

Quantitative analysis of multispectral fundus images

I. B. Styles^a, A. Calcagni^{a,b}, E. Claridge^a, F. Orihuela-Espina^a, J. M. Gibson^b

a: School of Computer Science, The University of Birmingham, Edgbaston,
Birmingham, B15 2TT, United Kingdom

b: Birmingham and Midland Eye Centre, City Hospital NHS Trust, Dudley Road,
Birmingham, B18 7QU, United Kingdom

Email: I.B.Styles,E.Claridge@cs.bham.ac.uk

This report summarises the work done during the EPSRC-funded project “Physics-based image interpretation to aid the detection of early signs of retinopathies”, grant number GR/S09906/01.

We have developed a new technique for extracting histological parameters from multi-spectral images of the ocular fundus. The new method uses a Monte Carlo simulation of the reflectance of the fundus to model how the colouration of the tissue varies with differing tissue histology. The model is parameterised by the concentrations of the five main absorbers found in the fundus: retinal haemoglobins, choroidal haemoglobins, choroidal melanin, RPE melanin and macular pigment. These parameters are shown to give rise to distinct variations in the tissue colouration. We use the results of the Monte Carlo simulations to construct an inverse model which maps tissue colouration onto the model parameters. This allows the concentration and distribution of the five main absorbers to be determined from suitable multispectral images. We propose the use of “image quotients” to allow this information to be extracted from uncalibrated image data. The filters used to acquire the images are selected to ensure a one-to-one mapping between model parameters and image quotients. To recover five model parameters uniquely, images must be acquired from six distinct spectral bands. Theoretical investigations suggest that retinal haemoglobins and macular pigment can be recovered with RMS errors of less than 10%. We present parametric maps showing the variation of these parameters across the posterior pole of the fundus. The results are in agreement with known tissue histology for normal healthy subjects. We also present an early result which suggests that, with further development, the technique could be used to successfully detect retinal haemorrhages.

1 Introduction

Pathologies in the ocular fundus are a major cause of blindness. These pathologies include diabetic retinopathy and age-related macular degeneration (ARMD). In both of these conditions, the distribution of blood and/or macular pigments is altered from the normal. These anomalies can impinge upon the foveal area, leading to a degeneration of central vision, with a severe impact on the quality of life. Such pathologies are frequently treatable, provided they are detected at an early stage. In the case of diabetes, modifications to the patient's diet and lifestyle can often control the severity of the condition. In wet/exudative ARMD, laser treatment can be used to destroy the pathological blood vessels that characterise the condition.

At present, primary screening for diabetic retinopathy is largely carried out by general practitioners and ophthalmic opticians, whilst ARMD is usually only detected once the patient has incurred some degeneration of visual function as there is no effective screening tool that can highlight early pathology. In the case of diabetic retinopathy, the primary care practitioners are equipped to detect more obvious pathologies, but can sometimes miss subtle changes or early pathologies where a referral to a specialist ophthalmologist would be advisable. The most common screening tool used by ophthalmic opticians is standard colour fundus photography, which provides an RGB image of the fundus. Modern digital fundus cameras can provide exceptional spatial resolution, allowing very fine detail to be seen. However, the spectral resolution of RGB imaging systems is very coarse, and yields only limited information about the nature of structures visible in an image. In order to determine the origin of a particular feature in an image, more detailed spectral information is required. This has traditionally been provided by spectroscopy, but this approach has its own limitations: spectroscopy yields information only about single points, and coverage of larger areas is impractical. Hence, spectroscopy has not been widely used as a diagnostic tool in the clinical setting and is largely limited to laboratory-based studies.

Recent advances in imaging technology have led to the development of so-called multi-spectral imaging systems. In theory, such systems combine the spatial resolution of conventional photographic systems with the spectral resolution of spectroscopy, although in practice, the spectral resolution is significantly lower than spectroscopy can provide, but still a great improvement over RGB systems. The most common multi-spectral imaging systems use one of two different methods to generate a multi-spectral image. Interferometric systems use an interferometer to build up the Fourier transform of the spectrum at each pixel of the image (a so-called "interferogram"). Typically, the interferometer will scan its interference pattern across the CCD in a number of finite steps, gradually forming the full interference pattern. This can then be inverse-Fourier transformed to obtain the spectrum at each pixel. An alternative method uses a "filter wheel" containing a set of narrow bandpass filters corresponding to the set of wavelengths that are required. The wheel is rotated rapidly in front of the camera and an image is taken for each filter, and the resulting images can be used to deduce the spectrum at each point in the image. Recent advances in technology have led to the development of liquid crystal programmable filters which can be configured into bandpass filters across a range of wavelengths. All of these system can provide high quality images with good spatial resolution, and fair spectral resolution. They are generally excellent at generating multi-spectral image data

for stationary objects, but the increased image acquisition time is problematic when the objects can move. For moving objects, the interferometric methods are generally unsatisfactory as the images must be registered before they are inverse-Fourier transformed: one must register the interferograms. The filter-based approaches can cope with moving objects by using modern registration methods to correct for the movement. This is a relatively straightforward task and there are a number of well-established techniques for achieving this.

It is clear that using these technologies a multi-spectral imaging system for the ocular fundus could be developed. One must correct for involuntary eye movements and so interferometric systems are unsuitable, but filter-based systems are relatively simple to construct and a number of groups have successfully done so (Harvey et al., 2002; Bessho et al., 2005). The quantity of information that multi-spectral imaging can provide could potentially be a very valuable aid in the early detection of pathologies, but it is not yet clear how one should interpret the data that these systems yield: human perception is limited to approximately RGB, and so it is likely to be rather hard for a clinician to interpret multi-spectral images directly. We propose a method for processing multi-spectral fundus images which yields quantitative information about the histological properties of the tissue. This information may be diagnostically useful in a clinical setting.

The basis for our proposed method is a physical model of light propagation in the ocular fundus. The model is parameterised by relevant histological variables (such as the concentration of blood haemoglobins). Using the model, we deduce a set of multi-spectral measurements which allow the histological variables to be computed from the multi-spectral images. The values of these variables (corresponding to quantities of pigment) can then be plotted at each pixel to generate a “parametric map” which shows how the pigment is distributed across the image. This could allow primary care clinicians to identify pathologies at a much earlier stage of development, allowing the patient to receive appropriate treatment in good time.

These ideas have previously been employed in dermatology, where the SIAscope system (Cotton et al., 1997) has been shown to be an effective screening tool in the diagnosis of pigmented skin lesions (Moncrieff et al., 2002). The SIAscope uses RGB and near-infrared images to map the distributions of dermal blood, dermal melanin and collagen in skin lesions, allowing suspicious lesions to be investigated non-invasively. This reduces the number of unnecessary biopsies from healthy subjects. Standard RGB images, plus one near-infrared image are sufficient for this information to be deduced from skin images, but the tissue of the ocular fundus is much more complex and a more careful choice of filters, based on some objective criteria, is required. In addition, the tissue of the ocular fundus is not directly accessible and it is not possible to calibrate the images using calibration patches. In this article, we propose methods by which these difficulties may be overcome, and show how these ideas allow parametric maps of the ocular fundus to be computed.

2 Modelling the ocular fundus

The methods of image analysis proposed in this work depend upon the construction of an accurate model of the reflectance of the ocular fundus. There have been several attempts to construct such a model, and Berendschot et al. (2003) provides an extensive summary of the most important works. The first model which took into account the whole

of the ocular fundus was van Norren and Tiemeijer (1986), who developed a model using a series of neutral reflectors and absorbing compounds. This was modified by Delori and Pfibsen (1989), who used Kubelka-Munk transport theory (Kubelka and Munk, 1931) to include the effects of scattering in the choroid. Hammer and Schweitzer (2002) used the adding-doubling method of radiation transport to include scattering in the retina and choroid, drawing on earlier work on the optical properties of these tissues by Hammer et al. (1995). Most recently Preece and Claridge (2002) used Monte Carlo simulation (Prahl et al., 1989) in combination with the data from Hammer et al. (1995) and the well-known properties of the main pigments in the fundus (haemoglobins (Horecker, 1942) and melanin (Anderson and Parrish, 1981)). The Monte Carlo method is more accurate than cruder approximations such as adding-doubling, but suffers from long computation times. Fortunately, modern computers allow Monte Carlo to be a reasonable practical method of performing transport calculations, although computation times are still significant.

Preece and Claridge’s model was parameterised by the concentrations of the main absorbers in the fundus:

- Concentration of melanin in the RPE
- Concentration of melanin in the choroid
- Concentration of haemoglobins in the choroid (assuming 95% oxygen saturation).

All other properties (layer thickness) of the tissue were assumed to be constant, taking the values from Rohen (1977). The specific absorption coefficients of these pigments are shown in Fig. 1.

This model was found to accurately reproduce reflectance spectra obtained from the perifoveal region of the fundus, away from any major blood vessels. In other parts of the fundus, other variable components are known to be present. Surrounding the fovea is the macula region containing macular pigment (Wyszecki and Stiles, 1982) (see Fig. 1(d)). Changes in the distribution of this pigment can occur in the early stages of several pathologies and it would therefore be very useful to be able to quantify variation. There is also a fine network of blood vessels within the retina itself. These are important since diabetic haemorrhages occur in these vessels and quantification of haemoglobins in the retina may allow us to identify such features. We have therefore added two parameters to Preece and Claridge’s model: concentration of macular pigments and concentration of retinal haemoglobins (where we assume 80% oxygen saturation (Alm and Bill, 1973)).

A schematic cross-section through the ocular fundus is shown in Fig. 2. The five variable parameters characterising our model represent concentrations of pigment in the different layers:

- C_{MP} : concentration of macular pigments in the retina.
- C_{RH} : concentration of haemoglobins in the retina.
- C_{RM} : concentration of melanin in the RPE.
- C_{CM} : concentration of melanin in the choroid.
- C_{CH} : concentration of haemoglobins in the choroid.

