

Durham E-Theses

Understanding complex phenomena in colour and development of a novel laundry metric

MARKIN, ANDREI

How to cite:

MARKIN, ANDREI (2023) *Understanding complex phenomena in colour and development of a novel laundry metric*, Durham theses, Durham University. Available at Durham E-Theses Online:
<http://etheses.dur.ac.uk/14840/>

Use policy

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a [link](#) is made to the metadata record in Durham E-Theses
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Please consult the [full Durham E-Theses policy](#) for further details.

Understanding complex phenomena
in colour and development of a
novel laundry metric

Andrei Markin

PhD Thesis



Submitted for admission to the degree of

DOCTOR OF PHILOSOPHY

Department of Chemistry

Durham University

January 2023

Abstract

Diffuse reflectance spectroscopy (DRS) is a powerful non-contact technique for probing the physical world. In this project it was applied to two main areas of study, colourimetry and chemometrics. Chapter 1 gives a review of the relevant concepts used throughout this study. Chapter 2 first explores the origin of the difference in perceived colour when compounds containing certain lanthanide ions are viewed under sunlight or fluorescent lighting. It was found that salts of neodymium and holmium are subject to the largest change in colour. The cause of this phenomenon was found to be the overlap of the concentrated green component in fluorescent lighting with absorption bands of the lanthanide ions, leading to illuminant metamerism. The phenomenon observed in neodymium chloride was then utilised in order to control the perceived appearance of a sample to any hue using a custom-build spectrally tunable light source, while maintaining near-white illumination. The process was herein termed *tunable illuminant metamerism*.

The colourimetry models which could accurately describe the complex colour appearance of lanthanide salts were then applied to study the workings of fluorescent whitening agents (FWAs) and hueing dyes (HDs) on the colour correction of naturally degraded fabric (yellowing). A platform was developed which could be used to model the appearance fabric under different lighting conditions, at different stages of yellowing while simulating the effects of FWAs and HDs.

In chapter 3, DRS is applied to study the lipid component of laundry stains, and a technique for the quantification of lipid on fabric substrates by short-wave infrared DRS was devised. A unit system for comparing the concentration of lipid on fabric, termed *wt%*, was first discussed and calibration samples of uniformly set lipid concentration were prepared and measured using DRS to yield a dataset for machine learning. A combination of preprocessing techniques, principal components regression and Gaussian process regression models were used to model lipid concentration from DRS spectra. The validity of the best performing model was tested on an external sample set with consistent

results, confirming that the proposed technique is capable of non-contact quantitative measurements of lipid concentration on fabric. The approach was applied to point-scan imaging where the concentration of lipid across a fabric could be quantitatively mapped and results were verified by gravimetry. Further validation of the developed machine learning models is still required as the method of producing calibration samples of low lipid concentration was found inaccurate at low levels of staining. A revised method for producing uniform calibration samples using inkjet printing was investigated to address this issue and found to be feasible.

Contents

1	Introduction	13
1.1	Photophysics and spectroscopy	13
1.2	Spectroradiometry	15
1.2.1	Light Intensity	17
1.2.2	Calibration and Correction	19
1.2.3	Diffuse reflectance spectroscopy	22
1.2.4	Transflectance spectroscopy	26
1.3	Colourimetry	28
1.3.1	Standard observer functions	29
1.3.2	Colour space	32
1.3.3	White point	36
1.4	Detergency	41
1.4.1	Yellowing	42
1.4.2	Fluorescent whitening agents	43
1.4.3	Hueing dyes	44
1.5	Machine learning	45
1.5.1	Principal component analysis (PCA)	45
1.5.2	t-SNE	49
1.5.3	Gaussian process regression	50
1.5.4	Validation	51

<i>CONTENTS</i>	4
1.5.5 Holdout validation	51
1.5.6 <i>k</i> -fold cross validation	52
2 Colourimetry and tuning colour appearance	53
2.1 Introduction	55
2.1.1 Discrepancies in colour appearance	55
2.1.2 The Alexandrite effect	57
2.1.3 Compounds containing rare-earth elements	60
2.1.4 Colorimetry of consumer fabrics	62
2.2 Results and discussion	64
2.2.1 Colourimetry of compounds containing rare-earth elements . . .	64
2.2.2 Tunable illuminant metamerism	82
2.2.3 Modelling the effects of fluorescent whitening agents and hueing dyes on the appearance of fabric	107
2.3 Summary and conclusions	126
2.4 Experimental	130
2.4.1 Rare-earth elements study	130
2.4.2 Tunable LED Light source	130
2.4.3 Modelling fabric appearance	132
3 Towards quantification of lipid residues on fabric	134
3.1 Foreword	135
3.2 Introduction	136
3.2.1 Chemistry of lipid stains	137
3.2.2 Quantitative assessment of laundry detergent performance . . .	140
3.3 Calibration sample set	143
3.3.1 Unit system	143
3.3.2 Model compounds	144
3.3.3 Sample preparation	144

3.3.4	Spectroscopy of lipid and fabric substrates	150
3.4	Modelling	153
3.4.1	Cotton model	153
3.4.2	Multi-fabric model	175
3.5	Model testing	183
3.5.1	External set	183
3.5.2	Depth penetration analysis	185
3.6	Comparison to colourimetric method	188
3.7	Imaging	193
3.8	Low-concentration calibration samples	197
3.9	Summary and conclusion	202
3.10	Experimental	204
3.10.1	Sample preparation	204
3.10.2	Reflectance spectroscopy	204
3.10.3	Imaging	205
3.10.4	Inkjet printing	206
4	Conclusions and future work	207
5	Appendices	212
5.1	Supplementary information	212
5.2	Electronic supplementary information	214

Declaration

The work described in this thesis was carried out at the Department of Chemistry, University of Durham between the years 2016-2022.

This thesis is the sole work of the author, except where acknowledged by reference, and has not been submitted for any other degree.

Statement of copyright

The copyright of this thesis rests with the author. No quotation from it should be published without the author's prior written consent and information derived from it must be acknowledged.

Aims

This study is centred around the applications of diffuse reflectance spectroscopy (DRS) in extracting and interpreting information about objects from their reflectance profile. The manuscript is divided into two main sections of work. These sections aim to:

Firstly, understand and utilise the complex metaphysical phenomenon of lanthanide compounds appearing completely different colours or hues, which has previously been reported in the literature yet lacked a thorough and detailed study to conclude its origin. The study of complex colour phenomena described herein is of importance in testing the accuracy of our pre-existing models of colour perception. These models find application in many industries, and one of which will be prevalent in this study is the ever-changing landscape of lighting. With each new technology offering a novel spectral power distribution that humans now spend using instead of natural illumination from the sun. Developers of lighting technology need to utilise colour appearance model to understand the impact of their products on consumers' colour perception.

The second section of this manuscript is focused on the development of a novel technique capable of quantifying the level of lipid present on a fabric. The units of such a quantity were first established, and the problem is approached using a combination of reflectance spectroscopy and machine learning on a model system. Such a technique would find application in the assessment of laundry detergents performance at cleaning lipids. This is key to development of more effective formulations which can function in a wider variety of situations.

Acknowledgements

As I sit here writing these final words that mark nearly a decade of my life at Durham University, I look back at my time with great fondness and I would like to thank all the people who made it possible.

I would foremost like to thank Andy Beeby for guiding me throughout this academic journey. ‘Welcome to the madhouse’ was the first message I received after accepting the position, and I can now say that it did not take long to realise what that meant; although not in any way I could have ever imagined, it’s hard to have a boring day when your to-do list consists of entries such as: ‘Check phasing on the acoustic levitator’ or ‘Measure a light signal with a microphone’. That is not to mention any of the amazing opportunities I’ve had to work out in *the field*, where I will always remember waving my time-resolved fluorimeter at the remains of a 3,000 year-old mummy.

The adventures of cg011 could fill a volume longer than this manuscript, and I would like to thank all the people who made the office such a fantastic place to be: Louise, Anna, Chris, Alice, Hannah, Jonny, Diana, Laura, Alex, Rob, and Paul (sorry if I’ve missed anyone), it’s been a pleasure working with you all, and I hope that you too will look back at our everyday struggles with fond memories.

To Rob Pal and Matt Kitching, I could never thank you enough for taking chances on me, and giving me opportunities to get involved with some of the most rewarding projects I have ever been a part of. Kathleen Bowes, Lewis MacKenzie, Ed Ward, and the rest of the Pal group, thank you for adopting me into your office, it was a serendipitous adventure every time I walked across the hall.

I’m grateful to Proctor & Gamble for sponsoring this work, and to all the people at the NIC: Joe Copsey, Andrew Moon, Steve Patterson, Silvia Ruscigino, and especially Euan Magennis, thank you for listening to me ramble about greasy fabrics.

To my family, I am eternally grateful for being so supportive throughout this whole process. I can't express how much I appreciated the alpine air which inspired some of this work. A special thanks to Adriano Parisi and Claire Holden for making sure that 'I'm still alive' after each my many MATLAB rabbit holes.

Finally, without you Arwa and all the walks we'd go for, this watermelon simply could never have come together. Now without further delay, let me introduce you to what has been the method behind the 'madness'.

‘Singularity is almost invariably a clue’

Sir Arthur Conan Doyle

List of abbreviations

Abbreviation	Definition
DRS	Diffuse reflectance spectroscopy
SWIR	Short-wave infrared
NIR	Near-infrared
UC	Ultra-violet
NUV	Near Ultra-violet
CCD	Charge-coupled device
ADC	Analogue to digital converter
PTFE	Polytetrafluoroethylene
SPD	Spectral power distribution
KM	Kubelka-Munk
CIE	International commission on illumination
FWA	Fluorescent whitening agent
HD	Hueing dye
ML	Machine learning
PCA	Principal component analysis
PC	Principal component
PCR	Principal components regression
t-SNE	t-distributed stochastic neighbour embedding
CCT	Correlated colour temperature
LED	Light emitting diode
CRI	Colour rendering index
PWM	Pulse width modulation
PCB	Printed circuit board
IC	Integrated circuit
CAD	Computer aided design
PET	polyethylene teraphthalate

SNR	Signal to noise ratio
SNV	Standard normal variate
MSC	Multivariate scatter correction
FT	Fourier transform
RMSE	Root mean squared error of prediction
GPR	Gaussian process regression
LOD	Limit of detection
LOQ	Limit of quantification

Chapter 1

Introduction

This chapter serves to provide relevant background into the general concepts which will be used and referred to throughout this manuscript. More specific background to each of the two main sections of research is provided at the beginning of proceeding chapters.

1.1 Photophysics and spectroscopy

Light or electromagnetic radiation can be considered as an oscillating electric field which generates an orthogonal oscillating magnetic field, and exists in discrete packets called *photons*. The energy (E) carried by photons in such a wave is proportional to the square of its amplitude, as well as the frequency of its oscillation (ν), and can take values in multiples of the Planck constant h :

$$E = h\nu = \frac{hc}{\lambda} = hc\bar{\nu} \quad (1.1)$$

where λ and $\bar{\nu}$ represent the wavelength and wavenumber of the oscillating field respectively. While photons of different frequencies may be similar in nature, the way

they interact with matter leads to a range of different properties for different regions on the electromagnetic spectrum. For this reason, the spectrum has been subdivided into 7 regions, ranging from low-energy radio waves to high-energy gamma-rays and their associated wavelength range can be seen below.

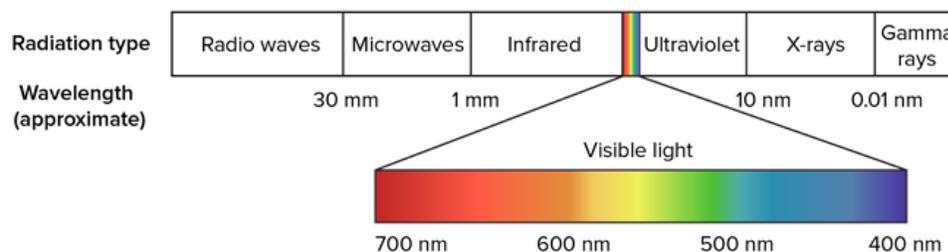


Figure 1.1 Representation of the different sections of the electromagnetic spectrum, with the visible portions approximated in colour.

The electric component of light means that it is able to interact with matter possessing a transition dipole moment, and in cases where the energy of the photon corresponds to the energy between two states in an atom, the photon may be absorbed. The absorption of a photon results in the corresponding atom, ion or molecule being prompted to an excited state, which can have different physical and chemical properties to the ground state. In this work, the main wavelengths of interest are the ones covering: the visible ($\lambda = 400 - 700$ nm) region, and transitions in this energy range typically taking place between electronic states of atoms; and a region between the infrared and red portion of the spectrum, which is referred to as the short-wave-infrared (SWIR: $\lambda = 1000-2500$ nm) region. Transitions in this region are typically of the right energy to excite vibrational overtones and combination modes, which have their fundamental absorption peak in the mid-infrared region at lower energy.

The absorption of light by an atom is followed by relaxation back to its ground state, and that may be radiation-less or radiative. Radiation-less relaxation occurs between vibrational states via *internal conversion*, with energy being transferred to internal molecular modes. Radiative relaxation from an excited state is called luminescence,

and it results with the emission of a photon. The process can be further subdivided into 2 types, fluorescence and phosphorescence depending on whether or not there is a change of spin between the initial and final state; but in general fluorescence is far more common. The presence of radiation-less relaxation means that the process of luminescence is never completely efficient in a bulk material, and the efficiency is defined by the so called *quantum yield*, which is measured as a ratio of the number of photons emitted by a system compared to the number absorbed.

1.2 Spectroradiometry

The measurement of light intensity as a function of frequency yields the optical spectrum, which forms the basis of the experimental presented in this thesis. An overview into the necessary considerations and instrumentation used to do so is therefore provided.

