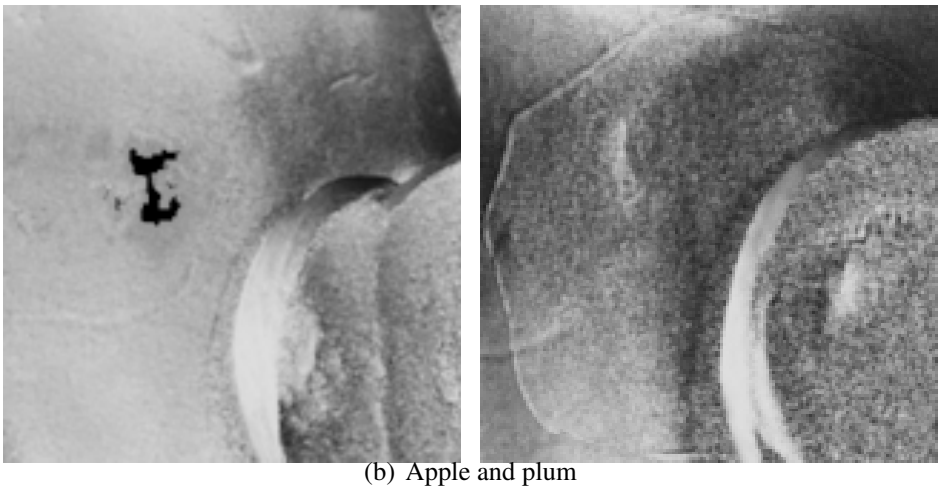

(b) Apple and plum

Figure 5.10: Polarisation phase on day 3 and day 4 showing changes in increasing stages of rotting

Figures 5.11 and 5.12 show the spherical harmonic coefficients for the two image sections. The main feature to note is that the different objects define different regions in the data. Also, the co-efficient variation within an object is greatest in the damaged areas. Here again, the PCA transform is applied on the coefficient vector covariance matrix Σ_A. The results for the two -image sections are shown in Figures 5.13 and 5.14. The figures show the first four principal components which account for 95% of the data variance. These are used to compute a block-by-block feature vector to describe the fruit surface. The features highlight the difference in the healthy and rotten fruit tissues as seen in Figure 5.13. The variation amongst the different fruits is also emphasized as shown in Figure 5.14, while the feature vectors remain unaffected by shape variation.

The results of segmenting the scene using normalized graph cuts algorithm from [41] are shown in Figure 5.15 and Figure 5.16. These results were obtained with 289 features from spherical harmonic coefficients up to degree 16 on a 240×240 image. The affinity matrix was computed by using the Mahalanobis distances between the feature vectors in

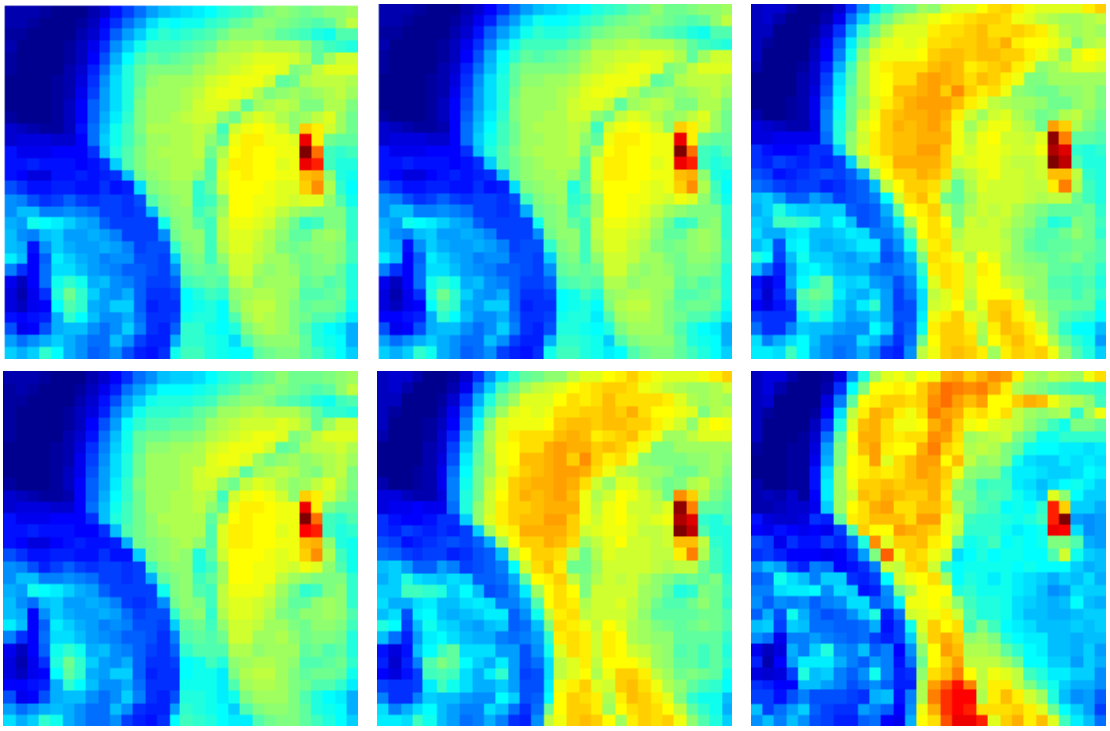

Figure 5.11: Spherical harmonic coefficients for pear and apple for l=0,1 and 2 and positive values of m.