The Monte Carlo model computes the reflectance $R(\mathbf{p}, \lambda)$ at wavelength λ of tissue described by a parameter vector $\mathbf{p} = (C_{MP}, C_{RH}, C_{RM}, C_{CM}, C_{CH})$. It is useful to examine how $R(\mathbf{p}, \lambda)$ varies as we change each of the parameters (holding the others constant). Fig. 3 shows how the reflectance of fundus tissues varies as its composition is changed. Note that a log scale is used to emphasise small changes. In each of these graphs, four of the five parameters are constant, and the fifth parameter is varied. Each of the identified histological components has a distinct effect on $R(\mathbf{p}, \lambda)$. In Fig. 3(a) we see that as the concentration of retinal macular pigment is increased, the reflectance is attenuated in the blue and green spectral regions, with no effect in the red since macular pigment does not absorb at $\lambda > 534\text{nm}$. As the concentration of retinal haemoglobins is increased (Fig. 3(b)) we see that the spectral reflectance attenuates the spectrum in blue and green regions, and introduces strong features due to peaks and troughs in the specific absorption coefficient of the haemoglobins. Increasing the concentration of RPE melanin affects the whole reflectance spectrum (Fig. 3(c)). A careful analysis shows that an increase in C_{RM} causes more attenuation in the blue spectral region than in the red, and the change is not a straightforward wavelength-independent scaling as might be first thought. Changes in choroidal melanin (Fig. /refsubfig:ccmvar) have a very strong effect in the red region of the spectrum, but a much weaker effect in the blue region. Note well the contrast between increase in RPE melanin and choroidal melanin. Whilst increases in RPE melanin attenuate the spectrum in roughly equal amounts in blue and red (there is a small difference), the effect of choroidal melanin is much stronger in the red, even though the pigments are identical. This difference arises from the different environments these components are found in. In the RPE, there are no other significant absorbers and there is relatively weak scattering from the RPE tissue itself. In the choroid, there are other absorbers present (haemoglobins), and scattering from the underlying tissue is very strong. It is the interplay between the varying component and its surroundings which gives rise to these differences, and this emphasises the importance of accurate modelling of scattering processes which have a strong effect on the behaviour of the model. Finally, increases in choroidal haemoglobins (Fig. 3(e)) lead to a reduction in reflectance in the green region, where absorption by haemoglobins is very strong. Note again the difference between choroidal and retinal haemoglobins, again due to the interaction with the surrounding tissues.

3 Analysis of the Imaging Process

3.1 Factors affecting the imaging process

Any attempt to optimise the imaging process must be based on a thorough understanding of that process. The imaging modality that concerns us in this work is visible optical photography. The ease with which such images can be obtained allows fundus photography to be a potentially powerful screening tool for a variety of pathologies, provided the imaging process has been optimised to allow the desired information to be extracted.

The process of obtaining an optical image is very simple. The object under study is illuminated using a light source with known spectral properties. A camera is focused on the object and one or more filters is used to select a band of wavelengths to capture (wavelengths outside this band are blocked by the filter). The resulting filtered image

is captured on a CCD photon detector which records the total intensity of light falling on each pixel of the detector, resulting in a monochrome image. Colour images are constructed using red(R), green(G) and blue(B) filters which approximate the response of the photo-receptors in the human eye, but this is only for the purpose of realistic display, and any set of filters can be used if this is not important.

There are a number of factors which affect the strength of the signal recorded by the detector at position \mathbf{x} :

- The spectrum and intensity of the illuminant, $I(\mathbf{x}, \lambda)$
- The reflectance of the tissue, $R(\mathbf{x}, \lambda)$
- The profile of the filter used, $F_n(\lambda)$
- The quantum efficiency of the CCD, $Q(\lambda)$

The signal (image value) recorded on the detector through filter n is then

$$i_n(\mathbf{x}) = C \int_{\Lambda} I(\mathbf{x}, \lambda) R(\mathbf{x}, \lambda) Q(\lambda) F_n(\lambda) d\lambda, \quad (1)$$

where the constant C scales the signal onto the range of the detector and is set by factory calibration. It is therefore possible to compute the signal recorded on the CCD provided that all of the terms are known. One could, in principle, take the image value recorded by the detector and deduce the model parameters which describe the reflectance of the tissue. In practice these terms are not all known, and we must consider the problem more carefully.

3.2 Formalising the imaging process

Preece and Claridge (2004) described a mathematical model of the image formation process. They assumed that the optical properties of the object under consideration could be described by a vector of K parameters

$$\mathbf{p} = \{p_k\}_{k=1\dots K}, \quad \mathbf{p} \in \mathcal{P}, \quad (2)$$

where \mathcal{P} is the space of all possible parameter vectors for a given object. The first stage in the imaging process is the formation of the reflectance spectrum of the tissue. The reflectance spectrum of the tissue is a continuous function of optical wavelength:

$$r = r(\lambda), \quad r \in \mathcal{R}, \quad (3)$$

where \mathcal{R} is the space of all possible reflectance spectra. We assume that there exists a mapping

$$a : \mathcal{P} \mapsto \mathcal{R} \quad (4)$$

which defines the relationship between object parameters and spectral reflectance. In Sect. 2 we discussed how such a model could be constructed for the tissues of the ocular fundus using Monte Carlo simulations. This model is an implementation of a .

Given the reflectance spectrum of an object, an imaging device typically acquires images by applying band-pass filters to the remitted spectrum and integrating over wavelength to generate a single image value per filter. For N optical filters $\{F_n(\lambda)\}_{n=1\dots N}$, an image vector

$$\mathbf{i} = \{i_n\}_{n=1\dots N}, \mathbf{i} \in \mathcal{I} \quad (5)$$

is recorded, where \mathcal{I} is the space of all possible image vectors from a given object. The process of image acquisition can be represented by a mapping

$$b : \mathcal{R} \mapsto \mathcal{I}, \quad (6)$$

and in the cases we will be studying, the value of the n -th component of the image vector is defined by Eqn. (1).

We are now able to compute a composite mapping $m = a \circ b$ from model parameters to image values: knowing the composition of the tissue at a given point allows us to deduce the appearance of the tissue in an image. In principle, therefore, we ought to be able to perform the inverse calculation, and deduce the composition of the tissue from its appearance in an image, subject to some mathematical constraints which we discuss in section 4.1. However, the ability to construct the inverse mapping depends on a complete knowledge of the forward mapping. We cannot always guarantee this: some of the terms in Eqn. (1) are unknown in certain situations.

When acquiring images of the ocular fundus, illumination and acquisition is done through an aperture (the pupil). In addition, the ocular fundus is strongly curved (a sphere of diameter $d \approx 24\text{mm}$). As a result, illumination of the fundus is uneven, with the centre illuminated more brightly, and the edges being significantly darker. It is not possible to determine the level of illumination in the fundus as the tissue is not accessible, and we cannot know $I(\mathbf{x}, \lambda)$. Fortunately there are some statements that we can make about $I(\mathbf{x}, \lambda)$ which lead to a resolution of this problem.

The key observation is that the only aspect of the illuminant which changes with position is its intensity: the relative contributions from different wavelengths do not change with position. We can therefore separate $I(\mathbf{x}, \lambda)$ into spectral and spatial components: $I(\mathbf{x}, \lambda) = A(\mathbf{x})S(\lambda)$, with no loss of generality. Substituting this into Eqn. 1 provides

$$i_n(\mathbf{x}) = A(\mathbf{x}) \int_{\Lambda} S(\lambda) R(\mathbf{x}, \lambda) Q(\lambda) F_n(\lambda) d\lambda, \quad (7)$$

where we have absorbed the calibration constant C into $A(\mathbf{x})$. This substitution allows us to remove the effects of spatially varying illumination by taking ratios of image values:

$$q_{mn}(\mathbf{x}) = \frac{i_m(\mathbf{x})}{i_n(\mathbf{x})}, \quad (8)$$

which are independent of $A(\mathbf{x})$. The use of “image quotients” has previously been employed in image segmentation problems (Maloney and Wandell, 1986; Healey, 1989) and is often referred to as “discounting the illuminant”, but this is the first time they have been used to provide quantitative information about the tissue itself.

We now formally introduce this new step into our mapping. Given an N -component image vector, we can form an $(N - 1)$ -component vector of independent image ratios,

$$\mathbf{q} = \{q_m\}_{m=1\dots N-1} = \{i_m/i_N\}_{m=1\dots N-1}, \mathbf{q} \in \mathcal{Q}. \quad (9)$$

This can be represented as a mapping

$$c : \mathcal{I} \mapsto \mathcal{Q} \tag{10}$$

The process of forming image ratios from an object can then be represented by a composite function

$$f = a \circ b \circ c : \mathcal{P} \mapsto \mathcal{Q} \tag{11}$$

relating the composition of the object described by a parameter vector \mathbf{p} to the image ratios \mathbf{q} formed from image captured by a set of optical filters $\{F_n(\lambda)\}_{n=1\dots N}$.

Having now defined a mapping from object composition to image ratios we can, in principle construct an inverse mapping

$$f^{-1} : \mathcal{Q} \mapsto \mathcal{P}, \tag{12}$$

(subject to some mathematical constraints which we discuss in Sect. 4.1) which allows us, given a vector of image quotients \mathbf{q} , to deduce the vector of parameters \mathbf{p} that describes the object at the point of interest.

The use of image quotients in the image analysis process holds several advantages over the use of image values. As we have already observed, it effectively removes the effect of uneven illumination on a object. Furthermore, it removes any dependence on the overall intensity of the illuminant: any scaling of the intensity can simply be absorbed into $A(\mathbf{x})$. When working with image values, one measures the intensity of light reaching the object of interest by placing an object of known reflectance on or very near the object, within the field of view. This is not required when working with image quotients, and allows these methods to be used when the geometry of the object makes the use of reflectance tiles impractical, for example when the object is irregular or very curved in shape, or the object must be imaged remotely as it cannot be directly accessed. Both of these are true of the ocular fundus, and so the use of image quotients is essential if we are to perform a quantitative analysis of fundus images.

The discussion so far has focused on what can be done in principle, and there are still a number of practical issues to consider. For example, we have not considered what kinds of filter must be used, or whether the spectrum of the illuminant and quantum efficiency of the camera must meet certain criteria. For a given $S(\lambda)$, $Q(\lambda)$, the only unanswered question is that of which filters should be used. In the following section we discuss objective criteria for filter selection which aim to ensure that the choice of filters allows f^{-1} to be computed, and to minimise the error with which model parameters can be computed from a set of images.

4 Selection of Optimal Filters

There are two factors which govern our choice of filters. The filters must ensure that the inverse mapping f^{-1} exists and is unique, and they must aim to minimise the error in the parameter recovery process. Let us deal with these separately.

4.1 Existence of f^{-1}

The method we have proposed depends upon the existence of f^{-1} . We can define some conditions which, if met by f , guarantee that f^{-1} exists, although they do not tell us what

f^{-1} is. We begin by noting that \mathcal{P} and \mathcal{Q} are vector spaces, with points in \mathcal{P} represented by some vector \mathbf{p} , and points in \mathcal{Q} represented by some vector \mathbf{q} . We represent the mappings $f : \mathcal{P} \mapsto \mathcal{Q}$ and $f^{-1} : \mathcal{Q} \mapsto \mathcal{P}$ as $\mathbf{q} = f(\mathbf{p})$ and $\mathbf{p} = f^{-1}(\mathbf{q})$. We will write $\mathbf{q} = \mathbf{q}(\mathbf{p})$ and $\mathbf{p} = \mathbf{p}(\mathbf{q})$.