Some of the earliest attempts in quantifying the electromagnetic spectrum used glass prisms as dispersive elements, and the light projected onto a screen where the operator would record the spectra using thermopiles,¹ or with their own visual system. Since then, the invention of more advanced optical sensors such as photodiodes and more versatile dispersive elements has lead to modern spectrometers. Recording optical spectra in different regions of the electromagnetic spectrum is achieved most effectively using a variety of different designs. Most commercial systems can be categorised into two groups: ones which disperse light into it's constituent frequencies before recording the intensity as a function of position, and another which utilises the interference properties of light to effectively record the entire spectrum simultaneously.²

In this work the majority of measurements are carried out in the NUV-SWIR region of the spectrum (350 - 2500 nm), as such the most suitable of the available spectrometers to use were deemed to be the CCD-spectrograph for UV-VIS region, and a Fourier

transform spectrometer for the NIR-SWIR region. The advantages and disadvantages of each system is beyond the scope of this review, as the performance of any spectrometer is typically determined by the implementation rather than the type of optical design. Although, the mechanistic workings of each type of spectrometer used in this work are still useful to consider:

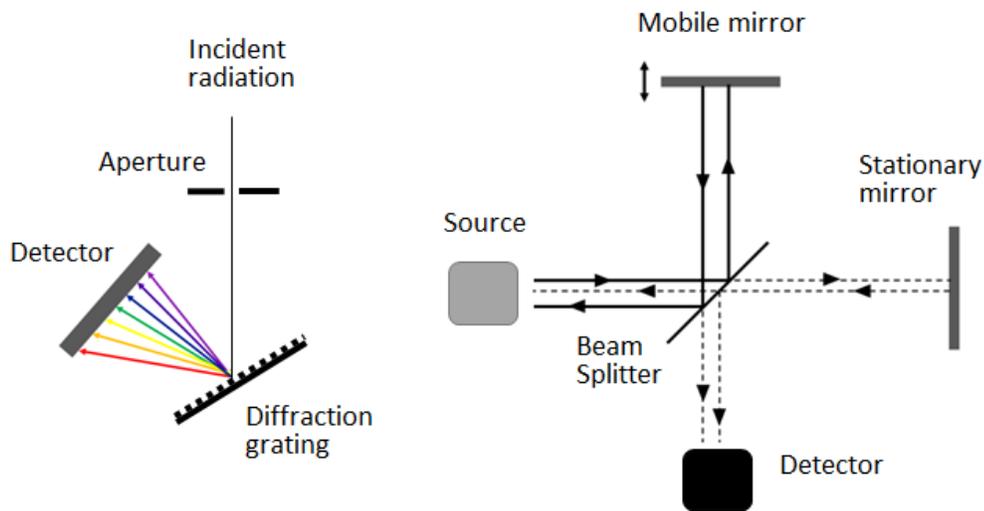


Figure 1.2 Basic schematic of two types of spectrometer used throughout this work: (Left) a dispersive spectrometer design, and (right) Fourier transform spectrometer employed for the SWIR region. Diagrams have been adapted with permission from ref. 3.

CCD spectrometers utilise a dispersive element, diffraction grating (or sometimes prism) and a set of optics which image the diffracted light onto a linear array of photodiodes, which form the detector and is known as the charge coupled device (CCD). A CCD can be used to measure the intensity of light as a function of position in a near-simultaneous manner. By measurement of a light source, a simple polynomial function can be found which relates position on the CCD to a wavelength of light that is incident on it.

Fourier transform spectrometers operate with no dispersive elements, instead light is diverted to two separate paths of varying length then recombined at a single point detector. Varying the length of a single path, typically by specially displacing and recording the intensity observed at the detector. The intensity of each optical frequency component can then be resolved via a Fourier transform. Provided that the difference

in pathlength between the two arms can be accurately measured, which is achieved by a reference laser travelling collinear or along a separate path to the probe beam. The detectors utilised in such designs are typically photodiodes due to their relatively high sensitivity and fast response.⁴

1.2.1 Light intensity

The output of a spectrometer is often quoted in units of ‘intensity’ as a function of frequency, however, there are several common practises used to report the intensity of light. Radiant flux, the intensity of light arriving at the detector per unit time can be represented either: as photonic flux, the number of photons arriving per unit time; or the perceived power, the energy falling on the detector. The interconversion is non-linear and related by equation 1.1. Perceived power is commonly used in radiometry while photonic flux is more useful when relating the spectra to physical properties of atoms and molecules.

When reporting the intensity of an optical spectrum, the units are commonly omitted, however, it is important to distinguish which unit system the intensity is proportional to in order to perform quantitative measurements with optical spectroscopy. In this work all units of intensity are reported proportional to the power of light incident on the detector.⁵

Determining the exact physical quantity reported by a spectrometer is difficult as it involves many non-trivial parameters such as the wavelength dependent efficiency of the optics, any aberrations or stray light effects within the optical assembly, as well as the quantum efficiency of the detector at each frequency. To complicate things further there may be non-linear thermal effects which skew these parameters with changes in temperature. To understand the nature of the signal, we therefore need to trace it back

to its physical source, the detector.

In each of the aforementioned spectrometers, the detector is typically a single point or array of photodiodes. Photodiodes are semiconductor devices with a defined band gap which corresponds to its spectral bandwidth. Incident electromagnetic radiation can cause excitation of electrons from the valence band to the conduction band within the depletion region of the photodiode. The net result is a charge separation which is proportional to the photon flux incident on the depletion region.⁶ Addition of a load across the charge separation leads to a current, referred to as the photocurrent to distinguish it from the current generated by excitation from thermal processes within the semiconductor, which is above absolute zero. The generated photocurrent can be measured using a variety of electronic circuits, for example, a trans-impedance amplifier based on a operational amplifier:

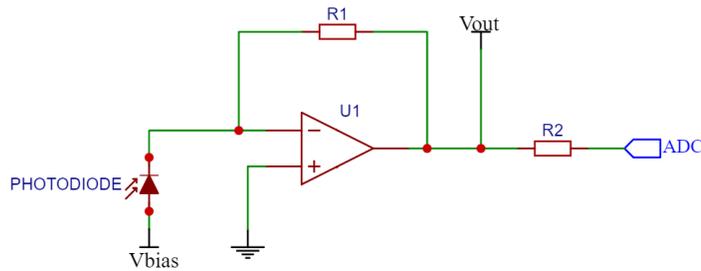


Figure 1.3 Photodiode current to voltage converter circuit. U1: operational amplifier, R1: feedback resistor, V_{bias} : reverse bias voltage, V_{out} : output signal, ADC: analogue to digital converter.

The circuit converts the photocurrent generated by a photodiode into a voltage directly proportional to the feedback resistor (R_1). In this configuration the photodiode is used in photo-conductive mode, meaning it is reverse biased and incident photons allow the diode to conduct in the usually blocked direction.

The wide spread use of this circuit comes from the high linearity with photon flux as well as increased bandwidth due to the decreased parasitic capacitance. In addition

commercially available low-distortion operational amplifiers are now available at low-cost.

The output (V_{out}) is a low impedance source which can be directly sampled by an analogue to digital converter (ADC). An ADC takes an analogue voltage as an input and outputs a digital number between 0 and 2^n , where n is the resolution in ‘bits’. Commercial devices are commonly available with resolution of: 8-bit (0-255), 16-bit(0-65536), and 24-bit (0-16777216). The output from the ADC is usually the quantity displayed as ‘intensity’ and it is therefore proportional to the photon flux at the detector, but lacks correction for sensitivity; a conversion to power units is possible using equation 1.1.

1.2.2 Calibration and correction

In order to have reliable data from a spectrometer the instrument needs to have a reliable calibration and an equally sensitive, linear response across its operating bandwidth. Usually these conditions are specified to within a given tolerance provided by the instrument manufacturer. Most systems, however, require additional correction steps to ensure the output is accurate.

A background correction is required if the detector produces a reading with no input light, such as thermal current in photodiodes or ‘hot pixels’, and this is achieved by recording the output when the detector is isolated from any optical input, and subtracting the recorded reading from subsequent measurements. The correction must be repeated on a time scale similar to the drift of the instrument.

Custom optical systems which introduce optical components in the beam path that are not accounted for by the manufacturer require further correction. The wavelength calibration is usually unaffected by introducing optics external to the spectrometer and

therefore usually does not need to be corrected, but each optical setup may have a complex frequency dependant attenuation. To account for this and achieve a uniform response across the spectrometers operating frequency range, it is therefore necessary to correct the intensity measured by each custom optical setup.

This can be achieved using a commercial calibration source, which has a well characterised output spectrum, and can be traced to reliable standard. The calibration source is guided into the optical setup in such a way as to imitate the sample being measured. A spectrum is then recorded and the intensity compared to the reference spectrum of the standard. A correction factor ($\xi(\lambda)$) is then determined by the ratio of the two spectra:

$$\xi(\lambda) = \frac{I_{cal}(\lambda)}{I_{measured}(\lambda)} \quad (1.2)$$

Where $I_{cal}(\lambda)$ is the calibrated intensity of the source and $I_{measured}(\lambda)$ is the intensity measured by the spectrometer. Provided the detector is operating within its linear region, the correction factor can be applied to any further reference-less measurements made to compensate for any frequency dependent attenuations within the optical setup. The units of corrected intensity become arbitrary due to ratiometry, however, it is important to note which unit system the supplied calibration source intensity follows as it will be incorporated into the correction factor and applied to subsequent measurements.

An incandescent source can also be used as a calibration source as long as it has a known, and stable temperature. The intensity can then be approximated using the Planck equation:

$$I(\lambda, T) = \frac{2hc^2}{\lambda^5} \left(e^{\frac{hc}{\lambda k_B T}} - 1 \right)^{-1} \quad (1.3)$$

where $I(\lambda, T)$ is the emission intensity of a black body at temperature T .⁷ This form of the equation provides the intensity in units $Wst^{-1}m^{-3}$, although a trivial conversion using equation 1.1 is possible if units of photonic flux are required.

In some applications it is useful to visualise spectra on a scale proportional to the energy of photons, rather than their wavelength. In such a scenario the conversion is not as trivial as re-assigning the x axis of the data. The non-linear relationship between wavelength and frequency means that the intensity recorded by the spectrometer is quantised to keep the spacing constant in wavelength units ($d\lambda$), but this results in an uneven spacing in units proportional to energy (dE), such as frequency or wavenumber. Therefore to perform the conversion and preserve the shape of spectra, the spacing needs to be maintained. If the spectrum is considered as a function of wavelength ($f(\lambda)$) or energy ($f(E)$), then the relationship $f(\lambda)d\lambda = f(E)dE$ must be maintained. Substituting with equation 1.1, yields the conversion factor:

$$f(E) = f(\lambda) \frac{d\lambda}{dE} = f(\lambda) \frac{d}{dE} \left(\frac{hc}{E} \right) = -f(\lambda) \frac{hc}{E^2} = -f(\lambda) \frac{\lambda^2}{hc} \quad (1.4)$$

The minus sign represents a change in direction and the factor hc is a constant which does not change the shape of the spectrum, and can be ignored as corrected intensity is usually scaled in ‘arbitrary’ units, in which a constant factor is included. The appropriate conversion from a wavelength quantised spectrum to one in frequency space therefore requires the intensity to be multiplied by the square of its associated wavelength in order to maintain even spacing on an energy scale.⁵

1.2.3 Diffuse reflectance spectroscopy

Measuring spectral properties of solids presents a range of difficulties when compared to transparent samples. Transmissive measurements are usually only possible on fluids or engineered materials with well defined dimensions and uniform properties. Solid samples can be opaque with varying surface texture. To perform spectroscopy on such samples, a useful approach is to illuminate the sample with a ‘probe’ beam of broadband radiation and collect the reflections from the surface, of which there are 2 types; specular and diffuse:

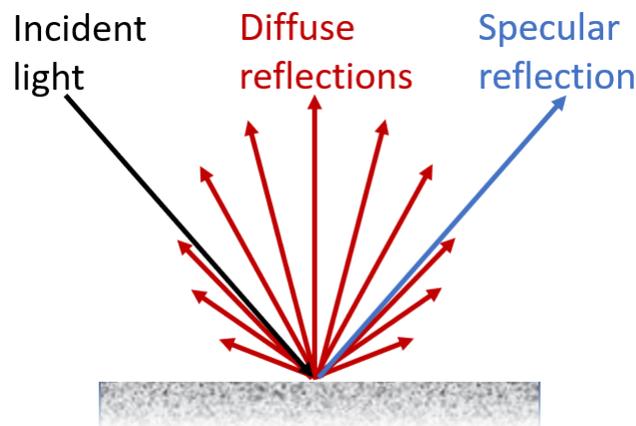


Figure 1.4 Representation of possible paths for light incident on a solid surface producing both dispersed diffuse reflections (red arrows) and mirror-like specular reflections (blue arrows).

Specular reflections result from mirror-like reflections from the top layer of the sample, the beam exits the sample such that the angle between the reflected beam and the normal of the surface is equal to that of the incident beam.

Diffuse reflections occur from light scattering in all directions. This may be caused by reflections from an irregular surface, or the incident light penetrating into the sample and after several reflections the light exits in a random direction. Depending on the sample, absorption can also occur meaning that the spectral profile of the scattered light is attenuated to some degree based on the dielectric properties of the sample. An ideally diffuse surface may be considered as exhibiting a Lambertian reflectance,⁸ where

the intensity of scattered radiation is directly proportional to the cosine of the angle between the incident beam and the normal to the surface, this has been demonstrated in figure 1.4 by the length of the red arrows.

Diffuse reflectance spectroscopy (DRS) aims to produce a reflectance spectrum that is independent of the light source and detector used to measure it. This is achieved by illuminating the sample with a light source which is continuous over the entire spectral region of interest. The diffuse reflections are then collected, collimated and input into a spectrometer. To account for the non-uniform spectral power distribution of the light source, as well as any losses in the collection optics, diffuse reflectance is typically reported relative to a reference material.

In practise, DRS is carried out by recording the spectral power distribution of reflected light from the sample and reference material under the same conditions. The relative reflectance (R) of sample is then defined as the ratio of intensities ($I(\lambda)$) between the two readings:

$$R(\lambda) = \frac{I_s(\lambda)}{I_r(\lambda)} \quad (1.5)$$

Where $I_s(\lambda)$ and $I_r(\lambda)$ represent the intensity of diffuse reflected light measured from the sample and reference respectively.