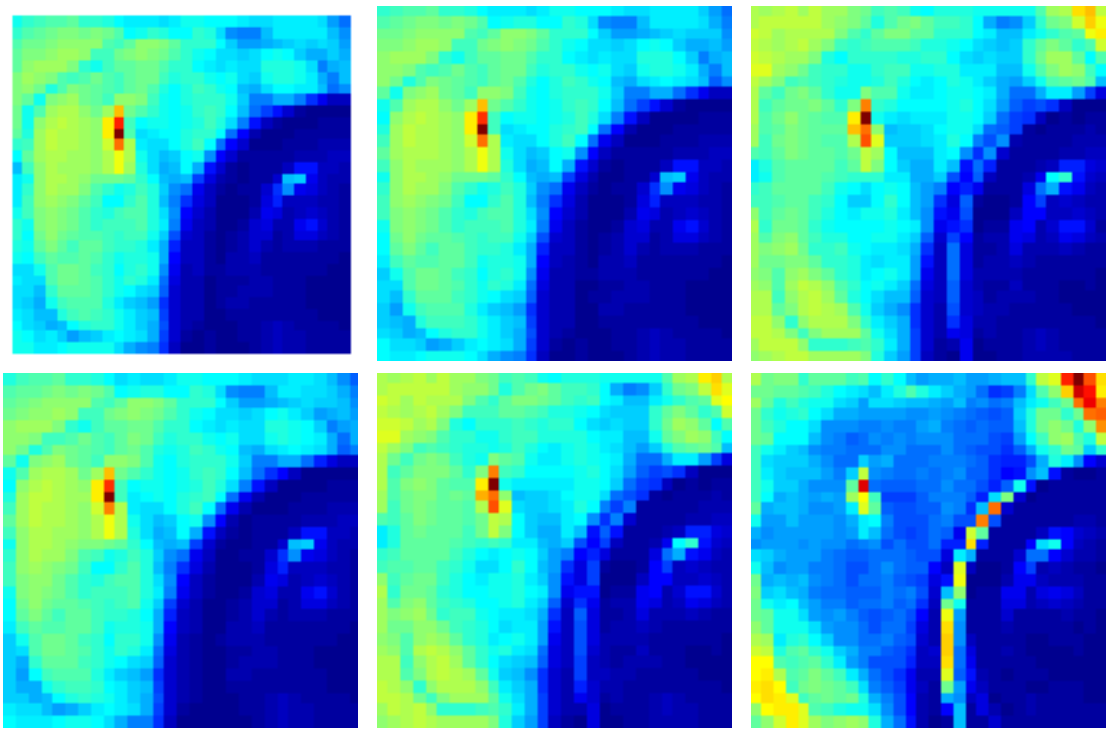

Figure 5.12: Spherical harmonic coefficients for apple and plum for l=0,1 and 2 and positive values of m.