The key quantity which allows us to determine the existence and uniqueness of f^{-1} is the Jacobi matrix, which for this problem is (in component form)

$$\mathbf{J}_{ij} = \frac{\partial \mathbf{q}_i}{\partial \mathbf{p}_j}. \quad (13)$$

From the Jacobi matrix, we compute the Jacobian $J = \det(\mathbf{J})$. The inverse function theorem (Schutz, 1980) states that f is bijective (one-to-one and onto) if and only if $J \neq 0 \forall \mathbf{p} \in \mathcal{P}$. In practical terms, this means that $\text{sgn}(J)$ must be strictly positive or strictly negative (thus never zero) throughout \mathcal{P} .

Note that in order to compute $\mathbf{J}(\mathbf{p})$ numerically, \mathcal{P} must be discretised on a rectangular grid to allow us to compute the partial derivatives using finite difference methods. In addition, we can use the Jacobi matrix to approximate the mapping near the grid of points: since $d\mathbf{q} = \mathbf{J}d\mathbf{p}$, we can write $\mathbf{q} = \mathbf{q}_0 + \mathbf{J}|_{\mathbf{p}_0}d\mathbf{p}$ in the vicinity of \mathbf{p}_0 , where $\mathbf{q}_0 = \mathbf{q}(\mathbf{p}_0)$. We can also compute a first approximation to f^{-1} using a similar method, since $d\mathbf{p} = \mathbf{J}^{-1}d\mathbf{q}$ and thus $\mathbf{p} = \mathbf{p}_0 + \mathbf{J}^{-1}|_{\mathbf{q}_0}d\mathbf{q}$.

A very important consequence of this result is that the \mathbf{J} must be a square matrix: to compute K model parameters, we must have K independent image quotients. This requires $N = K + 1$ image values. We choose one of the image values to act as a common denominator, and then

$$q_j = \frac{i_j}{i_N}, \quad j = 1 \dots N. \quad (14)$$

4.2 Minimisation of Error

It is important that any parameter values recovered from a set of images should be known to the greatest accuracy possible in order to be of use in a clinical setting. There are a number of sources of possible error in the procedure that we have described.

4.2.1 Error in the construction of f

The forward model of the imaging process described in Sect. 3.2 contains a number of sources of error which mean that the forward model f can only compute image quotients to finite accuracy. Let us consider each stage of the composite mapping.

The first step is the mapping a from model parameters to reflectance spectra. Assuming that the model is an accurate representation of the tissue in its construction (a point which we discuss in Sect. 7), we identify two main sources of error in this step. The first is error due to the stochastic nature of the Monte Carlo model of light transport. The model simulates the passage of a large number of photons N through the tissue, and then computes the average probability of the photon being reflected. In the limit $N \rightarrow \infty$ this calculation converges to the exact solution, but is otherwise an approximation. We cannot, of course, compute at the limit, but choosing a suitably large N can minimise any variation between simulations. We have found that $N = 100,000$ provides very stable solutions.

The second source of error in this step is that resulting from errors in the experimental data which characterise the tissue properties. The absorption and scattering cross-sections for the pigments, and for the underlying tissue are derived from experimental measurements (Hammer et al., 1995; Horecker, 1942; Anderson and Parrish, 1981; Wyszecki and Stiles, 1982). In our computations, the mean values stated in the literature are used, but there is an experimental error associated with each of the measurements. Using relatively simple error analysis (Kendall and Stuart, 1969), we can propagate the error in the physical data through our computations to calculate the error in the reflectance spectra. We can then propagate this through mapping i to give the errors in the image values computed by the model. Further information on this method, applied to skin imaging, can be found in Preece and Claridge (2004). We will denote the error (standard deviation) of the j 'th image value as $\sigma(i_j)$.

We now propagate the $\sigma(i_j)$ through to the image quotients. From Eqn. 14 we can compute the error in the k 'th image quotient using $(\sigma(q_k))^2 = \sum_{j=1}^N (\partial q_k / \partial i_j)^2 (\sigma(i_j))^2$ (Kendall and Stuart, 1969). It can be shown that the resulting error in the image quotients is

$$\sigma^f(q_j) = \frac{1}{i_N} \sqrt{\left(q_j^2 \sigma^2(i_N) + \sum_{k=1}^{N-1} \sigma^2(i_k) \right)}. \quad (15)$$

We can now use this to compute the error in the recovered model parameters due to these modelling errors, and we find that

$$\sigma^f(p_j) = \sqrt{\sum_{k=1}^N \left(\frac{\partial p_j}{\partial q_k} \right)^2 (\sigma^f(q_k))^2}. \quad (16)$$

The $\partial p_j / \partial q_k$ can be read directly from the inverse of the Jacobian of f . The $\sigma^f(p_j)$ are the errors in f^{-1} resulting from errors in the construction of the imaging model f , and do not take into account any other errors in the imaging and analysis process.

4.2.2 Error in the imaging process

The second source of error we have identified is error in the image acquisition process. In this paper, we are concerned with acquiring images using a CCD and the discussion will reflect this. However, the idea is quite general: the detector does not record the intensity of light it receives with complete accuracy.

When recording an image on a CCD, each pixel records an image value. On a typical 8-bit camera, the image values are integers in the range 0 (black) to 255 (white). The 12-bit camera used to acquire the images in this paper records integers from 0 to 4095. These ranges are set by the number of bits used to record each image value (8 or 12 in these cases), and their correspondence to different levels of brightness is set by a factory calibration process which leads to the constant C in Eqn. 1.

The value recorded by the camera is not exact. It is essentially the sum of the numbers of photons received at a given pixel during an exposure. Even in essentially steady-state conditions, this number can vary significantly. In addition, the CCD is a semiconductor device and as such is sensitive to thermal effects. The signal recorded by the CCD can be regarded as the mean intensity, and there is an associated error, which, for an 8-bit

camera, is typically ± 2 units. The exact value depends on the details of the camera. This error will affect the accuracy with which images can be used to recover model parameter, and we can quantify this. Denoting the “camera error” of the j ’th image value as $\sigma^C(i_j)$, we can show that

$$\sigma_{p_n}^C = \frac{1}{i_N} \sqrt{\sum_{m=1}^K \left(\frac{\partial p_n}{\partial q_m} \right)^2 \left[(\sigma_{i_m}^C)^2 + (q_n \sigma_{i_N}^C)^2 \right]} \quad (17)$$

This error is the error in parameter recovery due solely to errors in the image acquisition process.

4.2.3 Error in the construction of f^{-1}

The third source of error we have identified is due to the difficulties in constructing f^{-1} . The forward model is based on a numerical solution of the radiative transport equation, and so we cannot deduce an analytic solution. Instead, we must deduce an approximate solution, given the forward model.

The numerical nature of the forward model means that we can only compute f at a finite number of points in \mathcal{P} : f is effectively a mapping from N discrete points in \mathcal{P} to N discrete points in \mathcal{Q} . Assuming that the conditions in Sect. 4.1 are met, this mapping is bijective, and the inverse mapping from N points in \mathcal{Q} to N points in \mathcal{P} exists and is unique. Clearly we know f^{-1} at this set of discrete points, but not at points which lie in between. We need an approximation of f^{-1} .

A natural approximation to employ is a truncated Taylor series. The constraints imposed in Sect. 4.1 ensure that f and f^{-1} are analytic functions. We can approximate f in the vicinity of one of the known points, \mathbf{p}_0 as a Taylor series,

$$q_i(\mathbf{p}_0 + \mathbf{d}\mathbf{p}) = q_i(\mathbf{p}_0) + \sum_j \frac{\partial q_i}{\partial p_j} \Big|_{\mathbf{p}_0} dp_j + \frac{1}{2} \sum_{j,k} \frac{\partial^2 q_i}{\partial p_j \partial p_k} \Big|_{\mathbf{p}_0} dp_j dp_k + \dots, \quad (18)$$

and this can be computed to arbitrary precision, since all the derivatives are known. We can similarly write down a Taylor expansion of f^{-1} :

$$p_i(\mathbf{q}_0 + \mathbf{d}\mathbf{q}) = p_i(\mathbf{q}_0) + \sum_j \frac{\partial p_i}{\partial q_j} \Big|_{\mathbf{q}_0} dq_j + \frac{1}{2} \sum_{j,k} \frac{\partial^2 p_i}{\partial q_j \partial q_k} \Big|_{\mathbf{q}_0} dq_j dq_k + \dots. \quad (19)$$

We cannot compute this to arbitrary precision since we cannot numerically compute the derivatives: the points in \mathcal{Q} do not generally lie on an orthogonal grid. We can, however, compute the expansion to first order by using the inverse of the Jacobi matrix of f , as noted in Sect. 4.1. This provides a linear approximation for f^{-1} which is wholly inadequate, unless points in the model are close enough for the radii of convergence about each point to overlap. This requires a very large number of points which is computationally unfeasible. Another method for approximating the inverse function is required.

We have found that the use of a neural network to approximate the inverse function provides acceptable results. The particular type of network used is the radial basis network provided in the MATLAB Neural Network toolbox. Details of the implementation of the inverse function are provided in Sect. 5.2, and are not important in the context of the error analysis. The important point about the implementation of f^{-1} is that it is not

exact: we have to approximate. The approximation has an associated error which can be quantified and taken into account in the optimisation process. The choice of filters will affect the nature of f and f^{-1} , and hence will also affect the accuracy with which these functions can be approximated. We can define a measure of the approximation error: $\sigma_{p_n}^A$ is the error in the recovery of parameter n due to approximations in the construction of f^{-1} .

4.3 Objective procedure for filter selection

We are now in a position to define an objective procedure for the selection of optimal filters.

1. Discretise the $K = N - 1$ -dimensional parameter space using a rectangular grid to give M distinct parameter vectors $\{\mathbf{p}_i\}_{i=1\dots M}$
2. For each parameter vector \mathbf{p}_i , compute the spectral reflectance of the tissue (corresponding to mapping a as defined in Eqn. (4)), giving $\{R_i(\lambda)\}_{i=1\dots M}$.
3. Define a set of N independent (orthogonal) filters $\{F_n(\lambda)\}_{n=1\dots N}$, subject to any necessary constraints.
4. For each $R_i(\lambda)$, apply the N filters to provide image values.
5. Select one of the N filters to be the denominator in the image quotients, and compute $K = N - 1$ image quotients for each of the M parameter vectors. This defines f .
6. For each \mathbf{p}_i , compute the Jacobian J_i
7. If $\text{sgn}(J_i)$ is constant and non-zero $\forall i$, the mapping is unique and the procedure can continue; else the mapping is not unique and the filters $\{F_n(\lambda)\}_{n=1\dots N}$ are unsuitable.
8. If the mapping is unique, compute the total error in parameter recovery: $\sigma^T = \sqrt{(\Sigma^C)^2 + (\Sigma^f)^2 + (\Sigma^A)^2}$, where $\Sigma^C = \sqrt{\sum_n (\sigma_{p_n}^C)^2}$ etc.
9. Employ an optimisation procedure to find the set of filters which minimises σ^T .