A common reference material for UV, visible and near infra-red (NIR) reflectance spectroscopy is a sintered PTFE material called Spectralon[®]. It is used as it has the highest diffuse reflectance of any known material in UV-NIR, and behaves as a Lambertian surface between 257 nm to 10,600 nm.⁹

In DRS, it is difficult to directly extract quantitative metrics such as extinction coefficient. This is because the spectrum is heavily influenced by the surface of the sample, as well as other physical properties. A major problem is that it's nearly impossible to

tell how much of the sample, on average, the probe beam interacts with before reaching the detector, which is referred to as the effective path length. Furthermore, it is usually impossible to completely decouple the specular reflection and the diffuse reflection. Since the surface of the sample has a major influence on its reflectance spectrum, the sample preparation, which directly affects the particle size and shape, can have an extremely significant effect on the measured reflectance spectrum.

Due to the geometry of the optics, any luminescence caused by the probe beam will also be coupled with the collected diffuse reflections. This is particularly problematic as the intensity of luminescence depends not only on the sample but also the intensity and spectral power distribution (SPD) of the probe light source, meaning that the reflectance spectrum is no longer device independent. In such a scenario a single sample measured with different instruments will likely give a different reflectance spectra, leading to irreproducible results. Luminescence contribution can be reduced by utilising a spectrometer in which the probe light source is monochromatic and reflectance is measured at a single wavelength, however this requires scanning through the full spectrum, which is usually time consuming and limits sample throughput. In cases where this is not possible, luminescence can be reduced by filtering out the higher frequency components of the probe light source which causes excitation in the sample, although this limits the spectral bandwidth of the measurement.

To make matters even more complicated, the effective path length depends on the refractive index of a sample, which itself is a function of the absorption coefficient, and undergoes a significant change during a resonant transition. This is an example of a non-linear optical property, which can make interpretation of spectra extremely difficult. Mathematical pre-treatment of data, such as Kramers-Kronig relations,¹⁰ can be used to minimise this effect. However, due to the complex nature of reflections, it is only an approximation.

Quantification is further complicated when impurities are introduced. If chemicals of

dissimilar refractive index are mixed together then the changes observed in the reflectance spectrum are not only due to the different dielectric properties of each component like in transmission spectroscopy, but also the change in effective path length causes a distortion to the spectrum which is difficult to predict and compensate for.

Combination of all these effects means that quantification using DRS is not a trivial task and simple Beer-Lambert type analysis of data is impossible. However, pre-processing of reflectance spectra can be used to compensate for these factors to some extent and allow quantification with an acceptable degree of accuracy.

One model which is commonly used in reflectance spectroscopy is the Kubelka-Munk (KM) model. The KM model assigns variables to account for the scattering and absorption properties of the material, $k(\lambda)$ and $s(\lambda)$ respectively. The sample can then be divided into vertical layers, light penetrating the sample then interacts with each layer. The probability of light exiting the material and being detected is then calculated based on how many layers it penetrates, and the sum of these probabilities yields the KM function.

In the special case that the material is approximated to be infinitely thick and that light doesn't penetrate the full depth, among other approximation, the model can be written as:

$$F(R(\lambda)) = \frac{(1 - R(\lambda))^2}{2R(\lambda)} = \frac{k(\lambda)}{s(\lambda)} \quad (1.6)$$

By approximating that the scattering coefficient ($s(\lambda)$) is constant, the KM function ($F(R)(\lambda)$) is then proportional to the coefficient of absorbance. A property which makes it useful for identifying the centres of resonant transitions as well as quantitative measurements, since the absorption coefficient is proportional to the concentration of absorbing species.

1.2.4 Transflectance spectroscopy

Reflectance spectroscopy can be combined with transmission spectroscopy in a technique called *transflectance spectroscopy*, which can be utilised to characterise the optical properties of thin, translucent materials or even liquids.^{11,12} Transflectance spectroscopy is carried out in a similar method to reflectance spectroscopy: a probe beam is focussed on a reflectance standard, followed by collection and measurement of the reflected light ($I_r(\lambda)$); the sample is then placed in front of the reflectance standard and the measurement repeated ($I_s(\lambda)$) to yield a ratio of the two collected spectra (equation 1.5), otherwise called a transflectance spectrum (usually also denoted $R(\lambda)$). The geometry of the optics may be set to collect mainly diffuse reflections or specular reflections as required (figure 1.5).¹³

Transflectance measurements are subject to the same challenges which make quantification of absorbing species difficult in reflectance spectroscopy. However, the technique is usually employed due to the relaxed requirements for sample preparation, as the sample does not need to be of uniform thickness, or of small enough physical dimensions to fit inside a transmission measurement cavity.

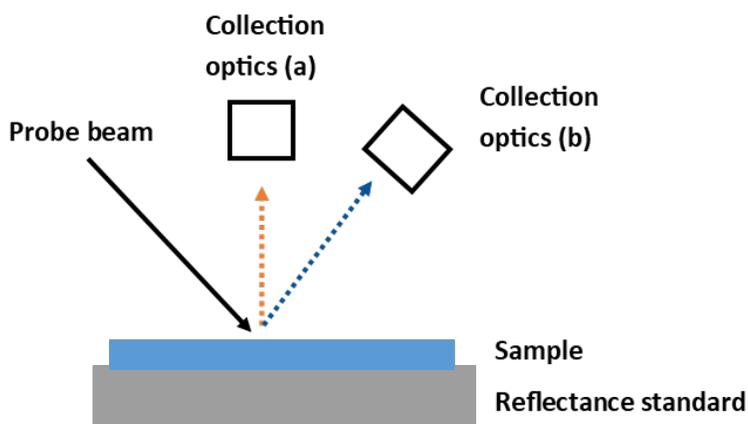


Figure 1.5 Diagram of transflectance spectroscopy setup with geometry of optics optimised for collection of diffuse reflections (a), and specular reflections (b).

Spectralon reflectance standards are not usually used for these measurements as the standard would be contaminated by the sample. A similar inexpensive standard may

be used as long as the measured spectrum is corrected for. Switching which material a reflectance or transmittance spectrum is referenced to may be done by measuring the new standard. The ratio of the two standards then forms a correction factor, which may be used to multiply a measured transmittance spectrum so that it may be comparable to measurements referenced to the Spectralon standard.

1.3 Colourimetry

The human visual system can distinguish around 100,000 - 7 million different colours.¹ From an evolutionary point of view, this ability has allowed us to find shelter and determine whether fruit was ripe enough to eat.¹⁴ Colour is often described as a photo-physical property, simply the spectral density of the light entering our eyes. However, this is not the case. The eye is complex, it contains mechanisms for image adjustment based on it's surroundings, one example which must be considered for colourimetry is 'colour constancy', which will be further explored.

Understanding the origin of colour has therefore puzzled a number of great scientists in the past such as: Newton, Helmholtz, Young and many others.¹⁵

From a spectroscopy perspective, the human visual system can be considered as a form of multi-spectral imaging, allowing us to gather basic information about the reflectance profile of objects in the visible region of the electromagnetic spectrum. It is this reflectance profile which contains useful information about the object, such as the 'ripeness' of fruit.

Modern instruments are capable of extracting vastly more information from objects than the human visual system. That may be through imaging systems which can easily surpass the spatial resolution, or spectrometers which can detect and quantify frequencies of light beyond either extreme of the spectrum detected by the human eye. The science of colourimetry has even allowed us to use such instruments in quantifying the metaphysical experience which we referred to as 'colour'. In the modern world, colour vision is arguably much less of a survival necessity than it is an additional sense for deciding what we like or dislike. The industrial applications of this are far-reaching, some industries are primarily based on accurate colour measurement and reproduction, with examples including: displays,¹⁶ imaging¹⁷ and paints.¹ It is for these reasons that a great deal of research has been carried out on our perception of colour and how it may be influenced.

1.3.1 Standard observer functions

It has been postulated as far back as 1802 that our eyes contain three distinct types of photoreceptors (Young-Helmholtz theory),¹⁵ and this is still the backbone of modern colourimetry.

The response of the human visual system to stimuli from the visible portion of the spectrum may be modelled by the International Commission on Illumination (CIE) system of colourimetry, which constitutes a wide range of formulas and algebraic functions for different metrics that will be introduced throughout this work.¹ The CIE developed their models to represent a ‘standard observer’, derived from the average of a series of psychological experiments which also included people with abnormal vision. Colourimetry systems developed by the CIE have since become widely used throughout industry and academic research for quantitatively describing colour appearance.¹

Colour appearance models require a base set of variables, in the CIE system these are defined as CIE X, Y, and Z. The variables are calculated using a set of numeric functions which could be applied on a SPD of visual stimulus (I_{stim}) to the human visual system within a small 1-4° field of vision.

The functions are referred to as *standard observer function* or similarly *colour matching function*, denoted \bar{x} , \bar{y} and \bar{z} , and they were provided by CIE in the form of a look-up table which has been plotted in figure 1.6. The output is a set of coordinates which, in-combination with a colour space (discussed in the following section), map the appearance of the visual stimuli being analysed.

The CIE 1931 colour matching functions are still widely used to this day, however, they have since been expanded to include corrections for numerous inconsistency found

through experiment; as well as more specific cases such as a wider field of vision.¹⁸

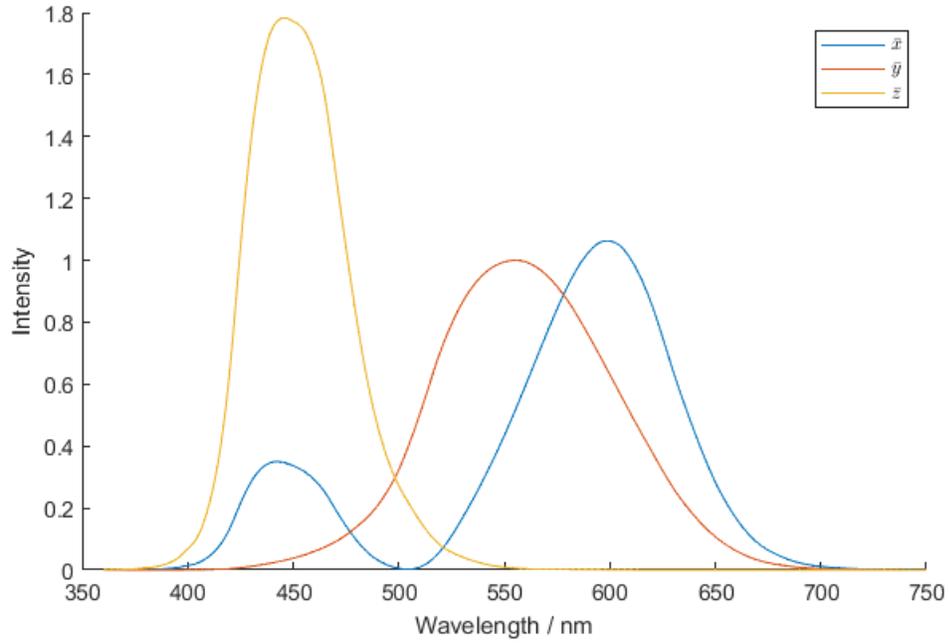


Figure 1.6 CIE 1931 2° standard observer functions used for quantifying human colour perception from the SPD of visual stimuli. The intensities are relative to the \bar{y} function which has been normalised, as defined by CIE.¹⁸

The standard observer functions are applied to calculate the base CIE X, Y, and Z variables by multiplying the SPD of visual stimulus (I_{stim}) with each function in-turn, followed by integration of each resulting sum over the visible region (equation 1.7). The SPD of visual stimulus can be measured directly if the sample under analysis is purely an emitter of light, such as an illuminant. For objects which get their colouration from absorbing species (majority of objects), a product of the DRS spectrum and the SPD of the illuminant ($R \cdot I$) is typically used in-place.

$$\begin{aligned}
 X &= \int_{380nm}^{780nm} I_{stim} \cdot \bar{x} \, d\lambda \\
 Y &= \int_{380nm}^{780nm} I_{stim} \cdot \bar{y} \, d\lambda \\
 Z &= \int_{380nm}^{780nm} I_{stim} \cdot \bar{z} \, d\lambda
 \end{aligned} \tag{1.7}$$

The physical origin behind the spectral shape of the colour matching functions, which were developed from psychology trials, became apparent in 1983, when the photoreceptors in human eyes were, for the first time isolated and their absorbance spectra measured using micro-spectrophotometry.¹⁹

The eye contains 2 types of photosensitive cells named *rods* and *cones*. Rod cells outnumber cone cells by a factor of 20, they contain the photosensitive carotenoid called rhodopsin, and are used in low light conditions. At high levels of illuminations, rhodopsin becomes bleached and our vision is dominated by cone cells.¹ There is also an intermediate level of illumination, in which the visual system operates using a mixture of stimuli from rod and cone cells.