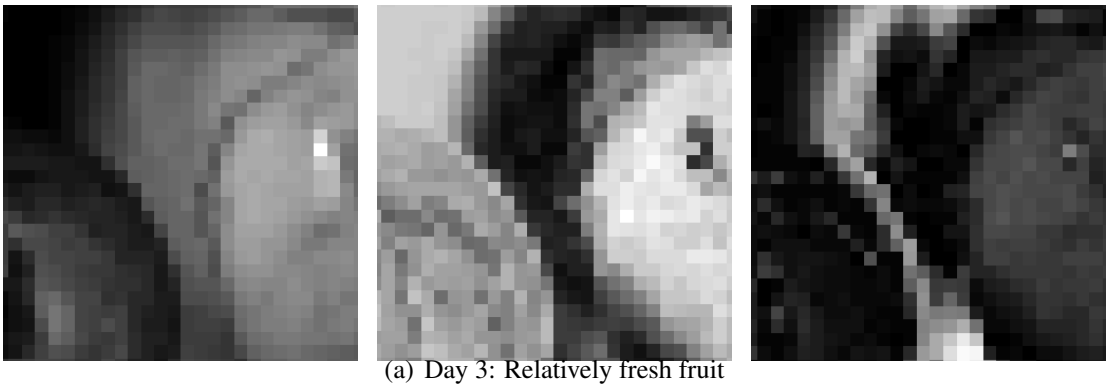

(a) Day 3: Relatively fresh fruit

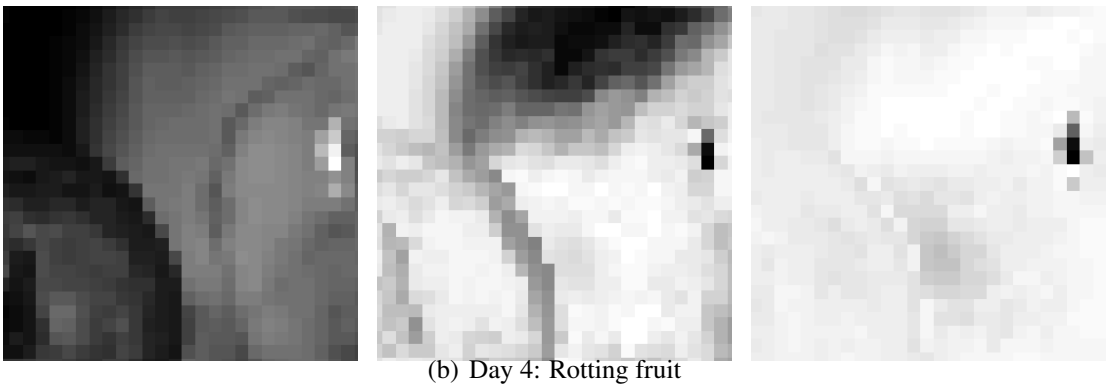

(b) Day 4: Rotting fruit

Figure 5.13: Feature vector for pear and apple

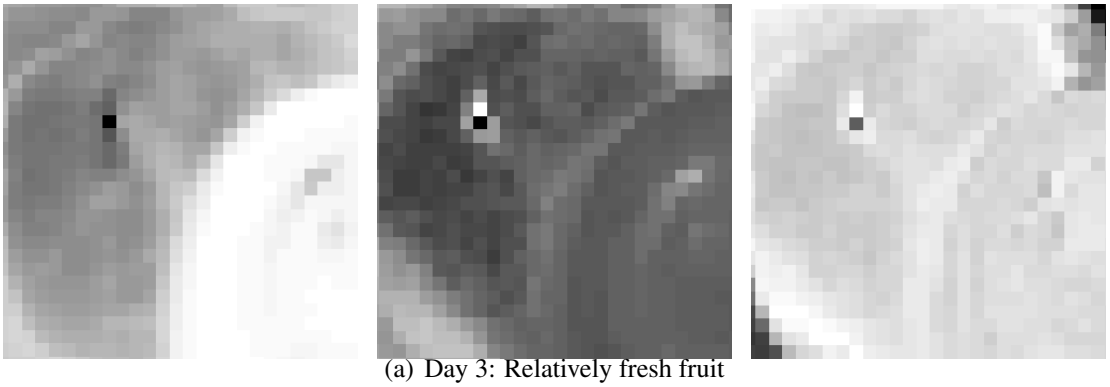

(a) Day 3: Relatively fresh fruit

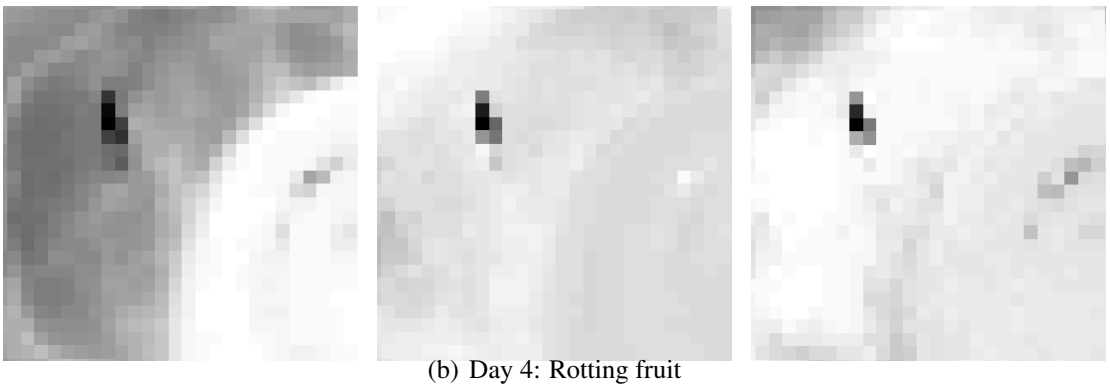

(b) Day 4: Rotting fruit

Figure 5.14: Feature vector for apple and plum

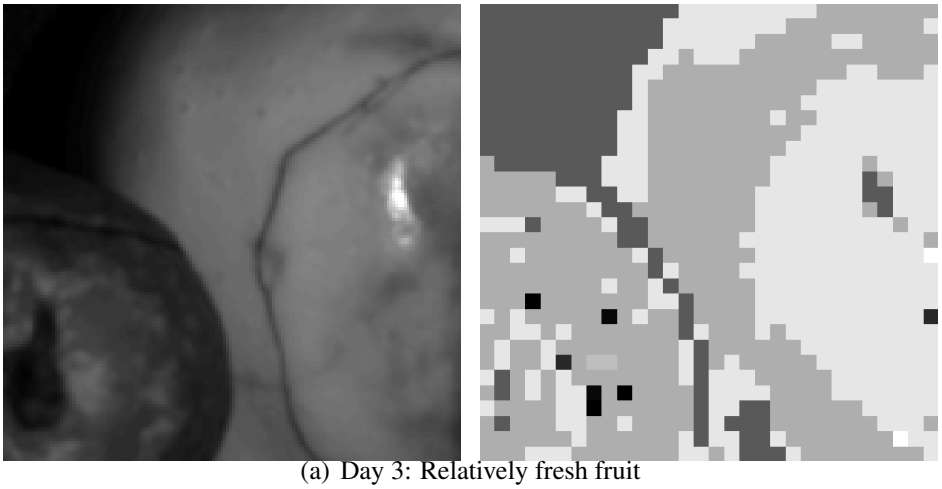