This procedure provides a set of filters which ensures that f and f^{-1} are well-behaved (bijective) mappings, and which minimises the error with which parameters can be recovered using the inverse model. In Sect. 5 we explain how these ideas are implemented in practice.

5 Implementation

Let us now discuss how these ideas are used in practice for imaging of the ocular fundus. We will begin by describing the equipment that is to be used for acquiring suitable images, since this imposes some constraints on our choice of filters. We then describe the implementation and optimisation of the forward and inverse models.

5.1 Image acquisition

Acquisition of ocular fundus images requires specialised equipment, since images must be taken through a small aperture and the ocular lens. We used a Zeiss RCM250 fundus microscope which contains the optics that are required to obtain fundus images, inbuilt sources of illumination, and the mechanics required to optimally position the camera. In standard form, the camera records images on 35mm film using a SLR camera body with modified bayonet lens fitting which is peculiar to Zeiss fundus cameras. Viewing illumination is via a halogen lamp which illuminates the fundus through the camera optics. Images are recorded using a xenon flash.

A number of modifications were made to the microscope to make it suitable for this work. The SLR body was replaced by a QImaging Retiga EXi 12-bit monochrome digital camera which allowed high quality digital images to be acquired. A Cambridge Research Instruments VariSpec LCD programmable filter was added to the system, in front of the Retiga camera. This device is a configurable interference filter which can be programmed to implement a range of narrow Gaussian bandpass filters with central wavelengths in the range 400-700nm. The width and transmittance of these filters is dependent on the central wavelength, and this provides a constraint on our choice of filters. An additional lens with focal length 75mm was included to shrink the image onto the 1/3 inch CCD which is much smaller than the 35mm film used in the original camera. An array of optical connectors was required to integrate these components with the microscope. The VariSpec filter and the Retiga camera were controlled using a standard notebook computer running custom software. In brief, the software pre-loads the required filters into the VariSpec filter. Images are then acquired (using appropriate exposure settings) with each filter sequentially. Illumination was provided by the viewing lamp at a suitable intensity. The xenon flash used during conventional photography was found to be unsuitable because it provides a transient illumination rather than the steady-state illumination assumed in our model, and because the spectrum of a xenon discharge lamp contains a number of sharp emission peaks which bias the filter selection process.

The lengthy exposure time required to obtain all of the necessary images means that image registration is required due to movement of the subject. Excellent results were obtained using an algorithm due to Stewart et al. (2003). The resulting registered images were then corrected to compensate for differing exposure settings, with all images normalised to an exposure time of one second, a gain of one, and zero offset.

5.2 Construction and optimisation of f and f^{-1}

The parameter space \mathcal{P} was discretised on a rectangular grid. The parameter values chosen for each of the five variables are shown in table 1. This discretisation of \mathcal{P} provides $M = 5 \times 5 \times 6 \times 6 \times 6 = 5400$ points at which the model is defined.

For each parameter vector $\{\mathbf{p}_i\}_{i=1\dots M}$, the corresponding reflectance spectrum was computed using the Monte Carlo model of tissue reflectance described in Sect. 2. The spectra were then used in the filter selection procedure described in Sect. 4.3. An evolutionary algorithm (Yao and Liu, 1997) was used to find the set of $N = 6$ filters which provided a unique mapping and minimised the error of parameter recovery. The filters were selected from the set of filters that can be implemented by the VariSpec filter and were specified only by their central wavelength, with all other properties dependent on

that central wavelength. The optimal filters were found to have central wavelengths 507nm, 525nm, 552nm, 585nm, 596nm, 611nm (to be used as the denominator in the image quotients), and these are shown in Fig. 4.

The optimal filters were applied to each of the M reflectance spectra computed using the model to give M vectors of six image values. From these, M vectors of five image values were computed by dividing the image values computed using the filters centred at 507nm, 525nm, 552nm, 585nm and 596nm by the image value computed using the filter centred at 611nm. This establishes a mapping between the original M parameter vectors in \mathcal{P} and M image quotient vectors in \mathcal{Q} and provides a definition of both f and f^{-1} at these M points.

The discretised model was used to train a neural network to approximate f^{-1} . The particular type of neural network used was the radial basis network provided in the MATLAB Neural Network Toolbox. The function `newrbe(P,T,SPREAD)` was used to train the network. The M image quotient vectors generated from the model were used as input data P , and the M parameter vectors were the target outputs T . The training parameter `SPREAD` was set to 3.0, which was found to give good results on unseen data. Approximately 350 random parameter vectors were generated and the corresponding image quotients were computed. These quotients were then given as input to the network which computed the corresponding parameter vectors. These were then compared with the original parameter vectors. The results of these tests are shown in Fig. 5. It is clear from these data that not all parameters are recovered equally well. We see quite clearly that whilst C_{MP} and C_{RH} are recovered with acceptable accuracy, the remaining parameters cannot be recovered. The reason for this is not fully understood but is believed to be because of the highly non-linear nature of f and f^{-1} . At low concentrations of pigment, small changes in concentration cause large changes in the reflectance of the tissue, and hence large changes in the image quotients. At high concentrations of pigment, a large change in the quantity of pigment is required to produce a small change in the tissue reflectance. This nonlinearity makes the functions f and f^{-1} very hard to approximate, and more work is needed in this direction. In our analysis and discussion in sections 6, 7, we will restrict our attention to the distributions of macular pigment and retinal haemoglobins.

After the image acquisition process described in Sect. 5.1, the images were processed by the neural network. For each pixel in the images, the appropriate image quotients were calculated directly from the image values recorded. The resulting vectors of image quotients were given as input to the neural network, which returned a parameter vector for each point in the image. The values of each parameter were displayed graphically as parametric maps which show how each of the parameters varies across the image.

6 Results

Images were obtained from a range of subjects spanning a wide variety of ethnic backgrounds including Caucasian, African, Indo-Asian, Afro-Caribbean and Mediterranean subjects. Images were initially taken only from subjects with no known pathology, and these were used to verify the technique. Further images were then obtained from subjects with known pathologies (for example, retinal haemorrhages, choroidal neovascular membranes) in order to investigate the ability of the technique to flag these pathologies. In order to provide maximum flexibility in the processing of the images, extra images were

recorded using additional filters chosen to highlight notable features in the reflectance spectra of the tissue. The central wavelengths of the chosen filters were at $\lambda = \{480\text{nm}, 507\text{nm}, 515\text{nm}, 525\text{nm}, 540\text{nm}, 548\text{nm}, 552\text{nm}, 560\text{nm}, 565\text{nm}, 569\text{nm}, 575\text{nm}, 585\text{nm}, 596\text{nm}, 611\text{nm}, 640\text{nm}, 675\text{nm}, 705\text{nm}\}$. Note that these include the optimal filters deduced from the properties of the model.

In order to verify the correctness of the model, image data taken from all normal subjects were collated and compared with the model to investigate the correspondence between model and data. Projections of the model and image data onto the q_i - q_j planes were used to check the correspondence. These are shown in Fig. 6. It is clear from these projected view of the data that the model and image data are not coincident. There is considerable overlap, and the shapes of the projections are roughly similar, but the image data clearly does not fall within the bounds of the model data. There are several possible reasons for this, although the causes of these reasons are not so easy to identify. Firstly, it may be that the model does not fully capture the whole range of variability within normal tissue. Extending parameter space to include higher concentrations of the variable pigments has little effect on the extent of the model. It may be necessary to introduce additional variable parameters (e.g. layer thicknesses) in order to model the full variability of normal tissue.

A further problem is that the data used to construct the model may be inaccurate or inappropriate. For instance, the only data on absorption and scattering by ocular tissues is from Hammer et al. (1995), who took ex vivo measurements on excised samples of bovine and porcine fundii. These are physiologically and histologically very similar to human fundi and much of the recent work on modelling of fundus reflectance (Preece and Claridge, 2002) has used these data as a starting point. However, this is clearly not an ideal solution, and data from human subjects would surely improve the results.

In Sect. 7, we offer some additional possible causes of this mismatch for which solutions require resources which are currently unavailable to us. In the absence of an ab initio method for resolving the problem with the resources available, we elected to introduce an empirical correction to the model, based upon the normal image data.

Although the image and model data are not coincident, the shapes of the projected datasets are similar. That is, where the projection of the model is “long and thin”, as in, for example, the q_3 - q_4 projection, the projected image data is also “long and thin”. This suggests that whilst the absolute position and size of the model in \mathcal{Q} is wrong, the behaviour of the model with respect to the constituent parameters may be correct. In other words, an increase in one of the parameters causes a move along a path in \mathcal{Q} . An increase in the same parameter in the real tissue causes a movement in a similar direction. This means that provided we can map the model and image data to be coincident, we can use the directionality within the model to extract parameter values from the image data. Although the absolute values of the recovered parameters may be incorrect, the relative values should be correct, and the parametric maps may reveal the distribution of parameters in the image.

We therefore choose a simple “shift and scale” operation to align model and image data. We first transform the image data such that the means of model and each set of image data are the same. We then scale the model such that it encapsulates all of the normal data. A scale factor of four was found to be sufficient to encompass the normal image data.. The two-dimensional projections of the scaled model are shown in Fig. 7.

The scaled model was then used as a basis for the interpretation of image data. A neural network was trained on the scaled data and this was used as the inverse function. Note that one important consequence of this scaling is that the effective parameter ranges become reduced, as shown in the test data of Fig. 8.

For all normal and pathological subjects, the mean of the image data was moved to the mean of the model data, and the translated image data was then used as input to the neural network. The resulting vectors of model parameters, one for each point in the image, were then displayed as parametric maps, each showing the variation of the corresponding histological parameter across the image. Due to the difficulties in recovering some parameters with sufficient accuracy, we will analyse only the parametric maps showing the distributions of macular pigment and retinal haemoglobins.

6.1 Analysis of macular pigment distribution

In a normal human fundus, there exist a family of pigments known as macular pigments which are present in the retina. Macular pigments are strongly localised in the area around the fovea, where the density of photoreceptors is highest. It is believed (Landrum et al., 1997; Beatty et al., 1999) that their principal function is to protect the photoreceptors from high intensities of blue and ultra-violet light which is well known to be damaging to the receptors. This pigment is therefore an ideal test for our methods, since the normal distribution is well established. In Fig. 9, we show several examples of parametric maps of macular pigment, alongside RGB pictures of the same subjects. These images shown a clear peak in the macular pigment distribution in the vicinity of the fovea. This is consistent with the known histology of the fundus. There is also a clear peak in the vicinity of the optic disk. This “false positive” occurs because the optic disk has a significantly different histology to the rest of the fundus. Since we assume that scattering and thickness of the retina, RPE and choroid and constant (and non-zero), the optic disk is not well represented by the model and parameters recovered in this region cannot be regarded as accurate. In Sect. 7, we examine the validity of this assumption.