Cone cells contain a different photosensitive pigment called iodopsin, which is far less prone to photo-bleaching. They can also be further subdivided into 3 types, classified as: long, medium and short (or L, M and S). These cells each have various responses to stimuli in the blue, green and red region of the spectrum, reaching peak absorbance at 420, 534, and 564 nm respectively as characterised through absorbance spectroscopy (figure 1.7).²⁰ The absorbance spectrum of cone cells may also be used to measure colour, instead of the standard observer functions, and this is known as the LMS colourimetry system. However, this is far less common than standard observer functions, although both can provide accurate results.¹

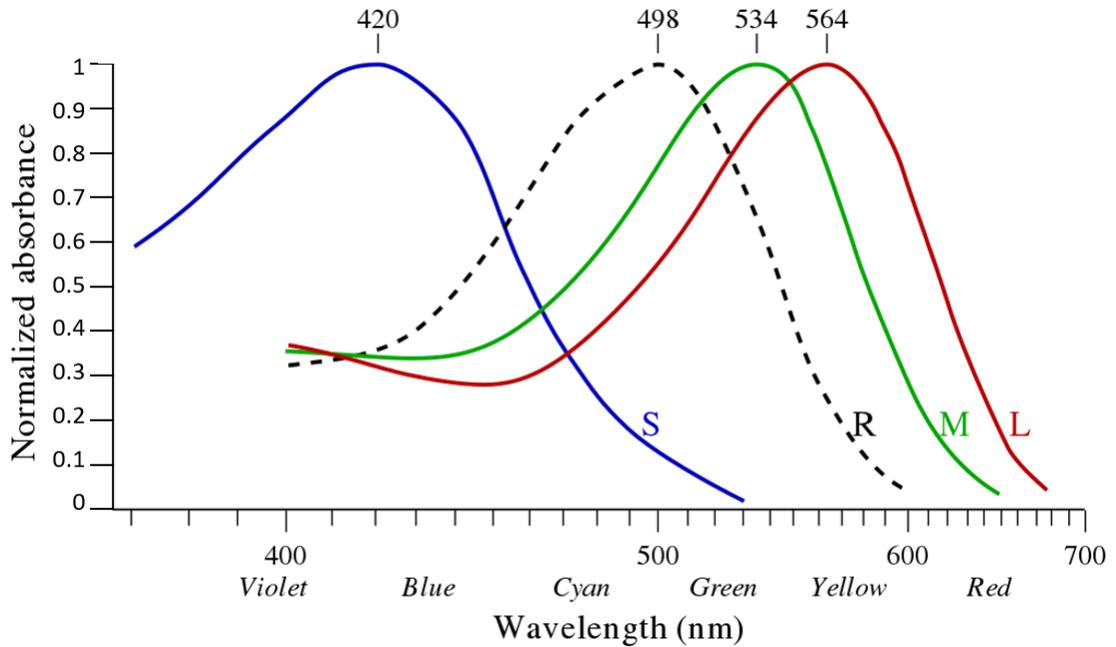


Figure 1.7 Normalised absorption spectra of human rod (R) and cone (L,M and S) cells, with peak maxima labelled as well as the approximate colour of the wavelength scale.

The subdivision of cone cells is the reason human beings have trichromatic vision. The response of these cells is similar, but noticeably different to the psychologically-derived CIE XYZ colour matching functions, and this dissimilarity shows that there is more to colour perception than simply the sensitivity of the photoreceptors. This is the result of biological ‘signal processing’ that takes place on the interconnections between the photoreceptors and the brain, as well as inside the visual cortex of the brain.¹ This has a significant input into our perception of colour and forms a major area of psychology research to this day.²¹

1.3.2 Colour space

To describe colour using the CIE XYZ parameters, they first need to be mapped onto a coordinate system where each point represents a unique colour, called a *colour space*. The coordinates are calculated using transformation functions which act on CIE XYZ

parameters. There are many colour spaces which have been created for specific purposes, for example, the sRGB colour space is widely used in displays and digital photography.²² The ones which are most relevant to this work are the CIE xy and CIE L*a*b* colour spaces.

The CIE xy colour space is a two dimensional plane in which the coordinates are derived directly from the base CIE XYZ values; The X, Y and Z parameters are normalised such that $X + Y + Z = 1$. This means that only two of the parameters need to be specified, as each parameter is a function of all three. The X and Y parameters are then directly used for the x and y axis which forms to the CIE xy colour space and can be seen in figure 1.8. Generally, X increases with an increasingly dominant red component of visual stimuli, the same can be seen for Y and green, as well as Z and blue dominant components. Including the effects of normalisation, the CIE xy colour space described blue hues with low values of x and y, green hues with high y and low x, and red hues with high values of x and an intermediate range of y. The area within the body of the CIE xy colour space defines all colours perceivable by humans with average colour vision.¹

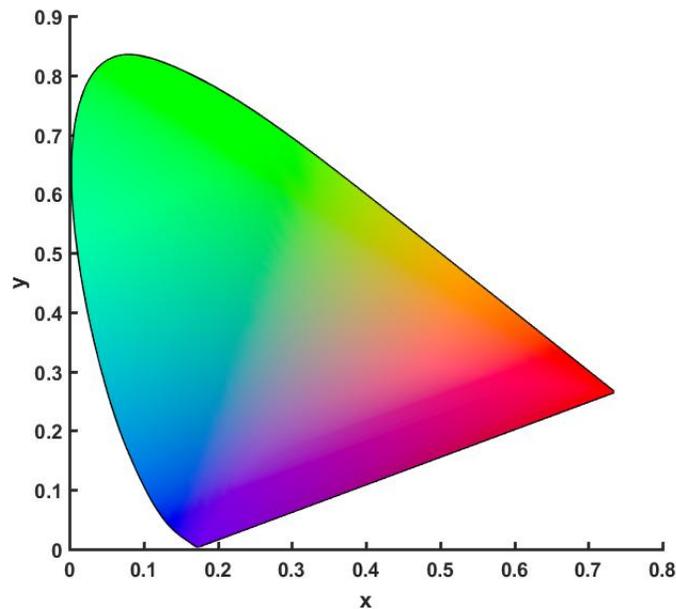


Figure 1.8 CIE 1931 xy colour space with the representative colours of the coordinates translated into the sRGB colour space.

The xy colour space is primitive as it only accounts for chromaticity of a perceived colour, completely ignoring the luminosity or perceived brightness. However, this is not an issue for objects which only emit light, therefore it is still widely used to describe the colour of illuminants. It's continued popularity is potentially due to it's simplicity, as it is 2-dimensional it is simple to display, therefore useful for human assessment of colour.

This colour model was later refined by Munsell,²³ who introduced the separation between chromaticity and luminosity in a 3-dimensional colour space. The theory has since been adapted by CIE who published their hugely popular version named the $L^*a^*b^*$ colour space, with the '*' used to distinguish the coordinates from other colour spaces using the same notation.¹

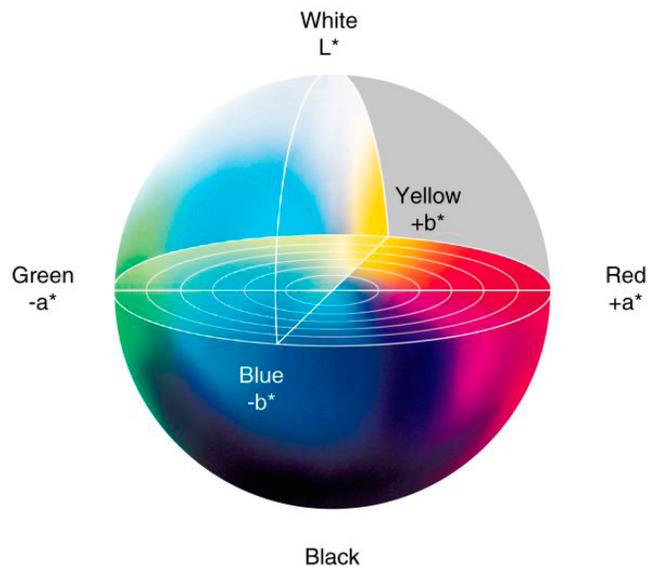


Figure 1.9 Cross-sectional representation of the 3-dimensional CIE $L^*a^*b^*$ colour space at the 50% neutral luminosity.

The aim of the $L^*a^*b^*$ colour space is to represent a perceptually linear coordinate system, meaning that Euclidean distances between colours or coordinates are far more tangible to visual perception and lend themselves to analysis and manipulation. To convert from the CIE XYZ colour space coordinate system to the $L^*a^*b^*$ colour space,

the following functions are used:

$$\begin{aligned}
 L^* &= 116f\left(\frac{Y}{Y_n}\right) - 16 \\
 a^* &= 500\left(f\left(\frac{X}{X_n}\right) - f\left(\frac{Y}{Y_n}\right)\right) \\
 b^* &= 200\left(f\left(\frac{Y}{Y_n}\right) - f\left(\frac{Z}{Z_n}\right)\right) \\
 f(t) &= \sqrt[3]{t} \text{ for } t > \left(\frac{6}{29}\right)^3 \\
 f(t) &= \frac{t}{3\left(\frac{6}{29}\right)^2} + \frac{4}{29} \text{ for } t < \left(\frac{6}{29}\right)^3
 \end{aligned} \tag{1.8}$$

where X_n , Y_n and Z_n specify the CIE XYZ coordinates of the light source used to illuminate the sample, and t the ratio X/X_n , Y/Y_n or Z/Z_n input into function $f(t)$. This sample-illuminant scenario is setup because the $L^*a^*b^*$ colour space is only used to describe the appearance of objects as opposed to illuminants, and the spectral profile of light reflected from an object is a function of the illuminant. In human colour vision, however, there are multiple mechanisms for altering the perception of visual stimulate based on external information. The impact of these mechanisms to colour perception, on a broader level, can be modelled to good level of accuracy; while the mechanisms which act on a local level in a complex scene are still difficult to model.²⁴

A major mechanism in colour perception on a broad-scale is called *colour constancy*. It is the focus of the following section, but it should be noted that an approximation of this mechanism is employed in the $L^*a^*b^*$ colour space. The colour of the illuminant must also be analysed using standard observer functions, and the outcome is referred to as the ‘white point’ in colourimetry.

1.3.3 White point

The human visual system is capable of adapting colour perception to surrounding environment of the observer. On a broad level the mechanism is called colour constancy, and is thought to function by globally adjusting colour appearance so that objects appear consistent colours under changes in the colour of the illuminant. Such changes in the illuminant can occur naturally, for example when the sun starts off an orange colour at sunrise, gradually changing to its brightest blue-white appearance during the middle of the day, followed by the reverse at sunset. The current state of artificial lighting, however, means that indoor illumination is typically a random mix of blue-white (*cool-white*) and yellow-white (*warm-white*) emitters.

The mechanism for colour constancy is still not well understood,²⁵ and recently even the existence of colour constancy has been called into question.²⁶ To understand the working of colour constancy and why this topic is currently so contentious, we can consider a set of objects (white, blue, red and purple) whose reflectance spectra resemble those in figure 1.10 (column 1), and these objects are viewed under daylight and incandescent lighting. The light reflected from the object and seen by the observer is calculated as the product of the source's SPD and the reflectance profile of the object. The resulting spectral intensity reaching the observer under each type of illumination can be seen in figure 1.10 (columns 2 and 3).

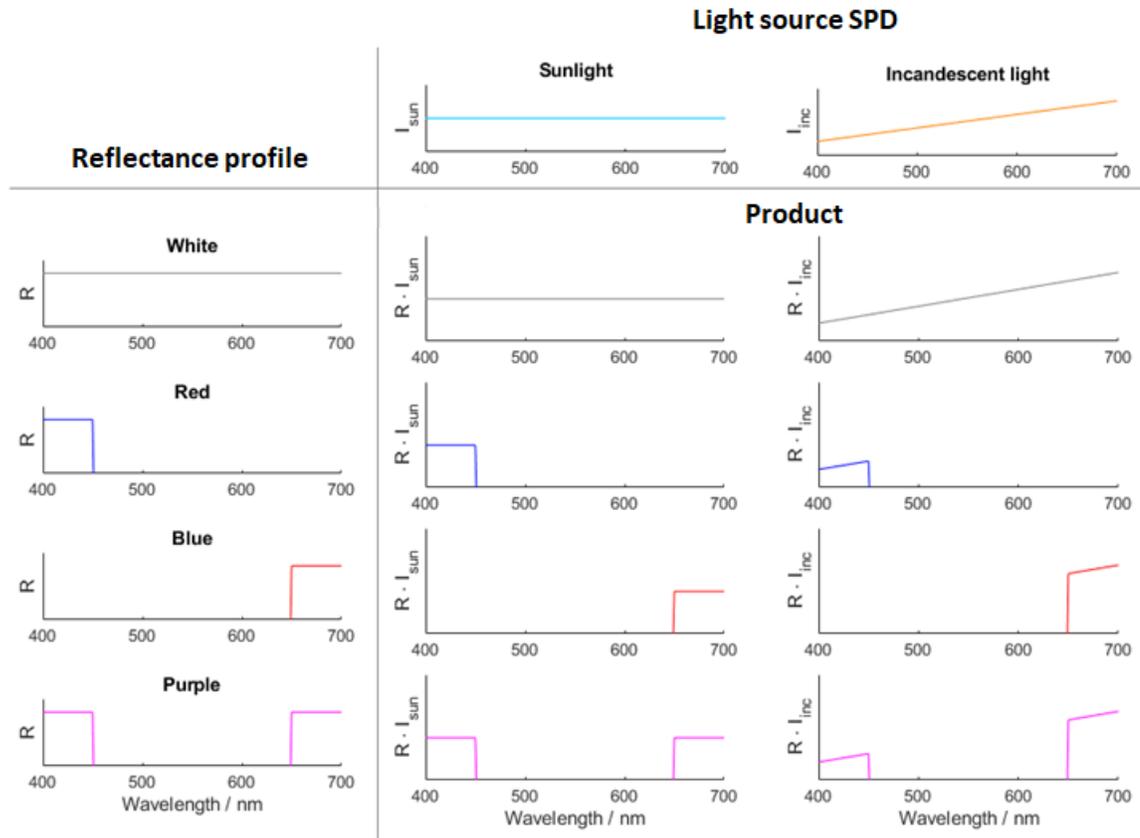


Figure 1.10 Diagrammatic explanation of the effects a light source’s colour temperature has on the SPD of light reflected from coloured surfaces (red, blue, purple and white for reference), represented with simplified spectral profiles to aid comparison. Each spectrum is a product of the SPD of the light source (shown at the top of each column), and the reflectance profile of the object (shown on left of each row). The product of the light source SPD and reflectance profile ($R \cdot I$) shows the SPD of light input into the human visual system while observing the coloured surface under high-colour temperature sunlight or lower colour temperature incandescent light. The relative intensity of the products shows that red surfaces reflect more light under illumination from lower colour temperature light sources, and the opposite can be seen for blue surfaces. Purple surfaces display a combination of the two effects.