(a) Day 3: Relatively fresh fruit

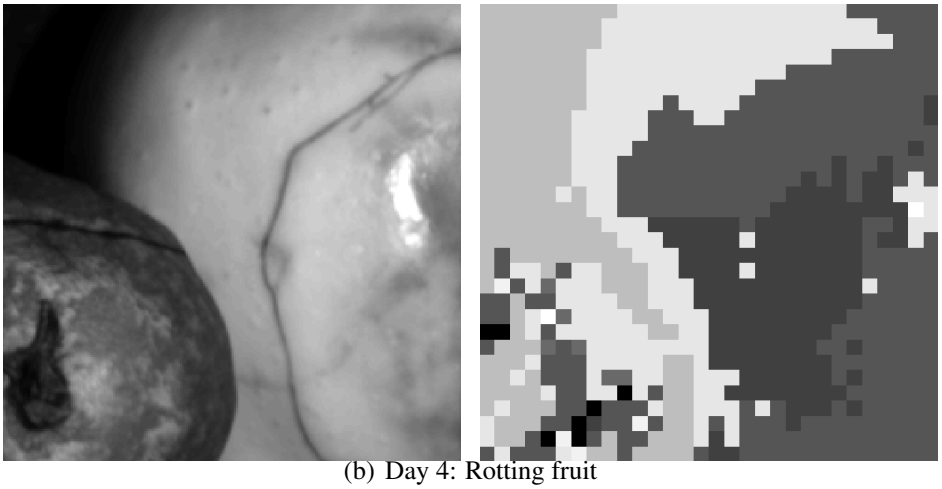

(b) Day 4: Rotting fruit

Figure 5.15: Segmentation results for pear and apple, segments coded by gray level

blocks of 8×8 pixels.

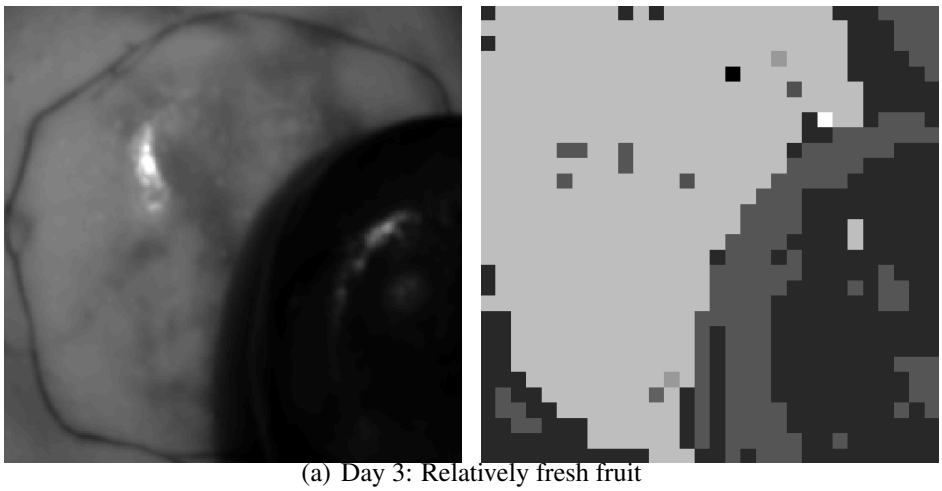

(a) Day 3: Relatively fresh fruit

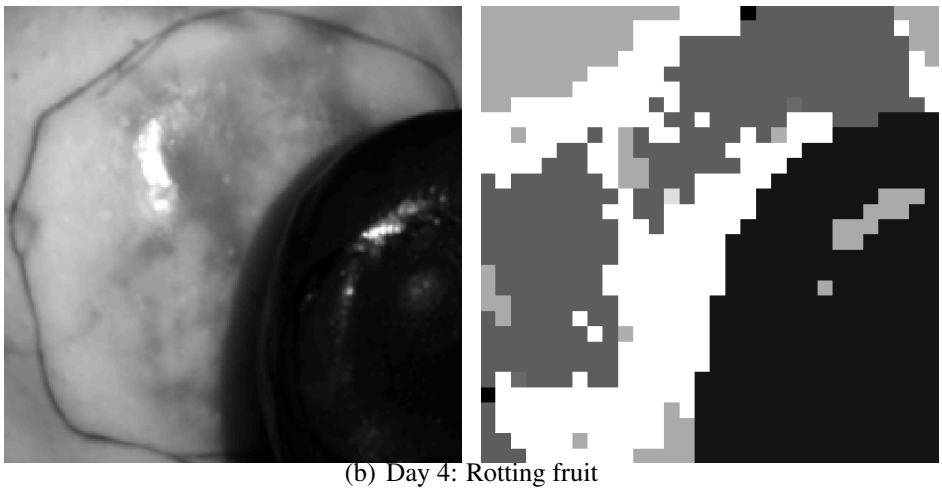

(b) Day 4: Rotting fruit

Figure 5.16: Segmentation results for apple and plum, segments coded by gray level