Not all subjects showed the same clear macular pigment peak in the fovea that is evident in Fig. 9. We believe that this is predominantly due to inherent problems with our imaging technology which we discuss in detail in Sect. 7. In brief, the prolonged and repetitive exposure used means that eye movement is inevitable. Whilst the images can be corrected spatially using advanced registration methods (Stewart et al., 2003), constancy of light intensity is vital for our techniques to succeed, and any movement necessarily compromises this.

6.2 Analysis of retinal haemoglobin distribution

Blood is distributed to the tissues of the retina via a network of arteries, capillaries and veins. The larger vessels (the main arteries and veins) are clearly visible to the human observer whilst the capillary bed provides a pinkish tint to the bulk of the fundus. When large amount of other pigments are present, this tint is rather less obvious than in relatively unpigmented subjects. In all subjects, the presence of macular pigment at the fovea tends to obscure this tint, making it harder for the human observer to discern any variations in blood concentrations near the fovea. This is important, since the fovea is normally

surrounded by an “avascular zone”, where the capillary network is absent. The presence of blood at or near the fovea can lead to obscuration of vision, and so detecting blood in this region is imperative if sight-threatening complications are to be prevented.

The standard method for determining blood flow across the fundus is through fluorescein angiography. The patient is injected with a fluorescent compound. Using a sequence of flash photographs, the propagation of this compound through the fundus is followed and used to observe the capillary network. Defects in the capillary network can be seen by this method and allow haemorrhages which may obscure vision to be detected. Diabetic macular ischaemia can also be seen due to the characteristic loss of blood flow to the macula. Fluorescein angiography is a very effective method, but is also invasive, involving the injection of the fluorescing compound, and is uncomfortable for the patient. In some subjects, the procedure can even be life-threatening. Any alternative method which allows these defects in the capillary network to be seen would be of great potential benefit. In Fig. 10 we show parametric maps of retinal haemoglobins for several subjects, together with the corresponding RGB images.

In Fig. 10(a), there is a clear depression in the quantity of retinal blood at the fovea. This corresponds to the well known foveal avascular zone which is a feature of normal subjects. The additional large depression to the left of the fovea is an artefact due to specular reflection from the inner limiting membrane. This could be removed by acquiring images with crossed polarising filters. In Fig. 10(a), there is a small region in the foveal area where there appears to be a large amount of blood. A comparison with the fluorescein images for this subject indicates that this corresponds to a significant and potentially sight-threatening haemorrhage which requires immediate treatment.

7 Discussion

We have shown how an understanding of the imaging process can allow us to extract histological information from multispectral images of the human ocular fundus. By establishing a mapping between histological parameters and “image quotients”, we are able to establish (subject to certain criteria) the corresponding inverse function which maps image quotients onto histological parameters. The use of image quotients (ratios of image values) allows us to compensate for the inevitable variations in illumination intensity across the fundus. The inverse mapping can therefore be used to compute tissue histology across the ocular fundus, yielding parametric maps which show the spatial distribution of each parameter. We have shown parametric maps for two histological parameters, macular pigment (Fig. 9) and retinal haemoglobins (Fig. 10) which are consistent with known histology. In the case of a known pathology (Fig. 10(b)), the parametric map for retinal haemoglobins is consistent with the fluorescein angiogram for the subject.

Whilst these positive results are very encouraging, there are a number of difficulties with our methods which still have to be resolved. Firstly, we note that not all model parameters can be recovered equally well. On test data, where parameters were known, only macular pigment and retinal haemoglobins could be recovered with acceptable accuracy, as shown in Fig. 5. We believe this to be due to the highly non-linear nature of the inverse mapping which means that construction of the inverse function is very hard. A more exhaustive analysis of inverse function construction is needed.

One of the main drawbacks of the image acquisition system used in this work is that

the multispectral image frames are acquired sequentially. The entire exposure time is of the order of five seconds and a subject cannot remain perfectly still during this time. We have introduced breaks between frames, during which subjects were allowed to close their eyes, but not move their head. They were then asked to re-fixate, and the next frame was acquired. All images can then be acquired without causing undue discomfort, but no matter how well the subject is able to re-fixate, the position of the fundus always changes during these breaks. Ideally, we would capture all images simultaneously. Standard RGB cameras capture three channels (frames) simultaneously, either by interleaving red- green- and blue-sensitive pixels on the same CCD and using interpolation, or by using a separate CCD for each channel. A similar system which simultaneously captures the six frames required in this work would greatly improve the image quality, and would resolve the following issues in the present work.

The first issue that simultaneous image capture would solve is the need for image registration. Whilst the registration methods used (Stewart et al., 2003) have produced excellent results, there is an inevitable loss of fine detail where perfect alignment is required. This would be especially important if we were studying early-stage diabetic retinopathies such as dot haemorrhages which may be only a few pixels in width.

The second issue results from the nature of the cardiac cycle. Since the images are recorded over several seconds, the blood concentration in the fundus is likely to vary between frames. A single “snapshot” would remove this variation.

The final issue is a result of subject movement, but cannot be corrected by post-processing. The use of image quotients to compensate for spatial variations in illuminant relies on those variations being identical in all frames. If this is not the case, then the function $A(\mathbf{x})$ in Eqn. (7) is also a function of wavelength, and may not be taken outside of the integral. Inter-frame movement will necessarily lead to changes in the illuminant. We have noted that the best results have been obtained where only small corrections have been required to register the images. Where large corrections have been necessary, our methods of interpretation fail. The effect of subject movement on illumination is much more serious than the movement itself. The registration methods adopted are able to spatially align images well (although fine detail may be blurred), but they cannot correct for changes in the illuminant. The simultaneous capture of all frames is essential if this difficulty is to be resolved.

A final technical problem that is relatively easy to solve is the removal of specular reflection from the inner limiting membrane. This is a well-known problem in younger subjects and the effects can be seen clearly in Fig. 10(a), where the reflection has caused an artifact in the parametric map. The removal of such reflections is fortunately quite straightforward, requiring only a set of polarising neutral filters. These should be included in any future investigations.

Perhaps the most concerning problem is the need for an empirical scaling factor to ensure that model and image data were coincident. The issues surrounding image acquisition are likely to be contributory factors to this discrepancy, but we can also identify potential problems with the model itself. The individual variation between fundi is likely to be important: we have assumed that the individual layers in the fundus are of constant thickness and are the same in all subjects. This assumption is not correct, and variation in the layer thicknesses may well lead to a significant spread of the data. There may also be variations in the underlying optical properties of the tissue which we have not taken

into account. Whilst the model appears to capture the correct trends, it does not take into account the full variation of normal fundi, and further investigations are necessary in order to determine the precise cause of this discrepancy.

In conclusion, we have demonstrated that a physics-based model of image formation which uses image quotients to compensate for the effects of spatially uneven illumination can be used to extract histological parameters from multispectral images of the ocular fundus. A set of filters was objectively chosen to minimise the error of parameter recovery and a neural network was used to compute tissue histology from image quotients. The method has been successfully used to map the distribution of macular pigment in the retina, and to compute the distribution of haemoglobins in the retina. Results are consistent with known histology and existing imaging methods. Further work is required to improve the image acquisition process to allow all frames to be captured simultaneously, and further refinements to the model are required in order to capture the full variation in fundus histology.

8 Acknowledgements

We gratefully acknowledge financial support from the EPSRC under grant number GR/S09906/01. We thank Mr James Osborne for his generous donation of the Zeiss RCM250 fundus camera used in this work. Mr Peter Lundh of the Institute of Ophthalmology provided technical help with the design of the imaging system. The Optometry department of the Birmingham and Midland Eye Centre kindly loaned us an artificial eye. Miss Stavrou of the Birmingham and Midland Eye Centre assisted with the selection of suitable subjects.

References

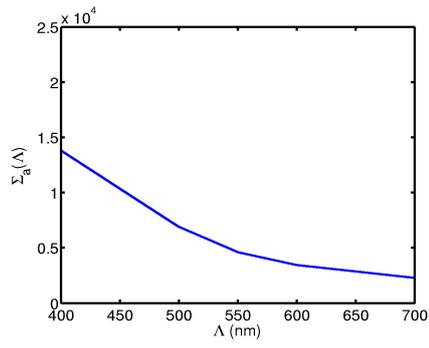
- Alm, A., Bill, A., 1973. Ocular and optic nerve blood flow at normal and increased intraocular pressures in monkeys (*macaca irus*): a study with radioactively labelled microspheres including flow determinations in brain and some other tissues. *Exp. Eye Res.* 15, 15–29.
- Anderson, R., Parrish, J. A., 1981. The optics of human skin. *J. Invest. Dermatol.* 77, 13–19.
- Beatty, S., Boulton, M. E., Henson, D. B., Hui-Hiang, K., Murray, I. J., 1999. Macular pigment and age-related macular degeneration. *Br. J. Ophthalmol.* 83, 867–877.
- Berendschot, T. T. J. M., DeLint, P. J., van Norren, D., 2003. Fundus reflectance - historical and present ideas. *Prog. Retin. Eye Res.* 22, 171–200.
- Bessho, K., Yamaguchi, T., Nakazawa, N., Mihashi, T., Okawa, Y., Maeda, N., Fujikado, T., 2005. Live photoreceptor imaging using a prototype adaptive optics fundus camera: A preliminary result. E-Abstract.
- Cotton, S. D., Claridge, E., Hall, P. N., 1997. Noninvasive skin imaging. In: *Information Processing in Medical Imaging (LNCS1230)*. pp. 501–507.