Comparing the light reflected from the same objects under the two light sources, it becomes obvious that under incandescent illumination the same objects reflect more red light than blue light. However, the human visual system is equipped to compensate for such changes with the colour constancy mechanism. A proposed basis for colour constancy states that white objects (those with a ‘flat’ reflectance profile) will continue to appear white when the colour of the illuminant is altered.²⁷ To understand how this occurs we need to consider the sensitivity of the cone cells responsible for human colour

vision (figure 1.11, row 1). For the white sample to appear white under both types of illumination, it must cause the same cone response in each scenario. The response of each type of cone cell can be considered as the integral of the stimulus SPD and cone response product across the visible spectrum (defined as 380 - 780 nm):

$$\begin{aligned}
 S(\text{short}) &: \int_{380\text{nm}}^{780\text{nm}} R \cdot I \cdot \beta \, d\lambda \\
 M(\text{medium}) &: \int_{380\text{nm}}^{780\text{nm}} R \cdot I \cdot \gamma \, d\lambda \\
 L(\text{long}) &: \int_{380\text{nm}}^{780\text{nm}} R \cdot I \cdot \rho \, d\lambda
 \end{aligned} \tag{1.9}$$

Where R is the reflectance spectrum of the sample, I is the SPD of the light source and β, γ & ρ are the spectral response function of the S, M and L cone cells respectively. To achieve colour constancy, the response of S, M and L cone cells can be individually increased or decreased by a multiplication factor but their spectral response cannot be altered, this is the basis of the von Kries colour constancy model.¹ This model forms the colour constancy correction in the L*a*b* colour space described earlier. In the case of these simulated illuminants, the luminosity or the perceived brightness of the light sources is set to be the same, consequently for colour constancy to be achieved: the S cones need to increase their response by a factor of 1.6, the M cones are used for reference therefore unchanged, and the L cones decrease their response by a factor of 0.9. Once this has been applied, the outcome is illustrated in figure 1.11 (Row 2) where the response of the cone cells to light reflected of the white sample is the same for each type of illuminant. However, it also becomes clear that the cone response of non-white samples is not compensated for by such a colour constancy mechanism.

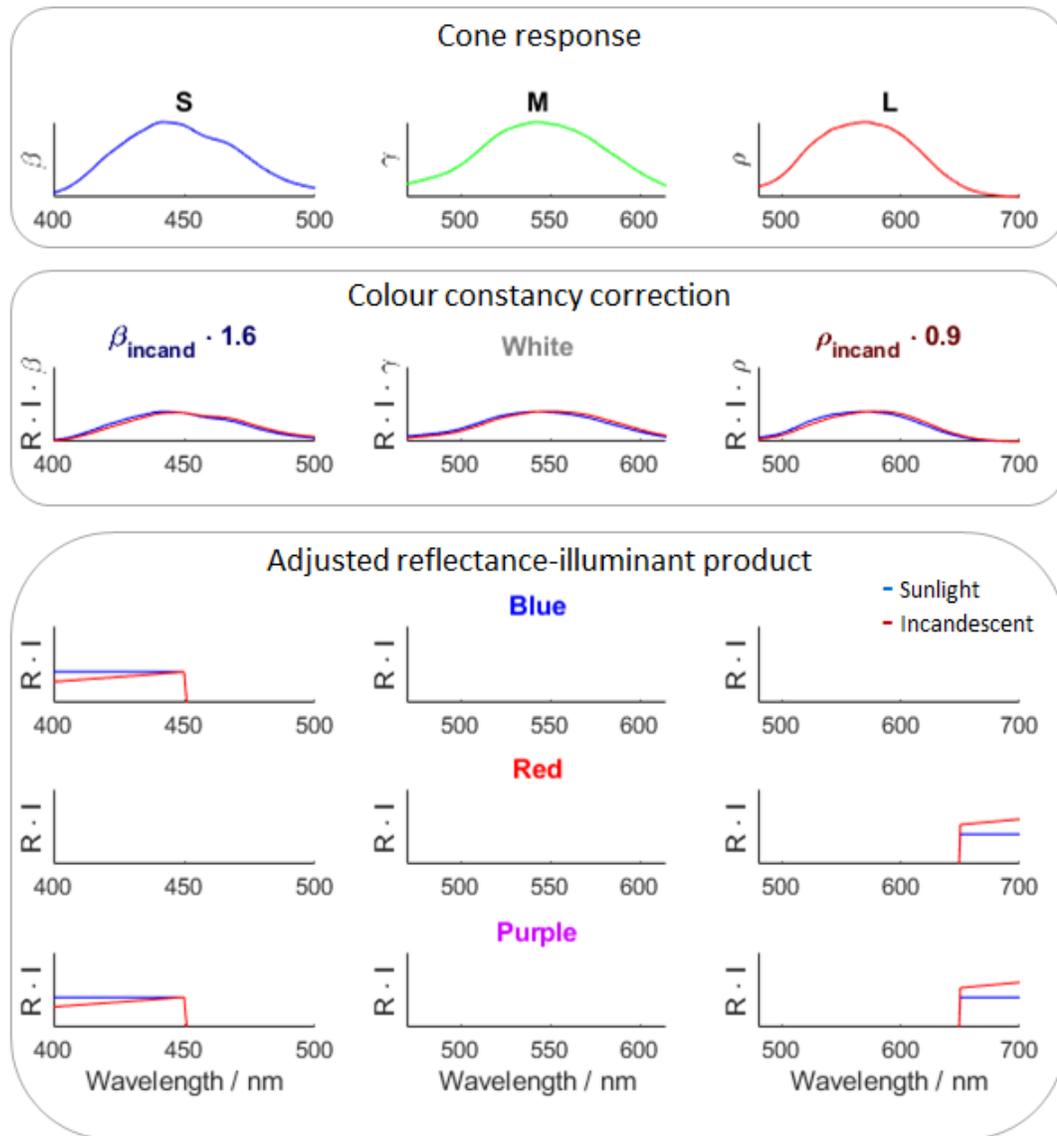


Figure 1.11 Diagrammatic explanation of why red, blue and especially purple colours cannot be adjusted for in colour constancy. The reflectance-illuminant ($R \cdot I$) products described in figure 1.10 are adjusted for colour constancy using the von-Kries model. The top row shows the response of human cone cell where β , γ & ρ denote the sensitivity of each type of cell. The middle row shows the SPD of light input into the human visual system while observing a white surface under sunlight and incandescent light, the SPD has been adjusted for the response of human cone cells. A colour constancy correction is derived by scaling the response of the S and L cells under incandescent light as to achieve the same cone response (defined in eq. 1.9) under each light sources, in this case the response of the S cells is increase by a factor of 1.6 and the L cells decreased by a factor of 0.6 to account for colour constancy. This colour constancy correction is applied to blue, red and purple surfaces illuminated by incandescent light by scaling the $R \cdot I$ product by 1.6 in the blue region and 0.9 in the red region. The correction appears to match the cone response of white objects under the two light sources of different colour temperature, but does not match the SPD of light reflected from coloured surfaces under the same conditions. Blue surfaces have a smaller $R \cdot I$ product under incandescent lighting, while red surfaces have a higher $R \cdot I$ product. Purple surfaces display a larger difference in appearance under these light sources due to a combination of the two effects.

The product of the simulated illuminants with the samples reflectance spectra show that: blue samples have a tendency to appear darker under incandescent lighting than under sunlight, while red samples generally appear brighter under incandescent illumination compared to sunlight. Purple samples exhibit an extra-skewed colour due to a combination of the effects seen in red and blue samples; and the overall effect is a large shift of the objects hue towards a more red appearance. In colorimetry, the effect of objects changing their colour appearance under different colour temperature illuminants is known as *colour inconstancy* and occurs in most coloured objects, although usually to an unnoticeable degree.¹

1.4 Detergency

It is claimed that 90% of clothing garments are disposed of before they become unfit for purpose.²⁸ This is a major problem for climate change as the clothing industry which manufactures them is claimed to be the second largest polluter after the oil industry.²⁹ Laundry detergents which can better clean garments are key to prolonging their usable life-span and lowering the environmental impact of producing new clothing. The laundry industry is, however, in itself a heavy contributor to climate change.

In 2005, It was estimated that residential laundry alone in the United States, which is one of the worlds most polluting countries,³⁰ accounts for over 8% of the sectors total carbon dioxide emissions.³¹ That includes the energy required to heat a large volume of water to a temperature at which laundry detergents can operate effectively; while the energy required to move, treat, and use the amount of water needed in the laundry process is estimated to make up a further 0.6% of total United States carbon-dioxide emissions.³²

Aside from the energy consumption associated with laundry, disposal of the waste wash water also forms a major environmental challenge. Typically, laundry detergents are formulated to form alkaline washing solutions, and following a wash cycle this solution is disposed of into the drain system or directly into the environment. It has been shown that in both scenarios this can cause serious adverse effects.^{33,34} Furthermore, the mechanical agitation involved in the wash process, especially on synthetic fabrics, has been found to release small fibers (micro-fibers) into the wash solution.³⁵ These fibers are around 12-16 μm in diameter with a length typically 360-660 μm , making them difficult to filter out at water treatment plants and ultimately they end up being released back into water streams; a study testing tap water around the world has found 83% of samples tested to be contaminated by fibres from the laundry process.³⁶ This issue is exacerbated as micro-fibers have now been identified as the main source of microplastics in the oceans,³⁵ which in-turn have been described as “the greatest

potential threat to marine environment for the whole world".³⁷

In total, these findings indicate that laundry detergent formulations, in addition to simply cleaning garments, need to be formulated to work: using less water, at lower temperatures and more neutral pH levels. These are extremely challenging problems which need to be solved in the next generation of laundry.

In addition to the above requirements for future laundry detergent formulations, they must also correct the colour of fabrics which arises from un-avoidable degradation mechanisms as the fabrics age. *Yellowing* is one example of such a process, and is a major contributor to degradation of fabric's visual appearance. This effect will be explored in more detail below.

Formulating better laundry detergents first requires a method of quantitatively assessing and comparing their performance. The methods currently used are primarily aimed at assessing the performance of formulations on stains which directly detract from the appearance of fabrics, such as visible soils and coloured food stuffs. However, many laundry stains may also contain lipids. Lipids found in laundry stains do not absorb strongly in the visible part of the spectrum, and therefore in low-moderate concentrations usually do not affect the colour of fabrics. For this reason the removal of lipids is usually overlooked in the washing process. Such residues can build up on garments and start a series of degradation processes, causing odour, or act as a site for bacteria growth, ultimately leading to a shortened usable lifespan of the garment.³⁸

1.4.1 Yellowing

Yellowing is a group term for a series of degradation processes that, over time, result in formation of blue-absorbing chromophores in or on the surface of textiles. The resulting 'dingy' yellow appearance is generally undesirable and can sometimes lead to premature disposal of garments.

In cotton fabrics, the yellow appearance is mostly due to two competing oxidative reactions. Exposure to the atmosphere and sunlight can cause oxidation of the multiple primary and secondary hydroxyl groups in cellulose. It was recently reported that yellow colour is exclusively a result of carbonyl group formation on the cellulose polymer chain, and the process is amplified by nearby carboxyl groups which find their origin also as products of oxidation.³⁹ Yellowing is also not exclusive to cotton, it is observed in virtually all commonly used textiles.^{40,41}

Tackling this issue has historically been the work of laundry detergents, which contain additives such as *fluorescent whitening agents* (FWA) or *hueing dyes* (HD). These optical active components seek to neutralise the yellow colour during the wash cycle and their function will be explored in greater detail. Clothes manufacturers are now considering more exotic solutions, such as coating the fibers with more stable polymers.⁴² However, until this technology becomes viable, the world continues to rely on FWAs and HDs.

1.4.2 Fluorescent whitening agents

Fluorescent whitening agents (FWA) or *optical brighteners* are blue-fluorescent chemicals added to most white paints and materials, as well as laundry detergents. They absorb in the invisible NUV region of the spectrum (below 400 nm) and emit blue light (400 - 450 nm), which compensates for the lack of blue light reflected from yellowed fabric, owing to the introduction of chromophores. The compounds used for this purpose are typically derivatives of stilbene, which is a simple molecule consisting of two phenyl rings connected by a conjugated carbon-carbon double bond.⁴³ Other FWAs based upon coumarins, imidazolines, diazoles, triazoles, and benzoxazolines have also previously been used for this purpose.⁴⁴ Interestingly, this mechanism is even found in some natural materials, with an example being human teeth.⁴⁵

1.4.3 Hueing dyes

A different approach to correcting for yellowing is to dye the degraded fabric purple-blue hueing dyes (HD) also during the wash cycle. This results in a more-flat reflectance profile, devoid of the ‘dingy’ yellow appearance at the expense of reduced overall brightness.⁴⁶ The chemistry of HDs is as broad as dyes themselves, and therefore the exact dyes used in formulations are not typically disclosed by the laundry industry.

1.5 Machine learning

The process of automated function fitting forms the core of machine learning. Algorithms used in machine learning can be sub-divided into two categories: ones which require a set of data containing both the input and correlated output variables, called training data and measured using an external technique; and others which require no prior knowledge about the input dataset. These are called supervised and unsupervised algorithms respectively.

The supervised type operate on a given set of input variables and desired output variables, a reflectance spectrum and concentration of analyte in a sample for example. There exists many algorithms which can ‘learn’ the relationship between the two sets of data. This relationship, also referred to as a ‘model’, and can then be used to infer the output from a new set of input data. The accuracy of these types of models is ultimately limited by the accuracy and reliability of the training data, as well as the complexity of the relationship between the input and output space. Unsupervised algorithms require no training data, however, unlike the former type, the output is subject to interpretation by the user or further analysis.