5.4 Evaluation of Results

In Figure 5.6, the image to be segmented is 462×264 pixels in size and is divided into blocks of size 24×24 for spherical harmonic coefficient calculation. The harmonic expansion is truncated at order 30. The image is divided into 7×8 blocks. Out of 56 blocks, For the scene in Figure 5.7(a), a block size of 66×66 for feature generation produced the best segmentation performance for the the image sized 370×420 pixels. With this window size, the image is divided into 7×4 blocks of pixel regions, generating a $7 \times 4 \times 4$ feature matrix after PCA. Out of 28 blocks, 2 blocks are classified incorrectly, and 3 blocks contain sections of both materials. Assuming a worst case scenario that all blocks containing more than one material are incorrectly classified gives an accuracy of 82% while a more optimistic estimation gives an accuracy of 93%. Images of the mixed scene in Figure 5.8 were captured in direct sunlight in late evening time. Sunlight is highly po-

larised early mornings and late evenings because of the large distances it travels through the atmosphere. Thus this scene is captured in highly polarised incident light. The image is 250×400 pixels in size and the coefficients were calculated over 16×16 pixel blocks and truncated at order 30. Thus the image is divided into 15×25 feature blocks. The segments in the image represent grass, a stone pavement and the tar road. Some grass growing around the corners of edges of the pavement is also classified correctly. A part of the pavement that is also classified with grass in the first segment is due to limitations of the block-wise feature calculation. In this image, out of 375 blocks 34 are classified incorrectly, giving an estimated accuracy of 91%

5.5 Image Reconstruction from Harmonic Coefficients

The previous section showed segmentation results for images using spherical harmonic expansions of different orders. In practice, it is common to truncate the harmonic expansion and ignore higher order terms. The truncation however introduces an inevitable error in the estimated function. The accuracy of the estimated intensity function can be checked by reconstructing the intensity at each pixel as a superposition of the weighted spherical harmonic functions $Y_l^m(\theta, \phi)$. The expansion, given in (4.20), is repeated here for reference:

$$\hat{I}(\theta, \phi) = \sum_{l=1}^{\infty} \sum_{m=-l}^{l} a_{l,m} Y_l^m(\theta, \phi), a \in \mathbb{R}$$

The reconstructed intensity $\hat{I}$ is an approximation to the low pass filtered original image, as the high frequency information contained in higher order terms is removed by the truncation. The histograms in Figure 5.17 show a comparison of the original histogram (a) with reconstructions (b) - (d) using coefficients of degree 20, 12 and 4, for a patch of the scene in Figure 5.1(a).

The reconstruction error shows a drop with increase in degree of expansion, as expected. The resolution of reconstruction however is limited by the choice of window size in coefficient calculation. The reconstructed images exhibit a smoothing of the intensities attributed to the truncation of spherical harmonic expansion and the windowing in calculation of coefficients. The mean error The following table summarizes the reconstruction errors in different images with change in window size w and expansion order l.

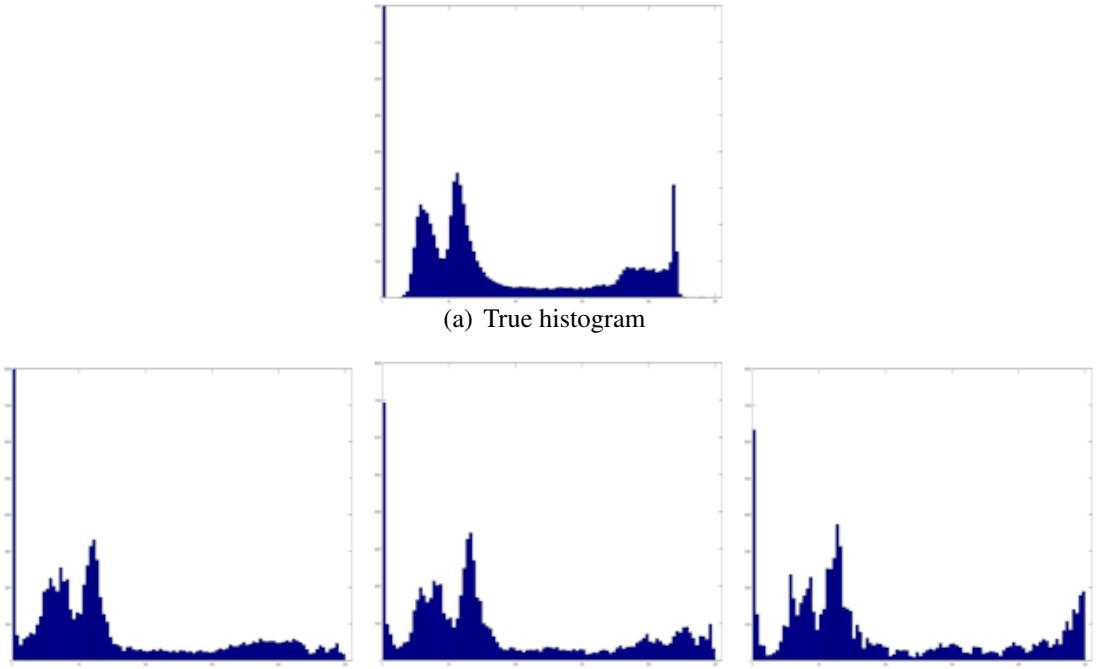

(a) True histogram

(b) Reconstruction at l=20 (c) Reconstruction at l=12 (d) Reconstruction at l=4

Figure 5.17: Histograms: (a) shows an image histogram and (b) to (d) show histograms of reconstructed pixel intensities with expansions up to $l = 20, 12$ and 4.