- Delori, F. C., Pflibsen, K. P., 1989. Spectral reflectance of the ocular fundus. *Appl. Opt.* 28, 1061–1077.
- Hammer, M., Roggan, A., Schweitzer, D., Muller, G., 1995. Optical-properties of ocular fundus tissues - an *in-vitro* study using the double-integrating-sphere technique and inverse monte-carlo simulation. *Phys. Med. Biol.* 40, 963–978.
- Hammer, M., Schweitzer, D., 2002. Quantitative reflection spectroscopy at the human ocular fundus. *Phys. Med. Biol.* 47, 179–191.
- Harvey, A. R., Lawlor, J., McNaught, A. I., Williams, J. W., Fletcher-Holmes, D. W., 2002. Hyperspectral imaging for the detection of retinal diseases. In: *Imaging Spectrometry VIII: Retinal Imaging*. Vol. 4816. pp. 325–335.
- Healey, G., 1989. Using colour for geometry-insensitive segmentation. *J. Opt. Soc. Am. A* 6, 920–937.
- Horecker, B. L., 1942. The absorption spectra of hemoglobin and its derivatives in the visible and near infrared regions. *J. Biol. Chem.* 148, 173–183.
- Kendall, M. G., Stuart, A., 1969. *The Advanced Theory of Statistics: Volume 1*, 3rd Edition. Charles Griffin & Company Limited, London.
- Kubelka, P., Munk, F., 1931. Ein beitrag zur optik der farbanstriche. *Z. Tech. Opt.* 11, 593–611.
- Landrum, J. T., Bone, R. A., Kilburn, M. D., 1997. The macular pigment: a possible role in protection from age-related macular degeneration. *Adv. Pharmacol.* 38, 537–556.
- Maloney, L., Wandell, B., 1986. Color constancy: a method for recovering surface spectral reflectance. *J. Opt. Soc. Am. A* 3, 29–33.
- Moncrieff, M., Cotton, S., Claridge, E., Hall, P., 2002. Spectrophotometric intracutaneous analysis - a new technique for imaging pigmented skin lesions. *British Journal of Dermatology* 146, 448–457.
- Prahl, S. A., Keijzer, M., Jacques, S. L., Welch, A. J., 1989. A monte carlo model of light propagation in tissue. *SPIE Insitute Series IS* 5, 102–111.
- Preece, S. J., Claridge, E., 2002. Monte carlo modelling of the spectral reflectance of the human eye. *Phys. Med. Biol.* 47, 2863–2877.
- Preece, S. J., Claridge, E., 2004. Spectral filter optimisation for the recovery of parameters which describe human skin. *IEEE Trans. PAMI* 26, 913–922.
- Rohen, J. W., 1977. Anatomie und embryologic. In: Francois, J., Hollwich, I. (Eds.), *Augenheilkunde on Klinik und Praxis*. Vol. 1. Thieme, Stuttgart, pp. pp 1.1–1.56.
- Schutz, B. F., 1980. *Geometrical Methods of Mathematical Physics*. Cambridge University Press, Cambridge, UK.

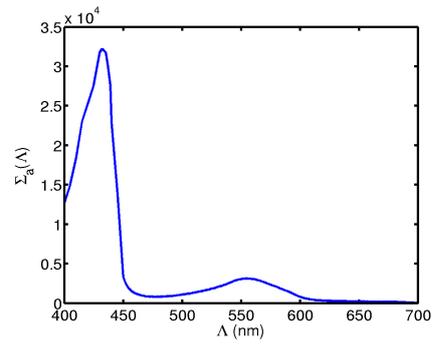
- Stewart, C. V., Tsai, C.-L., Roysam, B., 2003. The dual-bootstrap iterative closest point algorithm with application to retinal image registration. *IEEE Trans. Med. Im.* 22, 1379–1394.
- van Norren, D., Tiemeijer, L. F., 1986. Spectral reflectance of the human eye. *Vision Res.* 26, 313–320.
- Wyszecki, G., Stiles, W. S., 1982. *Color Science: Concepts and Methods, Quantitative Data and Formulae*, 2nd Edition. Wiley, New York.
- Yao, X., Liu, Y., 1997. Fast evolution strategies. *Control and Cyb.* 26, 467–496.

Parameter	Values (mmol/L)
Macular Pigment	0.0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6
Retinal Haemoglobins	0.0, 0.25, 0.50, 0.75, 1.00, 1.25
RPE Melanin	1.64, 3.11, 4.58, 5.56, 6.54, 8.00
Choroidal Haemoglobins	1.90, 3.80, 5.70, 7.60, 8.50
Choroidal Melanin	0.00, 0.33, 0.66, 1.00, 1.33, 1.66

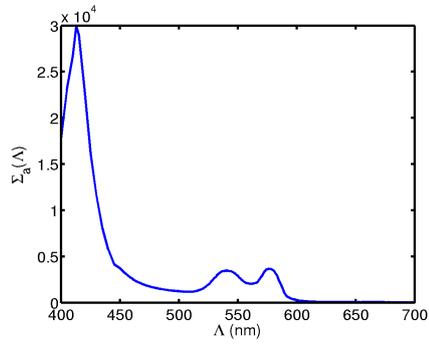
Table 1: Parameter values used in the calculation of f .



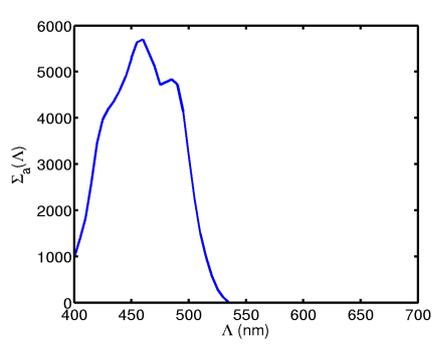
(a) Specific absorption coefficient of melanin (Anderson and Parrish, 1981)



(b) Specific absorption coefficient of haemoglobin (Horecker, 1942)



(c) Specific absorption coefficient of oxyhaemoglobin (Horecker, 1942)



(d) Specific absorption coefficient of macular pigment (Wyszecki and Stiles, 1982)

Figure 1: Specific absorption coefficients (in units of $\text{m}^{-1}/(\text{mmol.L}^{-1})$) of the main absorbers in the fundus.

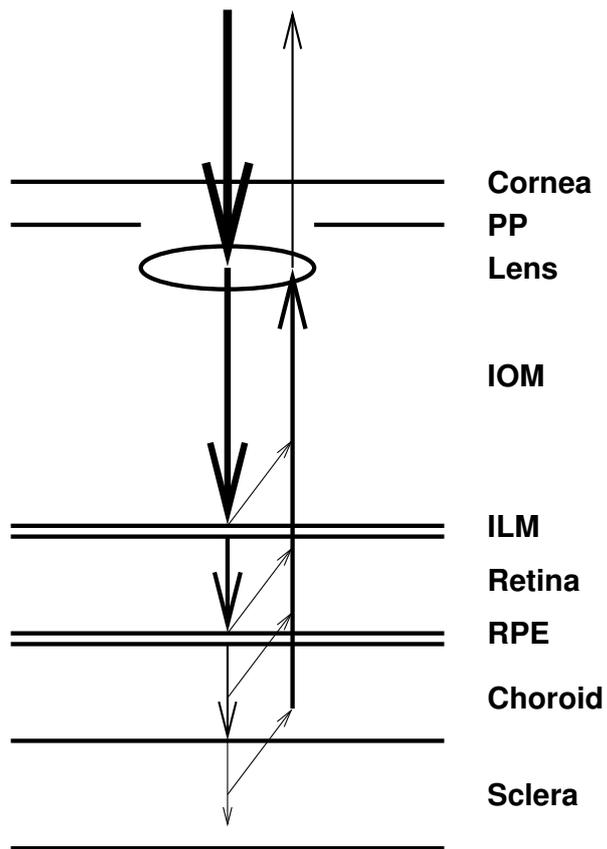
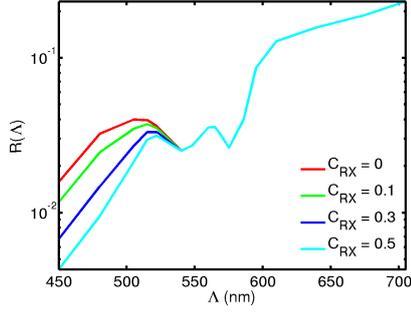
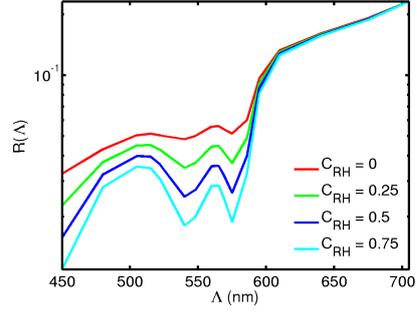


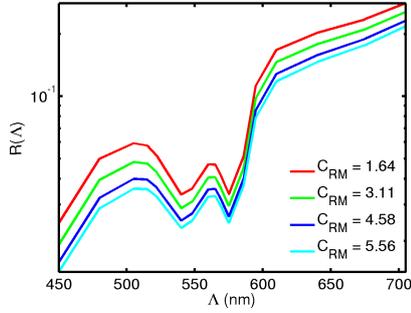
Figure 2: Schematic cross-section of the ocular fundus showing the layered structure of the tissue.



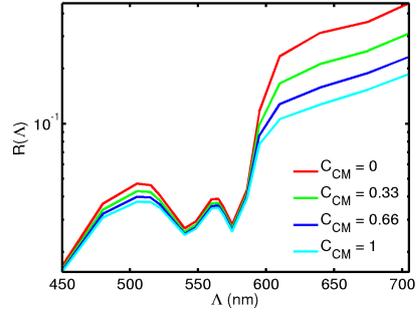
(a) Variation of $R(\mathbf{p}, \lambda)$ with changes in C_{MP} .



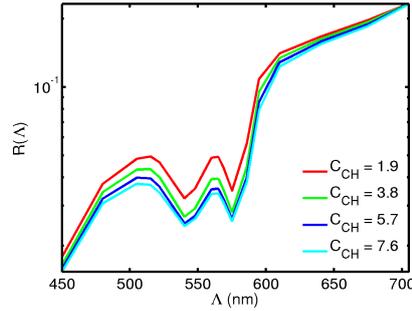
(b) Variation of $R(\mathbf{p}, \lambda)$ with changes in C_{RH} .



(c) Variation of $R(\mathbf{p}, \lambda)$ with changes in C_{RM} .



(d) Variation of $R(\mathbf{p}, \lambda)$ with changes in C_{CM} .



(e) Variation of $R(\mathbf{p}, \lambda)$ with changes in C_{CH} .

Figure 3: Variation of $R(\mathbf{p}, \lambda)$ with respect to changes in (a) Macular Pigment (b) Retinal Haemoglobins (c) RPE Melanin (d) Choroidal Melanin (e) Choroidal Haemoglobins. Non-varying parameters were fixed at: $C_{MP} = 0.0\text{mmol.L}^{-1}$, $C_{RH} = 0.75\text{mmol.L}^{-1}$, $C_{RM} = 4.58\text{mmol.L}^{-1}$, $C_{CM} = 0.66\text{mmol.L}^{-1}$, $C_{CH} = 5.70\text{mmol.L}^{-1}$.