A brief introduction to the specific functions used in this work will now be provided. The functions were used from the MATLAB Statistics and Machine Learning add-on package, and further information about the computation is available elsewhere.⁴⁷

1.5.1 Principal component analysis

Reflectance spectroscopy can generate large amounts of data, with a spectrum for a single sample usually consisting of over a thousand variables. In cases where multiple samples need to be compared or analysed together, data analysis can become computationally intensive, or sometimes even impossible with the available hardware. However, spectra usually contain a large amount of multicollinearity, where spectral bands are

likely to be correlated between each other. These correlations can be removed to simplify the dataset using a dimensionality reduction technique such as principal component analysis (PCA).⁴⁸ PCA creates a new set of variables called principal components (PCs) which are linear combinations of the original variables. The PCs are formed by rotating the original axes in order to describe the maximum amount of variance in the dataset, it is therefore necessary to mean-center data prior to application of PCA.

Mathematically, PCA can be described by considering a set of n samples with K number of variables associated with each sample. The dataset is then formed by defining each sample as a vector in K dimensions. PCA can be performed on the dataset by firstly forming a $K \times K$ covariance matrix (\mathbf{C}) using the whole dataset. The off-diagonal entries of \mathbf{C} describe the correlations between variables. They can be subsequently removed by diagonalisation, which is carried out by solving the characteristic equation:

$$\det[\mathbf{C} - \sigma^2 \mathbf{I}] = 0 \quad (1.10)$$

where σ^2 is an eigenvalue and \mathbf{I} is the identity matrix. The eigenvalues (σ^2) are the projections of the variance of each variable onto the corresponding eigenvectors, which can be found by solving:

$$\mathbf{C}\mathbf{v}_j = \sigma^2 \mathbf{v}_j \quad (1.11)$$

for each eigenvalue (indexed by j). Each eigenvector contains the coefficients, also called loading or weighting, of the original variables which can be used in linear combinations to produce what are called PC scores.⁴⁸

Once the PC scores have been determined, they may be used for data analysis in place of the original variables with the added advantages of: a single PC usually accounts for the majority of the variance in the total dataset, therefore only a few PCs are

needed to describe the entire dataset. Secondly the PCs are orthogonal by definition, making them ideally suited for regression analysis. The combination of using PCA followed by linear regression on the resulting scores is a popular supervised machine learning technique called principal components regression (PCR), and the process will be described in more detail as it is applied.

The use of PCA in the context of spectroscopy can be further demonstrated with simulated data; in this example, a spectrum with a single Gaussian peak can be defined by the expression:

$$f(x) = ae^{-\frac{(x-b)^2}{2c^2}} \quad (1.12)$$

where a defines the peak's amplitude, b the center position, and c the standard deviation or the peak width. An arbitrary x axis is defined with a set of integers 1:1000 in unit steps. A set of Gaussian curves can be defined by $a = 0.2, 0.4, 0.6, 0.8,$ and 1 , with b and c as constants of value 500 and 50 respectively (arbitrary). The peak height and integral both scale linearly with a , which can be observed in figure 1.12 (a-left). It can be seen that PC1 describes 100 % of the variance in the dataset, with the weightings for PC1 forming a Gaussian curve. The initial 1000-dimensional dataset can therefore be completely described by a single variable, the score of PC1. Furthermore, a PCR type model can be applied by fitting a first order polynomial to the score of PC1 against peak amplitude, leading to perfect fit with gradient 9.4 and intercept -5.3 (2 sf).

To further demonstrate the effect of spectral features in PCA, a more complex example is given. A set of spectra have been defined as a sum of two Gaussian functions, one increasing and the other decreasing linearly in magnitude throughout the dataset. The functions take the form:

$$f(x) = \alpha e^{-\frac{(x-300)^2}{2(50)^2}} + \beta e^{-\frac{(x-700)^2}{2(50)^2}} \quad (1.13)$$

where α and β have been arbitrarily set to: $\alpha = 0.2, 0.4, 0.6, 0.8, 1$, and $\beta = 2, 1.6, 1.2, 0.8$, and 0.4 . The resulting simulated curves can be seen in figure 1.12 (b-left). The PCA coefficients, shown in figure 1.12 (b-middle) now take the shape of two horizontally offset Gaussian peaks, with one showing twice the amplitude and opposite sign. PC1 again describes 100 % of the variance in the dataset, showing how multiple correlated features from 1000 input variables may be decomposed to a single variable. An ideal PCR type model can also be fitted with a linear trend of gradient -21 and intercept 13 (2sf).

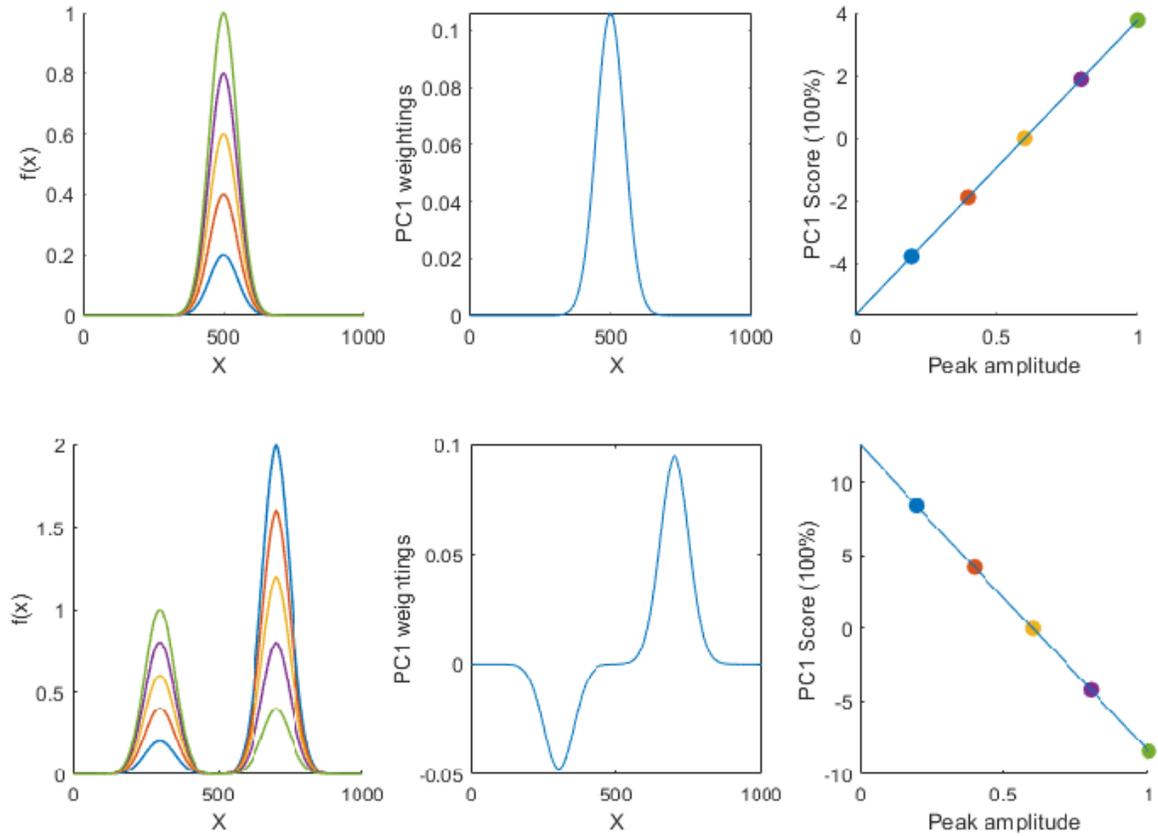


Figure 1.12 Demonstration of PCA on an example dataset consisting of Gaussian functions. The functions $f(x)$ in the left column have been set up using equations 1.12 (top) and 1.13 (bottom). The top row demonstrates a function consisting of a single Gaussian curve which increases linearly in magnitude set by variable $\alpha = 0.2, 0.4, 0.6, 0.8$ and 1 . The bottom row demonstrates a function that consists of two summated Gaussian curves, one increasing and the other decreasing in magnitude linearly, as defined by the variables $\alpha = 0.2, 0.4, 0.6, 0.8, 1$, and $\beta = 2, 1.6, 1.2, 0.8$, and 0.4 respectively. The middle column shows the weightings used to define PC1, calculated using PCA on the corresponding set of functions to the left. The right column shows a plot of the PC1 scores as a function of peak amplitude (defined by α). The scores are calculated by linear combinations of each curve and the PC1 weightings, each point corresponds to a single curve shown in the same colour on the left-hand set of functions.

1.5.2 t-SNE

Another method of visualising datasets with many variables is t-distributed stochastic neighbour embedding (t-SNE), which is an unsupervised technique primarily used for data exploration. In t-SNE, the output is two dimensional unit-less scatter plot, with the axes formed to again maximise variance, only in this case local similarities between

variables are preserved. This means that the output can be analysed in a similar manner to a score plot in PCA, however, one consideration is that t-SNE is not deterministic, meaning that the exact shape of the output will be different with every computation.⁴⁹

1.5.3 Gaussian process regression

Gaussian process regression (GPR) is a non-parametric approach to regression that has been shown to be highly effective in extracting quantitative information from NIR spectral data, and works well on small datasets.⁵⁰ A Gaussian process is a collection of random variables such that any finite subset exhibits a joint Gaussian distribution. In GPR, this idea is extended to a distribution of functions that is used to represent the training data, and inference can then be carried out by taking the mean of the function distribution sampled at a single point. A major advantage of this approach is that the distribution of functions provides a measure of the uncertainty in each prediction, meaning that the model can evaluate how confident it is in relating the input to the training data used to create the model. A detailed explanation of the maths behind this method is available elsewhere⁵¹ and the algorithms used to fit GPR models vary between implementations; the approach used in this work is documented in ref. 52.

1.5.4 Validation

In fitting a function to a dataset, it is important to prevent over-fitting. Over-fitted functions are very good at describing the dataset presented (training data) but fail when the function is tested on external data (validation data). Over fitting can generally be reduced using the following techniques: (a) increasing the size of the training dataset which increases the complexity requirements for over fitted functions; or (b) using validation techniques to supervise the fitting procedure. Validation aims to prevent over fitting by assessing the quality of the fit on a dataset that is external to that used in parameter fitting. Mathematical, this is defined by calculation of the loss function using the fitted model on a dataset separate to that used in minimisation the residuals vector, which is used to optimise the coefficients which define a model. Validation is key to producing robust models which are tolerant to random fluctuations which are always present in training data, as well as real-world data.

An external dataset can be obtained from a subdivision of the total acquired dataset by using two common validation techniques: holdout validation and k -fold cross validation.

1.5.5 Holdout validation

As the name suggests, this a method which is carried out by excluding a specified proportion of the overall dataset from the training data to be used as validation data. The fitted model can then be used to predict the responses of the excluded samples to assess the model performance on data not present in its training set. A measure of model performance, is then calculated from the difference between the predicted and known responses in the validation dataset. Holdout validation is usually used with large datasets as the excluded data is not used in the fitting process. The amount of data excluded needs to be statistically significant to represent the entire dynamic range of the model, which may impact model performance if only a small number of samples are present in each region of the input range.

1.5.6 k -fold cross validation

Cross validation can be considered as a more rigorous version of holdout validation. It is carried out by splitting the total dataset into k number of groups. Then a series of k iterations removes a single group to be used as the validation data, with the remaining groups used as training data. Model performance is then calculated from the difference between the predicted and know responses in the excluded group during each iteration. This iterative approach is more suited to smaller datasets as the entire dataset is used in model construction. Furthermore, computational requirements are increased compared to holdout validation, which maybe not be feasible with more complex models and larger datasets.

The error in fitted models is assessed using a parameter called RMSE, which is calculated during the validation process as to include the effect of over-fitting in assessing model performance. RMSE is calculated as the quadratic mean of residuals when the fitted function is applied to the external validation data:

$$RMSE = \sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - y(x_i))^2} \quad (1.14)$$

where n is the number of samples in the validation dataset, y_i is the set stain concentration, x_i is the model input, and $y(x_i)$ is the model output of the i^{th} sample. This empirical parameter provides an estimate of model accuracy, however the error associated with model prediction requires a more rigorous analysis with external data.

Chapter 2

Colourimetry and tuning colour appearance

In studying the appearance of objects using colourimetry models, it is important to investigate their limits, and particularly phenomena which may not be possible to predict using our existing models of human colour perception. In this chapter, the mechanisms which can affect human colour perception will be the object of study through a combination of spectroscopy and computational modelling.

Initial studies were focused on the unusual range of appearances which can be attributed to compounds containing lanthanide ions, and origin of these discrepancies identified a misnomer commonly referred to in the literature.

Following this, the attention turned to tuning colour appearance, and the ideas developed were then applied to control the colour change seen in one particular lanthanide salt (neodymium chloride), which has led to the development of an entirely new mechanism for controlling the appearance of an object.

Finally, the colour appearance models which were rigorously tested in the sections described above were then applied to colourimetry of consumer laundry. A model used for predicting the appearance of fabric under different lighting conditions, as well as levels

of degradation and colour-correction by FWAs or HD is described and implemented.

2.1 Introduction

2.1.1 Discrepancies in colour appearance

The background of why objects may appear different colours is foremost discussed. There are many examples of objects which can appear different colours under different conditions, some objects are artificial while, others are completely natural and have been commonly known since early civilisation such as the Uluru landmark in Australia, that can appear to change from deep red to orange-brown at sunrise; or the reverse at sunset.⁵³ Similarly, the less common *Alexandrite effect* describes the difference in appearance of the mineral ‘Alexandrite’, which can appear any hue ranging from purple to green depending on the illumination source and chromatic adaptation of the observer.⁵⁴ This effect has also been attributed to the cause of the colour change seen in compounds containing rare-earth elements, and will be discussed in more detail later in this section.

There are many well known mechanisms for which a single object can appear different colours, the most common of which can be summarised as follows:

Luminescence

In luminescence, photons of higher energy are absorbed by atoms, typically leading to emission of lower energy photons, the process is also sometimes referred to as *down-conversion*. It may cause a drastic change to the appearance of an object if the absorbed light is in the ultra-violet (UV) region, which is otherwise invisible and therefore does not contribute to the colour of the illuminant regardless of intensity. If a material can down-convert the absorbed energy to a photon in the visible portion of the spectrum, then it’s colour shifts towards the hue of the emitted light. This is a particularly challenging area of colourimetry to model as discussed in a previous section (1.3).