Table 5.1: Reconstruction error in intensity recalculated from spherical harmonic functions

No.	Image size	w	l	Mean error
1	320x320	4	20	0.030091
2	320x320	4	12	0.030091
3	200x200	4	30	0.033601
4	200x200	4	12	0.033601
5	200x200	12	30	0.086619
6	200x200	12	12	0.086619
7	120x120	12	30	0.19924

5.6 Conclusions

The experimental results presented in this chapter show successful application of polarisation theory to the problem of image segmentation based on surface reflectance properties. The results presented show how polarisation information can be used to detect damaged surface regions for soft fruit and different materials. The technique segments out sections of the image that contain objects of different material, as evident from the segmentation of plastic leaves and foliage, camouflage and foliage. It also segments sections of same material that differ in quality, as demonstrated with the segmentation of the rotting fruit. The segmentation seems to be more effective in unpolarised light than in polarised light, attributed to the fact that the spontaneous polarisation of remitted light is more easily measured in initially unpolarised light. The suitability of the technique is demonstrated in controlled laboratory conditions and uncontrolled outdoor conditions, yielding successful segmentation in both scenarios. The histogram of the image reconstructed by superimposing the weighted spherical harmonics bears increasing similarity to the original image as the degree of expansion is increased. A satisfactory reconstruction is achieved at $l > 16$ although there is some inevitable blurring due to effects of windowing and truncation.

Chapter 6

Conclusions and Further Research

6.1 Conclusions

There has been substantial work over the last two decades in polarisation vision with applications in specularity removal, shape from polarisation and material discrimination. Wolff suggested in [21] that polarisation degree can be used to differentiate between metals and dielectrics. This thesis concentrates on the narrower problem of polarisation analysis of dielectric material surfaces. The contribution of this work is in suggesting a technique to segment dielectric materials using polarisation information, in cases that are otherwise difficult to segment using colour and intensity based techniques. The novel aspects of this research are:

- Use of spherical harmonics to represent polarisation information contained in intensity distributions.
- Using spherical harmonic coefficients as a feature vector to segment an image based on material properties.

The method characterizes differences in refractive index within dielectric materials using the Fresnel theory and uses spherical harmonics to represent the spatial intensity variation. The harmonic coefficients are then used as features to segment the image based on material properties. Two application areas are selected to test the algorithm and demonstrate its effectiveness. Results from the experiments show promising segmentation accuracy in both artificial and natural light. The first application is in target detection, where camouflage net and plastic leaves are placed with natural leaves and an attempt is made to separate them using the algorithm. The results in Section 5.2 demonstrate that the method

is effective in polarised and unpolarised artificial light, as well as in direct and diffuse sunlight. The thesis also demonstrates the applicability of the algorithm to quality inspection of degrading fruits in Section 5.3. Various fruits are bruised using a soft hammer and allowed to rot. Periodic images are taken and the proposed technique is used to separate out the damaged and degenerating areas of the fruit. Segmentation separates out areas that appear at a later stage as skin lesions and visible rot.

The segmentation results are slightly more accurate under controlled lighting conditions than in natural sunlight. This is attributed to lack of control over experimental parameters in outdoor settings. For example fluctuations in the intensity of sunlight due to atmospheric variations or passing clouds can cause significant jumps or drops in the recorded pixel intensities for a given set up. Another source of error is mis-registration of images taken at successive polariser orientations, due to movement of leaves and branches in outdoor scenes. These drawbacks can easily be overcome using more sophisticated equipment like liquid crystal based polarisation cameras or PLZT (Polarised Lead Zirconium Titanate) [42]. Moreover, despite these inevitable variations, the overall segmentation is remarkably accurate.

It is observed that the polarisation image for experiments performed in unpolarised incident light contain more details than those performed in polarised light. This is due to the fact that the analysis relies on information contained in diffuse polarisation. Diffuse polarisation arises when light enters a material, undergoes scattering and is remitted through refraction at the material surface. The light acquires a spontaneous polarisation due to refraction at the surface, which is measured for analysis. Since this is a spontaneous phenomenon, the degree of polarisation is very weak (usually < 0.4). This range of polarisation is easier to detect in initially unpolarised light than in light that is already strongly polarised.

It is also observed that the segmentation accuracy is somewhat dependent on window size. It appears that spherical harmonic decomposition may be sensitive to window size and is more accurate when the window size matches some measure of the texture elements in the image. However more time and effort needs to be allocated to further investigate the factors that affect the choice of cut and to identify the dependency of segmentation results on window size and choice of cut off frequency.

6.2 Further Research

The work presented in this thesis is meant to serve as an exploration into the use of polarisation for material discrimination and camouflage detection. The thesis inevitably raises questions that require further attention to arrive at satisfactory answers. Some of these are listed here as suggestions for further research.