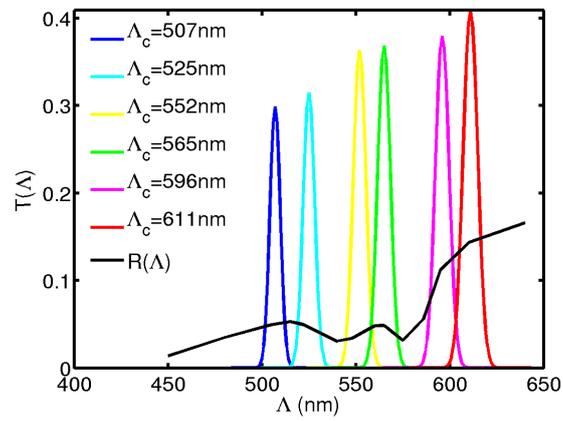
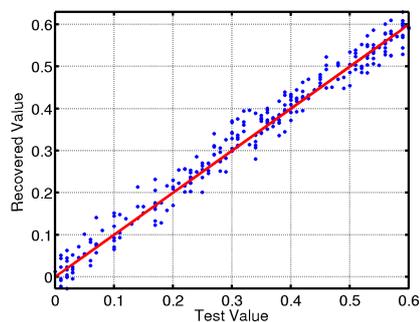
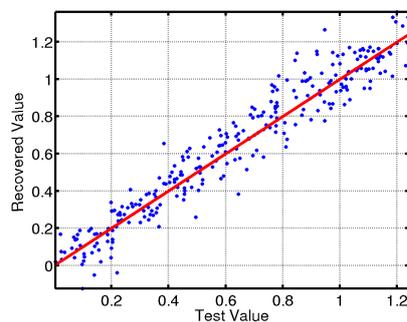


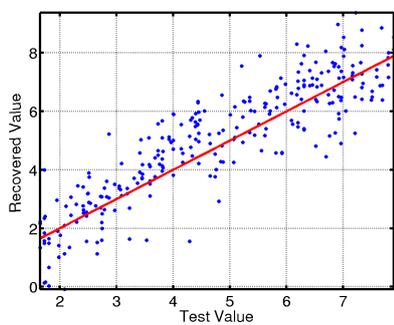
Figure 4: Optimal filters for parameter recovery. A sample reflectance spectrum is shown for reference.



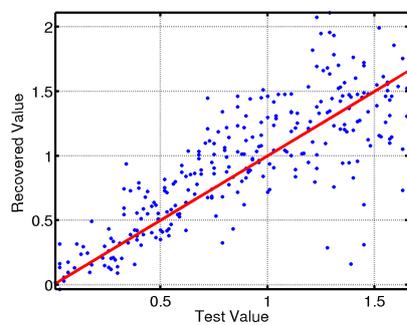
(a) Macular pigment



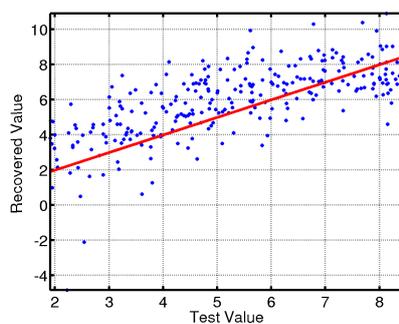
(b) Retinal haemoglobins



(c) RPE melanin

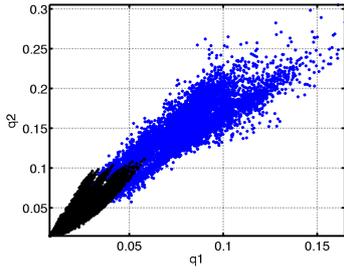


(d) Choroidal melanin

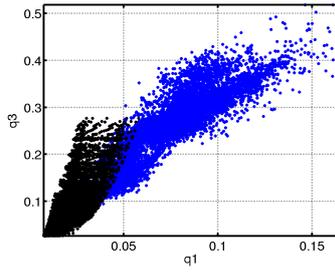


(e) Choroidal haemoglobins

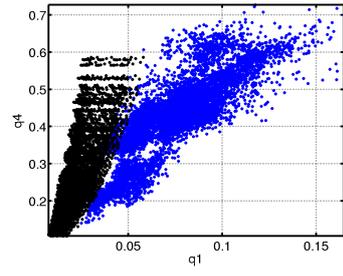
Figure 5: Parameter recovery from generated data. Image quotients were generated from random parameter vectors and a neural network (trained on a different dataset) was used to recover the parameters. The standard error of parameters recovery based on these data is: macular pigment: $\pm 0.13 \text{mmol.L}^{-1}$; retinal haemoglobins: $\pm 0.37 \text{mmol.L}^{-1}$; RPE melanin: $\pm 2.85 \text{mmol.L}^{-1}$; choroidal melanin: $\pm 0.87 \text{mmol.L}^{-1}$; choroidal haemoglobins: $\pm 3.7 \text{mmol.L}^{-1}$.



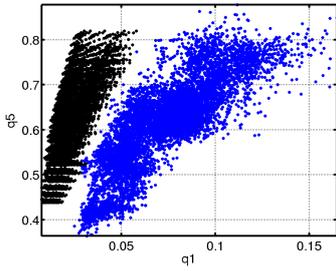
(a) q_1 - q_2 plane



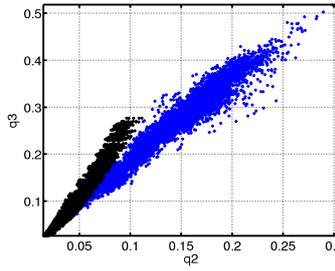
(b) q_1 - q_3 plane



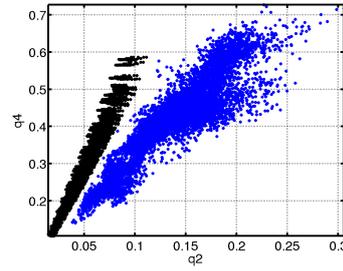
(c) q_1 - q_4 plane



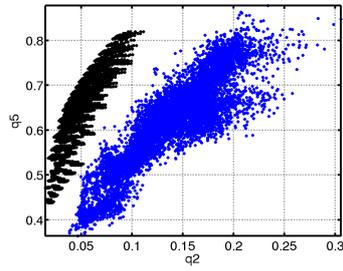
(d) q_1 - q_5 plane



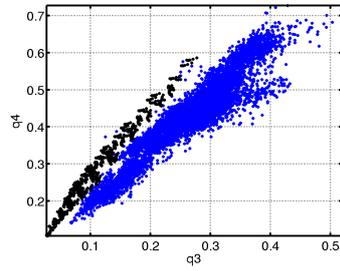
(e) q_2 - q_3 plane



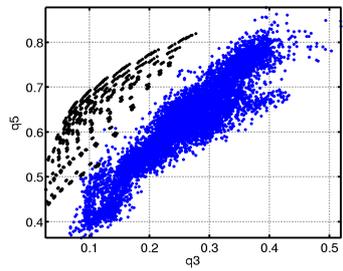
(f) q_2 - q_4 plane



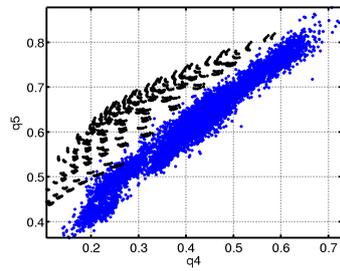
(g) q_2 - q_5 plane



(h) q_3 - q_4 plane



(i) q_3 - q_5 plane



(j) q_4 - q_5 plane

Figure 6: Projections of model (black points) and image (blue points) data onto each of the ten possible q_i - q_j planes.

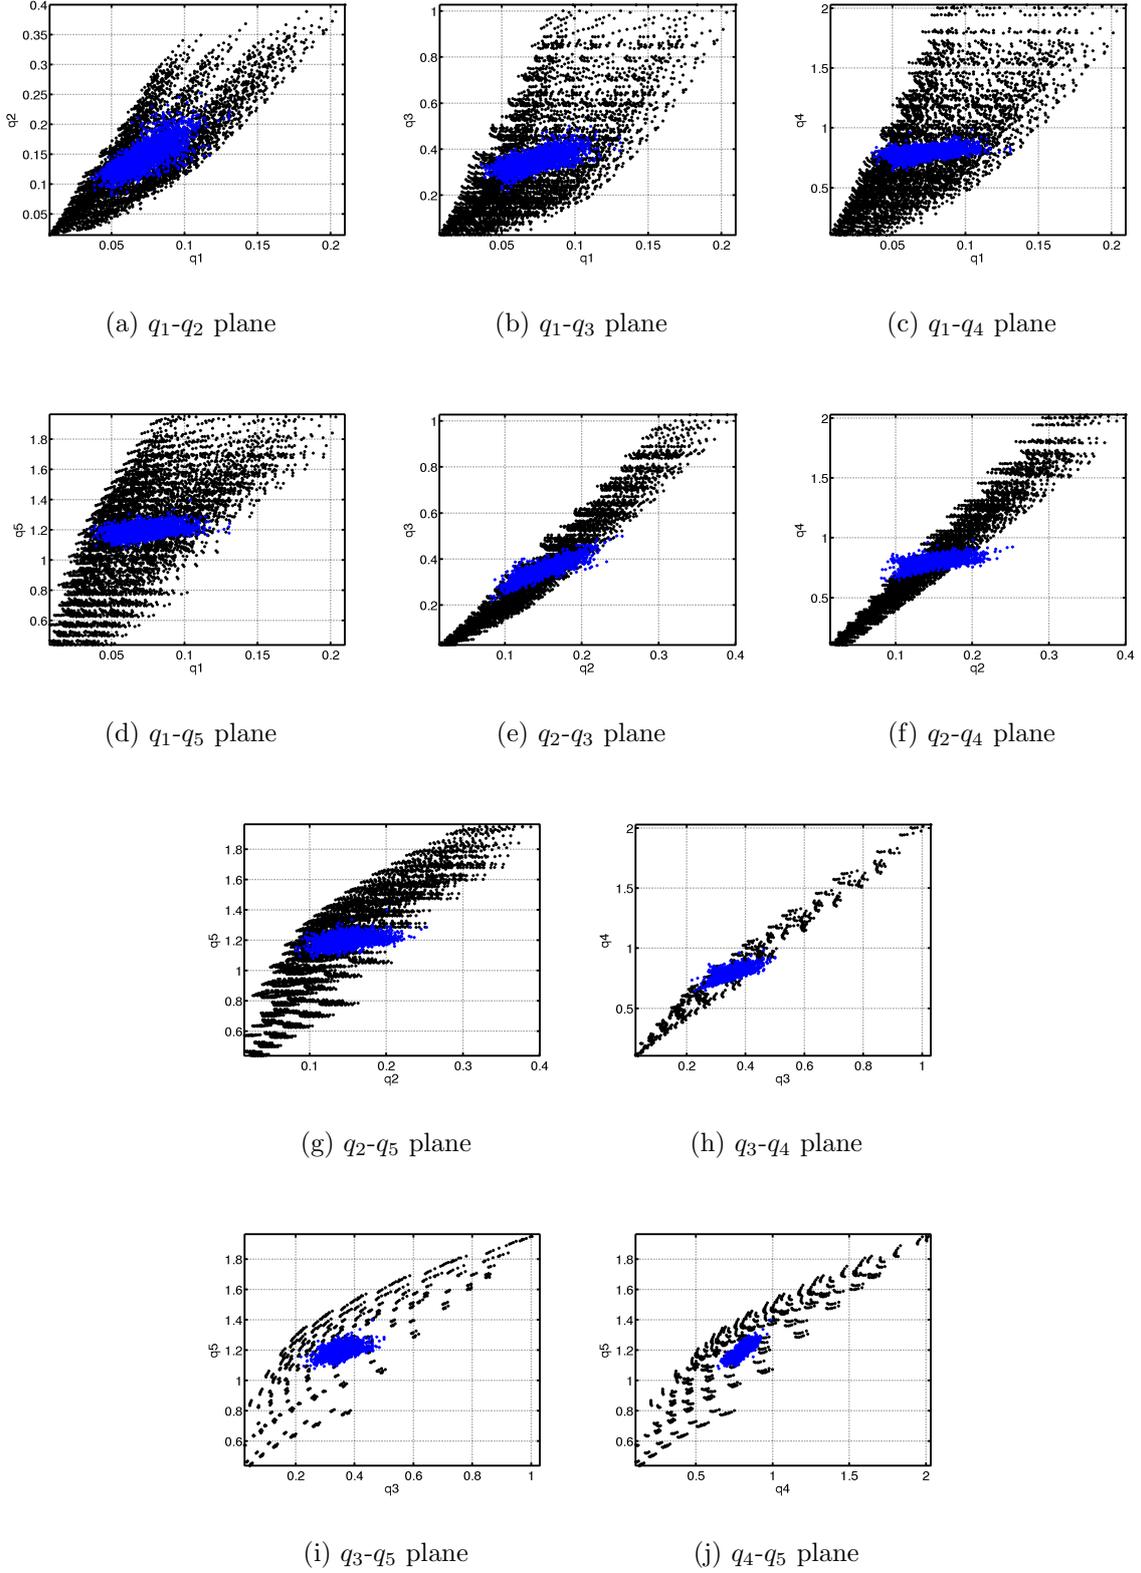
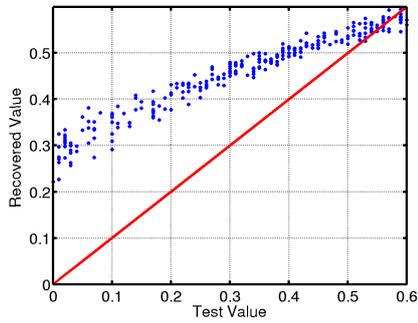
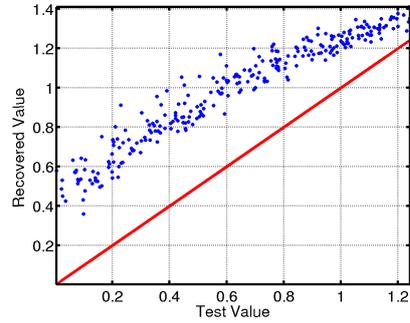