Photochromism

The absorption of light by matter results in excitation of the irradiated atom to an excited state, the de-excitation pathway which occurs in-turn can result in a physical or chemical change within the material. This could be a chemical reaction⁵⁵ or a change in conformation,⁵⁶ for example. The result may alter the materials frequency-dependant dielectric properties which, if in the visible region affects its perceived colour. This effect may be permanent or reversible and is widely used in commercial applications. An example of this effect in everyday life is utilised in auto-darkening sunglasses.⁵⁷

Dichromatism

The colour of a material depends not only on its dielectric properties, but also on how much light is absorbed when passing through the material before reaching the observer. This is affected by the effective path length in the measurement as well as concentration of absorbing species. A common example where changing these factors leads to a large difference in colour is seen in pumpkin seed oil, which appears green in a thin layer and changes to deep red as the thickness or concentration increases.⁵⁸ Most translucent materials which are coloured will display this property, although the difference in hue is usually too small to be noticeable compared to the difference in saturation. Dichromatism can be explained by the Beer-Lambert law and by the response of the three types of cone photoreceptors in the observer. It is particularly prevalent in substances that have an absorbance spectrum with two features, one wide but weak in intensity, and another narrow but with high intensity. As the absorption of the sample increases, the perceived hue changes from that defined by the position of the wide but shallow feature to one defined by the deep but narrow feature.⁵⁹ In pumpkin seed oil this is caused by two absorbance bands, one high in magnitude peaking around 405 nm, and the other low in magnitude but broad peaking in the yellow region around 575 nm.⁶⁰

The change in hue stems from the non-linear relationship between absorbance and perceived colour. In gemstones, this effect is sometimes also referred to as *Usambara effect*.⁶¹

2.1.2 The Alexandrite effect

Some natural minerals exhibit a colour change when viewed under different light sources. Perhaps the most striking of these is the mineral Alexandrite which can appear any hue from green to deep purple (figure 2.1). This change in colour occurs when the illuminant under which the sample is observed changes between a high colour temperature source such as daylight, to lower colour temperature one such as incandescent lighting. The phenomenon is called the ‘Alexandrite effect’ and was first observed in the mineral Alexandrite which is a type of Chrysoberyl (BeAl_2O_4), originally discovered in the Ural Mountains of Russia.²⁰

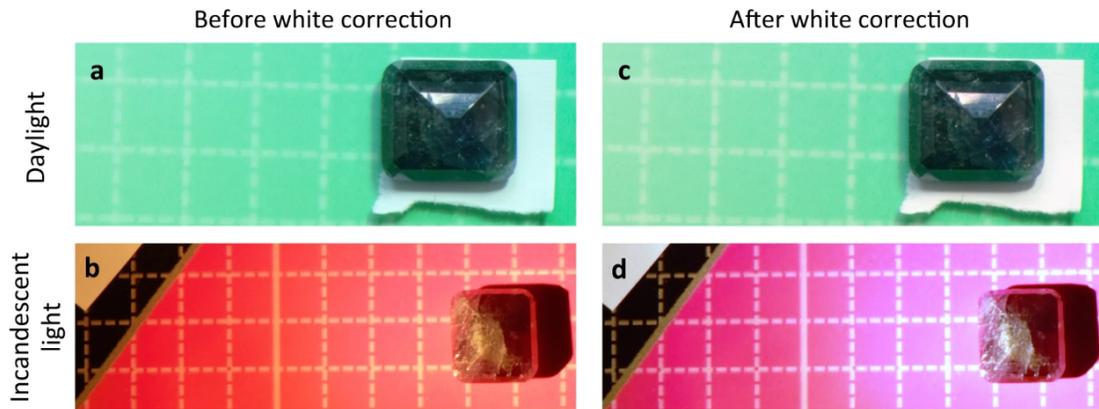


Figure 2.1 Mineral Alexandrite observed under daylight (top) and incandescent lights (bottom) displaying the Alexandrite effect. Reproduced with permission from ref. 54.

Pure Chrysoberyl is actually colourless and the colouration in Alexandrite originates from the absorption of light by chromium impurities, which replace a small portion of aluminium atoms in the beryllium aluminate lattice with chromium. This is present as Cr^{3+} ions which have two transitions (absorbance bands) prevalent in the visible region: one in the UV-blue region (380-450 nm) and the other in the yellow region (565-590

nm).^{62,63} The resulting optical spectrum contains two bands of high transmission, one centred at around 490 nm in the blue/green region and another in the red region, which peaks in the NIR region:

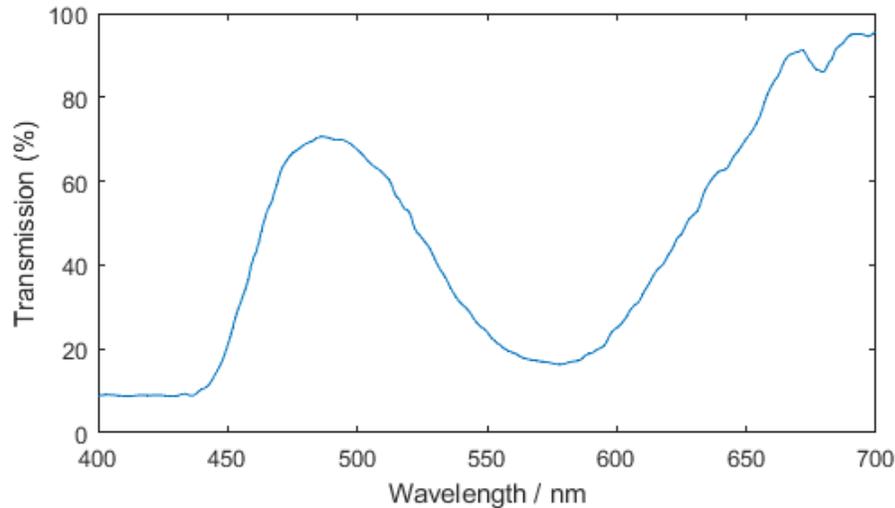


Figure 2.2 Transmission spectrum of a mineral Alexandrite sample which displays a strong colour change between daylight and incandescent lighting.⁶⁴

The result is that this mineral appears green under broad spectrum illumination; which may seem contradictory due to the strong transmission feature in the red part of the spectrum. However, due to the green-centred sensitivity of the human visual system, the blue/green transmission peak dominates the apparent colour. In translucent materials, transmission properties are closely related to reflectance properties, as the light typically passes through the material before reaching the observer, meaning that the two types of spectra can be analysed in a similar manner.

The colour change of Alexandrite has been the subject of numerous studies. Early studies concluded that the phenomenon was a ‘psychophysical’ one involving: the adaptive aspect of the human vision system, the two transmission bands in the blue-green and red regions, and the spectral power distributions of the different kinds of light sources.^{20,65–67}

Colourimetric analysis of Alexandrite has previously been carried out by applying the CIE colour matching functions on spectra of light reflected from the sample.⁶⁴ The

results showed that the observed colour change is also concordant with the CIE system, it is therefore not outside of our current understanding of human colour vision.⁶⁸ The ‘Alexandrite effect’ is not exclusive to mineral Alexandrite, in fact samples of this gemstone which exhibit such a strong perceived colour change are extremely rare and valuable. Certain ruby, sapphire and other gemstones that feature trivalent chromium also exhibit a colour change when the light source is changed from sunlight to incandescent light, although not as pronounced as Alexandrite.^{20,64,69} Synthetic gemstones which exhibit this colour change are now routinely synthesised; a popular approach involves doping sapphire with low levels of vanadium, which results in similar optical properties to that of Alexandrite.^{20,54,70} It has recently been shown that the petals of purple *Torenia* flowers can also exhibit this colour change property to a small degree.⁶⁸

With so many different objects showing a change in colour between sunlight and incandescent lighting, the classification of what is the Alexandrite effect has been challenging. Early studies concluded that simply the presence of two transmission bands in the blue/green and red regions of the spectrum, caused by an absorption band in the yellow region is enough to classify a material as possessing an Alexandrite or Alexandrite-like effect.⁶⁹ A later study looked at a collection of four materials: Alexandrite, coloured glass, garnet and sapphire. The study then categorised the observed colour changes into four types based on the CIE Lab colour space.⁶⁸ The defining metric was the hue angle measured between the a^* and b^* coordinates under a set of standard illuminants. This classification system has since been adopted for identifying if materials which change colour under different illuminants exhibit the true Alexandrite effect; which is surprising given the study only included four samples.^{71–73}

The true cause of the Alexandrite and Alexandrite-like colour changes was only recently published in 2020, despite the concept having already been well established in colorimetry literature.

In the case of Alexandrite, while it does not appear purple under broad spectrum

illumination, the transmission/reflectance spectrum contains two bands which expressly interact with the S and L cone cells, causing an overall hue shift towards red under lower colour temperature illuminants due to colour inconstancy (described in section 1.3.3). The position and width of the two transitions responsible for the transmission spectrum of Alexandrite have been found to be critical to achieve such extreme colour inconstancy which gives the mineral its sought-after appearance.⁵⁴ These spectral band shapes are primarily governed by the chromium levels in the beryllium aluminate lattice,²⁰ which explains why only a small fraction of naturally occurring Alexandrite with the right level of chromium doping exhibits a strong colour change effect.

The reflectance or transmittance spectrum of Alexandrite also satisfies the criteria for dichromatism, which is sometimes misinterpreted as the origin of the colour change,²⁰ although it does have an effect on the magnitude of colour inconstancy observed. The anisotropic orthorhombic crystal structure of Alexandrite also facilitates a change in observed colour due to pleochroism,⁷² but this is unlikely to be observed under natural conditions.

2.1.3 Compounds containing rare-earth elements

One phenomenon that may challenge our current understanding of colour perception is the uncertain colour of particular lanthanide compounds. Anhydrous neodymium chloride as an example, can be described as appearing pale green or lilac according to different chemical suppliers.^{74,75} While this may be attributed to impurities or isomerism, it is actually found that the appearance of a single sample can vary between these two colours depending on the light source used as the illuminant. Under sunlight (or other broad spectrum light sources) neodymium chloride appears lilac, yet under fluorescent lighting the observed colour is pale green. These observations can be confirmed by multiple observers, and even replicated with digital photography as

shown in figure 2.3. The phenomenon has been reported in a range of chemical compounds that contain rare-earth elements,^{73,76–79} although the colour change is yet to be fully explained. Recently the same phenomenon has also been reported in actinides, particularly compounds of americium and neptunium.^{80,81}



Figure 2.3 A single sample of anhydrous neodymium chloride under sunlight (left) and fluorescent lighting (right) including a QP 203⁸² colour standard card for reference. Photographs were taken using a Canon EOS 60-D digital camera.

Upon closer inspection of this phenomenon, it becomes clear that the colour change observed in rare-earth elements cannot be due to the ‘Alexandrite effect’, or even colour inconstancy more generally. The colour change in Alexandrite or ‘Alexandrite-like’ materials occurs when the colour temperature of the light source is altered. In this novel class, however, the difference in colour is never reported when the illuminant is changed from sunlight to an incandescent source. Furthermore, fluorescent lighting is sometimes used as a ‘daylight simulator’ in studies of Alexandrite-like materials, apparently because samples appear similar colours under both illuminants, yet in compounds containing rare-earth elements, changing between daylight and fluorescent lighting triggers a large and noticeable change in perceived colour.⁶⁴

If the new type of colour change observed in rare-earth compounds can be understood,

then there is a possibility that it can be tuned. Controlling the appearance of objects using the illuminant has significant commercial applications. In the retail industry for example, colour is a key aspect of a products aesthetic; and while retailers may not always be able to control the colour of their products, they can chose the light sources used to display them. The sale of red meat is a prime example of a scenario where this is utilised, warmer illumination causes a more-red appearance of the product⁸³ which is typically considered more appealing.⁸⁴ The tuning of the colour change seen in rare-earth compounds will therefore also form a part of this study.

2.1.4 Colorimetry of consumer fabrics

The laundry detergent industry strives for products which better clean fabrics and to replicate the users-experience, this is usually assessed through colourimetric assessment of fabrics in combination with wash testing. Describing the colour of consumer garments may seem like a trivial exercise, however, if reliable quantitative results are required then the exercise becomes far more complex, and the most challenging aspects will herein be outlined.

The challenges to colourimetry of fabric stem from the contamination and mechanical abrasion that functional fabrics are typically subject to, as well as the numerous degradation mechanisms which impact the appearance of fabrics, these have been discussed in a previous section (1.4.1). Colourimetry of consumer fabrics is further complicated by laundry detergents themselves, which introduce optically active additives such as FWAs and HDs to the fabrics during a wash cycle. Consumers can choose between many laundry detergent brands and therefore the nature of these additives is subject to vary between washing cycles, as well as the concentrations in which they are used in.

Overcoming these challenges to yield a method which can reliably assess the performance of laundry detergent formulations is key to optimising laundry detergents to

better correct the colour of garments for the effects of yellowing, and ultimately extend the usable life span of a garment as far as possible. The future of current technologies to combat yellowing is, however, becoming more uncertain as advances in lighting technology mean that the spectral profile of light illuminating the washed fabrics is going to change. This is particularly challenging for FWAs, which rely on the light source having some high-energy near-UV (NUV) component to give the renowned ‘whiter-than-white’ appearance.

LED based lighting is quickly becoming the most wide-spread lighting technology due to its relatively high efficiency and colour rendering compared to incandescent and vacuum fluorescent alternatives. The efficiency of LED lighting, among other factors, stems from its spectral output, which is typically very low in the NIR and NUV regions where the human visual system is least sensitive.