- Determination of zenith and refractive index: The results presented in this thesis have been derived with the assumption that refractive index of materials in the study varies in the range of 1.3 to 1.5, which is reasonable for dielectric materials. Since an assumed refractive index of 1.45 is used for calculations, there is a small error in the estimated zenith angle. This can be corrected using a multi-view set up. Changing light source direction allows the capture of two views of the object. The polarisation images calculated for each light source can then be used to write equation (3.13) as a system of simultaneous non-linear equations which are then solved for the exact values of n and θ.

- Shape recovery for features of non-planar objects: The results in Section 5.2 show segmentation for bruised fruits. Although the method is effective, the application does not strictly satisfy the initial planar assumption for estimation of zenith angle from refractive index using Fresnel theory. A more accurate solution would be to use a shape recovery method, for example shape from shading, to estimate the surface normals and use these estimates in the spherical harmonic expansion. This solution unfortunately is beyond the scope of the current thesis.

- Quantitative analysis of refractive index variations: The results developed and presented in this thesis are based on the hypothesis that slowly varying refractive index and planar samples allow the determination of the zenith and azimuth angles of remitted light. Experimental observations support the hypothesis that changes in material characterize the intensity distributions of remitted light and this can be used for material based image segmentation. However a differential analysis of the relation between zenith angle and the degrees of diffuse and specular polarisation could yield a richer theoretical and mathematical explanation for the phenomenon.

6.3 Merits and Drawbacks

The merits of the technique suggested in this thesis lie in its accuracy for correct selection of window size and expansion order. Although the data collection method used in the experiments is cumbersome and time consuming, there are polarisation cameras available in the market that make data collection much swifter and easier. One major drawback of the suggested technique is the use of an eigenvector-based image segmentation method which makes it too slow for practical and real-time applications. The error in the approximations introduced by Shi and Malik for computation of normalized cuts is still not full understood and the applicability is still limited to small image patches with computation time typically in minutes. An alternative would be to attempt segmentation using a more efficient segmentation like [43]. In summary, this thesis presents a new technique for segmentation of camouflaged textures in images, with novel contributions to the analysis and detection of the textures.

List of Acronyms

BRDF Bidirectional Reflectance Distribution Function

BTF Bidirectional Texture Functions

E-field Electric Field

FFT Fast Fourier Transform

KLT KarhunenLoève transform

MRF Markov Random Field

N-cut Normalized graph cut

PCA Principal Component Analysis

PLZT Polarised Lead Zirconium Titanate

Bibliography

[1] G. Atkinson and E. R. Hancock, "Recovery of surface orientation from diffuse polarization," *IEEE Transactions on Image Processing*, vol. 15, no. 6, pp. 1653–1664, 2006.

[2] H. Cheng, X. Jiang, Y. Sun, and J. Wang, "Color image segmentation: Advances and prospects," *Pattern Recognition*, vol. 34, no. 12, pp. 2259 – 2281, 2001.

[3] N. R. Pal and S. K. Pal, "A review on image segmentation techniques," *Pattern Recognition*, vol. 26, no. 9, pp. 1277 – 1294, 1993.

[4] M. J. Chantler, "Why illuminant direction is fundamental to texture analysis," *IEE Proceedings: Vision, Image and Signal Processing*, vol. 142, no. 4, pp. 199–206, 1995.

[5] D. A. Forsyth and J. Ponce, *Computer Vision - A Modern Approach*. Prentice Hall, 2003.

[6] M. Saito, Y. Sato, K. Ikeuchi, and H. Kashiwagi, "Measurement of surface orientations of transparent objects using polarization in highlight," *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, vol. 1, pp. 381–386, 1999.

[7] S. Rahmann and N. Canterakis, "Reconstruction of specular surfaces using polarization imaging," in *Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, vol. 1, pp. 149–155, 2001.

[8] D. Miyazaki, M. Kagesawa, and K. Ikeuchi, "Transparent surface modeling from a pair of polarization images," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 26, no. 1, pp. 73–82, 2004.

[9] B. Julesz, "Textons, the elements of texture perception, and their interactions," *Nature*, vol. 290, pp. 91–97.

[10] J. Malik, S. Belongie, J. Shi, and T. Leung, "Textons, contours and regions: Cue integration in image segmentation," vol. 2, pp. 918–925, 1999.

[11] T. Leung and J. Malik, "Representing and recognizing the visual appearance of materials using three-dimensional textons," *International Journal of Computer Vision*, vol. 43, no. 1, pp. 29–44, 2001.

[12] K. J. Dana, B. van Ginneken, S. K. Nayar, and J. J. Koenderink, "Reflectance and texture of real-world surfaces," *ACM Transactions on Graphics*, vol. 18, no. 1, pp. 1–34, 1999.

[13] M. Chantler, M. Petrou, A. Penirsche, M. Schmidt, and G. McGunnigle, "Classifying surface texture while simultaneously estimating illumination direction," *International Journal of Computer Vision*, vol. 62, no. 1, pp. 83–96, 2005.