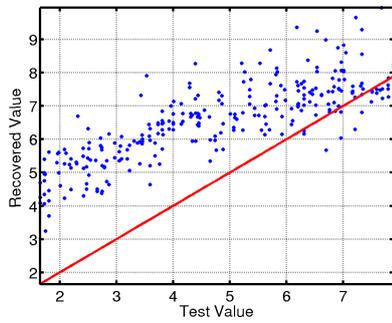
Figure 7: Projections of model (black points) and image (blue points) data onto each of the ten possible q_i - q_j planes. The datasets have been mean-aligned and the model has been scaled by a factor of four.



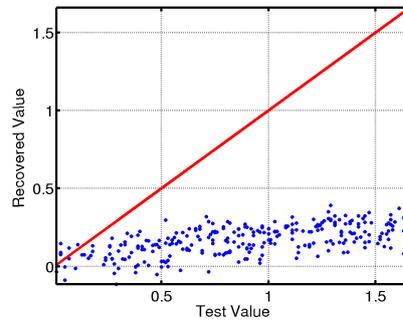
(a) Xanthophyll



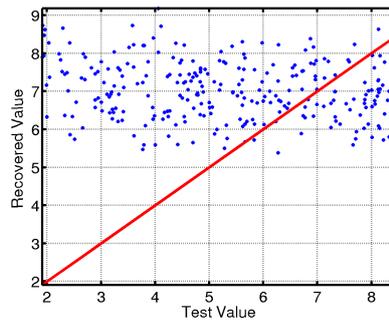
(b) Retinal haemoglobins



(c) RPE melanin

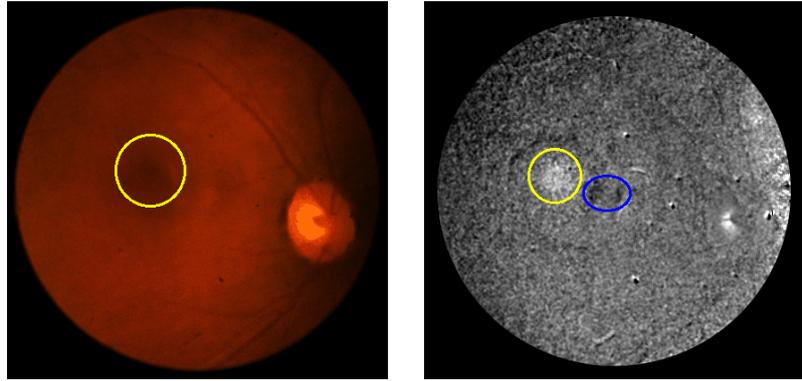


(d) Choroidal melanin

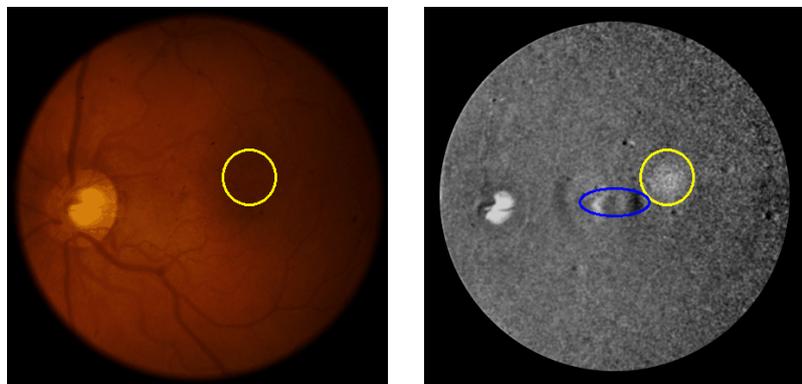


(e) Choroidal haemoglobins

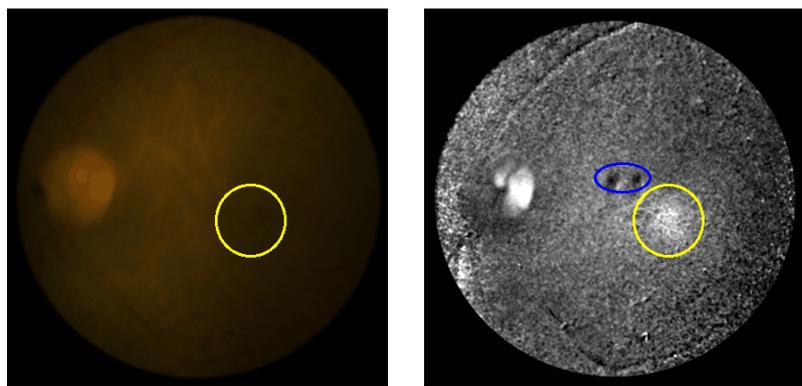
Figure 8: Parameter recovery from generated data. Image quotients were generated from random parameter vectors and scaled as described in the text. A neural network trained on the scaled data was used to recover the parameters. Note well that although the trends for RPE melanin and retinal haemoglobins are correct, the effective range onto which these parameters can be recovered is significantly reduced.



(a) Young Mediterranean female

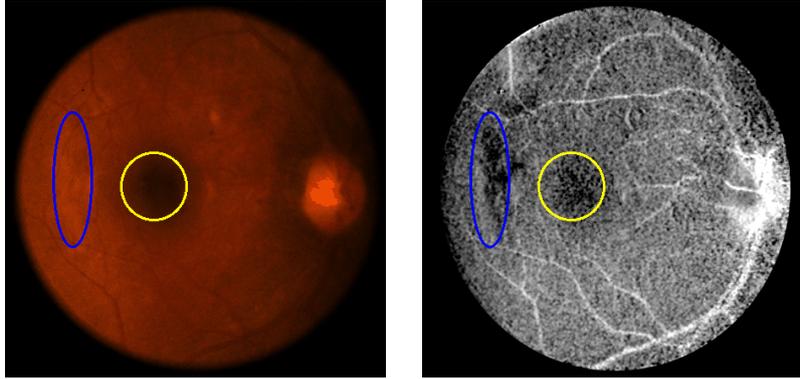


(b) Young Afro-Caribbean male

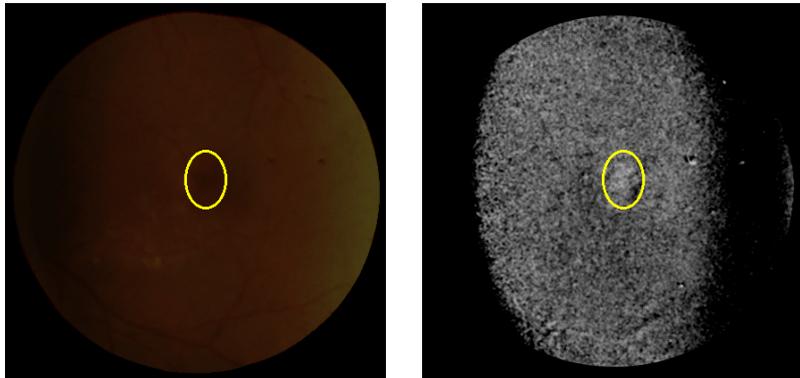


(c) Elderly Afro-Caribbean female

Figure 9: RGB images and the corresponding parametric maps for macular pigment. In the parametric maps, lighter areas correspond to higher concentrations of pigment. The foveal region is circled in yellow. The artifacts at the centres of the parametric maps (blue oval) are due to a reflex from the camera optics.



(a) Young African female. The foveal region is circled in yellow, an artifact due to reflection from the ILM is highlighted in blue.



(b) Elderly Caucasian female showing a foveal haemorrhage (circled in yellow)

Figure 10: RGB images and the corresponding parametric maps for retinal haemoglobins. In the parametric maps, lighter areas correspond to higher concentrations of pigment.