[14] Y. Xu, H. Ji, and C. Fermüller, "Viewpoint invariant texture description using fractal analysis," *International Journal of Computer Vision*, vol. 83, no. 1, pp. 85–100, 2009.

[15] W. T. Freeman and P. A. Viola, "Bayesian model of surface perception," in *Proceedings of Conference on Advances in Neural Information Processing Systems*, pp. 787–793, 1998.

[16] M. Bell and W. T. Freeman, "Learning local evidence for shading and reflection," in *Proceedings of IEEE International Conference on Computer Vision*, vol. 1, pp. 670–677, 2001.

[17] M. F. Tappen, W. T. Freeman, and E. H. Adelson, "Recovering intrinsic images from a single image," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 27, no. 9, pp. 1459–1472, 2005.

[18] S. Padilla, O. Drbohlav, P. R. Green, A. Spence, and M. J. Chantler, "Perceived roughness of $1/f_\beta$ noise surfaces," *Vision Research*, vol. 48, no. 17, pp. 1791 – 1797, 2008.

[19] A. Clarke, P. Green, M. Chantler, and K. Emrith, "Visual search for a target against a $1/f_\beta$ continuous textured background," *Vision Research*, vol. 48, no. 21, pp. 2193–2203, 2008.

[20] L. Itti and C. Koch, "Computational modelling of visual attention," *Nature Review Neuroscience*, vol. 2, no. 3, pp. 194–203, 2001.

[21] L. B. Wolff, "Polarization-based material classification from specular reflection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 12, no. 11, pp. 1059–1071, 1990.

[22] H. Chen and L. B. Wolff, "Polarization phase-based method for material classification in computer vision," *International Journal of Computer Vision*, vol. 28, no. 1, pp. 73–83, 1998.

[23] G. Atkinson and E. R. Hancock, "Two-dimensional brdf estimation from polarisation," *Computer Vision and Image Understanding*, vol. 111, no. 2, pp. 126–141, 2008.

[24] G. Atkinson and E. R. Hancock, "Polarization-based surface reconstruction via patch matching," pp. 495–502, 2006.

[25] E. Hecht, *Optics*. Addison-Wesley, 4th ed., 2002.

[26] M. Born and E. Wolf, *Principles of Optics*. Cambridge university Press, 7th (expanded) ed., 1999.

[27] L. B. Wolff and T. E. Boult, "Constraining object features using a polarization reflectance model," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 13, no. 7, pp. 635–657, 1991.

[28] G. Atkinson and E. R. Hancock, "Shape estimation using polarization and shading from two views," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 29, no. 11, pp. 2001–2017, 2007.

[29] G. Sun, E. Onoichenco, Y. Fu, Y. Liu, R. Amell, C. McCandless, R. Reddy, G. Kumar, and M. Guest, "High-throughput polarization imaging for defocus and dose inspection for production wafers," vol. 6518, 2007.

[30] F. Meriaudeau, M. Ferraton, C. Stolz, O. Morel, and L. Bigué, "Polarization imaging for industrial inspection," in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, vol. 6813, 2008.

[31] A. Papoulis and S. U. Pillai, *Probability, Random Variables and Stochastic Processes*. McGraw-Hill, 4th ed., 2002.

[32] E. W. Weisstein, "Spherical harmonics," undated. Accessed 20 March 2010.

[33] R. Courant and D. Hilbert, *Methods of Mathematical Physics*. interscience, 7th ed., 1966.

[34] L. Shen, J. Ford, F. Makedon, and A. Saykin, "A surface-based approach for classification of 3d neuroanatomic structures," *Intelligent Data Analysis*, vol. 8, no. 6, pp. 519–542, 2004.

[35] M. K. Chung, K. M. Dalton, and R. J. Davidson, "Tensor-based cortical surface morphometry via weighted spherical harmonic representation," *IEEE Transactions on Medical Imaging*, vol. 27, no. 8, pp. 1143–1151, 2008.

[36] D. Saupe and D. V. Vranić, "3-d model retrieval with spherical harmonics and moments," in *Proceedings of the 23rd DAGM-Symposium on Pattern Recognition*, pp. 392–397, 2001.

[37] R. Duda and P. Hart, *Pattern Classification and Scene Analysis*. Wiley, 1973.

[38] E. Oja, *Subspace Methods of Pattern Recognition*. Research Study Press, 1983.

[39] K. Fukunaga, *Introduction to Statistical Pattern Recognition*. Academic Press, 2nd ed., 1973.

[40] Z. Wu and R. Leahy, "An optimal graph theoric approach to clustering: Theory and its application to image segmentation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 15, no. 11, pp. 1101–1113, 1993.

[41] J. Shi and J. Malik, "Normalized cuts and image segmentation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 22, no. 8, pp. 888–905, 2000.

[42] P. E. Shames, P. C. Sun, and Y. Fainman, "Modelling of scattering and depolarizing electrooptic devices. i. characterization of lanthanum-modified lead zirconate titanate," *Applied Optics*, vol. 37, no. 17, pp. 3717–3725, 1998.

[43] P. F. Felzenszwalb and D. P. Huttenlocher, "Efficient graph-based image segmentation," *International Journal of Computer Vision*, vol. 59, no. 2, pp. 167–181, 2004.