FWAs will fail to correct the appearance of fabric for degradation under a certain type of LED lighting, which is quickly becoming the most wide spread form of lighting. The type of LEDs primarily affected are the blue-pumped white LEDs, which consist of a 450 nm GaN LED coated with a cerium-based (typically $Y_3Al_5O_{12}:Ce$) down-converting phosphor,^{85,86} which are the most common type currently used in lighting technology. This was the conclusion of a study published by Houser and co-workers,⁸⁷ which has since spurred several news stories warning about the effects of a ‘dingy’ appearance even after washing.^{88,89} The authors of the study even suggested that lighting manufacturers should include violet LEDs into lighting products, which result in an overall reduction in the efficiency of the lighting,⁸⁹ but still allow FWAs to function. Such light sources are now commercially available for applications where the whitening effect of FWAs is crucial.⁹⁰

2.2 Results and discussion

2.2.1 Colourimetry of compounds containing rare-earth elements

Initial studies into the true cause of the colour change of rare earth compounds used anhydrous neodymium chloride as the representative sample. It is readily available and was found to exhibit a strong colour change between natural sunlight and fluorescent lighting. After these initial studies, the concepts will be expanded to cover a broader range of rare earth compounds.

Anhydrous neodymium chloride appears lilac when viewed under sunlight but appears pale green when viewed under fluorescent lights; this is the opposite of the colour change observed with the Alexandrite effect. To investigate the currently accepted connection to Alexandrite, the CIE system (introduced in section 1.3.2) may be employed to predict the colour under sunlight and incandescent lighting. The first step in this procedure was to measure the reflectance spectrum of anhydrous neodymium chloride using DRS. The SPD of sunlight and a 2800 K incandescent lamp were also recorded (figure 2.4). The light diffusely reflected off the samples surface can be calculated as the product of the light source's SPD and the reflectance of the sample, both as a function of wavelength. The coordinates of the reflected light in the CIE Lab colour space can then be calculated using colour matching functions. In this study the L^* coordinate (or perceived brightness) of the reflected light is of little relevance since only the change in hue is of interest, which places the focus onto the a^* and b^* coordinates. The calculated Lab coordinates are therefore displayed as projections onto a two dimensional plane (figure 2.4).

In this calculation, 'chromatic adaptation' or the process responsible for colour constancy, was estimated using the von Kries transformation (described in section 1.3.3). More complex and accurate corrections do exist, however, these usually use the von

Kries hypothesis as their basis.²⁷ Furthermore, in this case the simple von Kries correction was found to give results concordant with observations.

Colourimetry is usually carried out with standard illuminants as the light source; in this case we are interesting in explaining the observed colour change, therefore the measured SPD of two broadband sources, sunlight and an incandescent lamp, which are known to cause this phenomenon will be used instead.

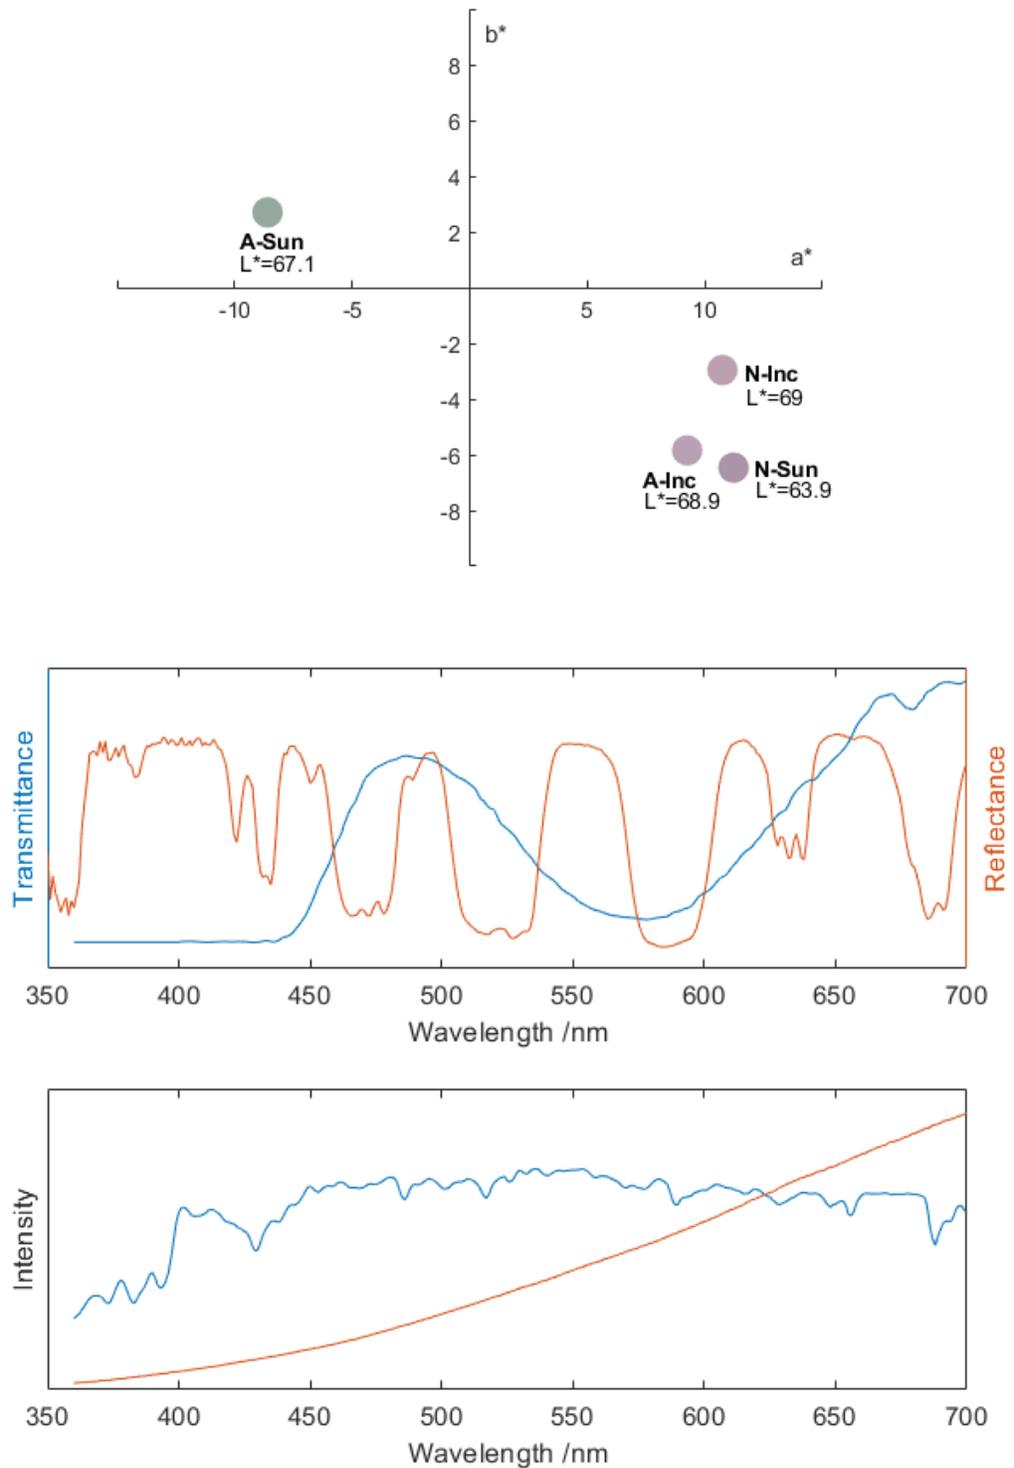


Figure 2.4 (Top) CIE $L^*a^*b^*$ colour space plot showing chromaticity change of Alexandrite (**A**) and neodymium chloride (**N**) as the illuminant is exchanged from sunlight (**Sun**) to incandescent lighting (**Inc**), the label next to each point corresponds to the sample and illuminant whose spectral properties are displayed below. The colour of the points is a conversion of the $L^*a^*b^*$ coordinates into the sRGB colour space. (Middle) Transmittance spectrum of Alexandrite⁶⁴ (blue) and reflectance spectrum of neodymium chloride (orange). (Bottom) SPD of midday sunlight recorded in Durham, UK (blue) and a 2800 K incandescent lamp (orange).

The $L^*a^*b^*$ coordinates of Alexandrite correctly predict the green to purple hue shift as the illuminant is exchanged from sunlight to incandescent lighting (all $L^*a^*b^*$ coordinates are provided in section 5.2), which is consistent with literature colourimetric studies.⁶⁴ By contrast, the hue of neodymium chloride is almost unaffected by this change of illuminant. This is consistent with observations and further supports the argument that there is another mechanism, separate from the Alexandrite effect, responsible for the colour change phenomenon in rare-earth compounds.

The reflectance spectra of Alexandrite and neodymium chloride show little similarity. While the colour of Alexandrite arises from broad transitions within the d subshell of the chromium ions, broadened by vibrations in the crystal field; the colour of neodymium chloride is a result of a myriad of parity-forbidden transitions in the $4f$ subshell. The transitions become allowable due to non-centrosymmetric interatomic and crystal field effects on the ions. However, the intensity of $f-f$ transitions are still much weaker than those in the d subshell. The ‘shielding’ of the unfilled $4f$ subshell by the filled $5s$ and $5d$ subshells greatly reduces the interatomic and crystal field effects on the $4f$ transitions, producing the relatively narrow bands.²⁰ This effect also means that the reflectance spectra of the majority of chemical compounds containing neodymium ions show little variation, hence the colour change phenomenon is reported in a wide variety of compounds containing neodymium.

The SPD of sunlight and the incandescent emitter are both continuous over the visible spectrum with only a few sharp features towards the NIR portion of sunlight, which can be attributed to Fraunhofer lines.⁹¹ These light sources can both be considered as black-body radiators, emitting a continuous frequency spectrum that depends on the body’s temperature. Conversely, fluorescent lighting utilises a completely different mechanism for generation of light; stimulated (electron-bombardment induced) emission from gaseous mercury is used to initially generate UV radiation, and a mixture of phosphors is then used to down-convert a portion of the mercury emission to visible

light through luminescence. Two common phosphor technologies include: rare earth element based, which are typically used to produce cool white lights, and halophosphate based which are typically used to produce warm white illumination. SPDs from both types were recorded for colorimetric measurements (figure 2.5).

The colour of cool white fluorescent lamps is designed to mimic that of natural sunlight, however the SPD is completely different. While sunlight is effectively continuous through out the visible spectrum, fluorescent lights constitute high intensity peaks which carry the main components necessary to generate ‘white’ light as perceived by the human visual system. In colourimetry, the effect of different SPD’s producing an identical cone response is called *metamerism*. Even in the case of warm white fluorescent lamps, which are orange in colour, the effect of colour constancy is that their colour is still perceived as white.

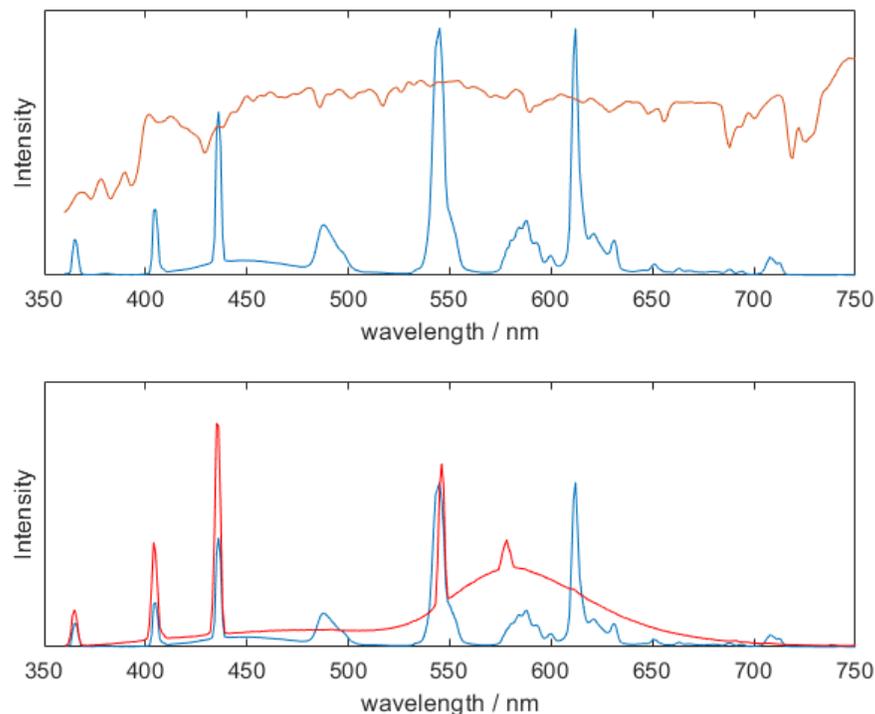


Figure 2.5 (Top) SPD of sunlight (orange, recorded at midday in Durham, UK) and white fluorescent lighting (blue) which responsible for the colour change in neodymium chloride. (Bottom) SPD of typical cool white (blue) and warm white (red) fluorescent lights.

Metamerism can be considered as an artefact of the human visual system having less channels than required to represent the spectral information; it appears in several forms and presents many challenges to colourimetry.²⁶ For example, the two light sources described above may appear the same colour to one person but different to another. This is classed as *metameric failure* and in the example case it is caused by *observer metamerism*, as the observers' visual systems have subtly different spectral responses. If a sample has a reflectance profile which is not even across the visible spectrum then it is also subject to *illuminant metamerism*, which describes the differences in observed colour of an object under different white light sources. The ability of a light source to maintain the colour of objects relative to a reference illuminant, such as D65, is referred to as colour rendering. A commonly used metric to assess this is the colour rendering index (CRI), which may be calculated for white light sources based on the colour of a set of CIE standard under the light source (discussed further in this section).¹ Fluorescent lighting typically scores relatively low in CRI compared to incandescent or LED illuminants.⁹²

The interaction of different light sources with neodymium chloride can be visualised by considering the spectral composition of the diffuse reflections off the surface, which have been calculated as the product of the light source SPD and reflectance spectrum (figure 2.6). For the sake of comparison, the intensity of the light sources has been scaled as to maintain an equal luminosity, which was estimated as the integral of the SPD- \bar{y} product.

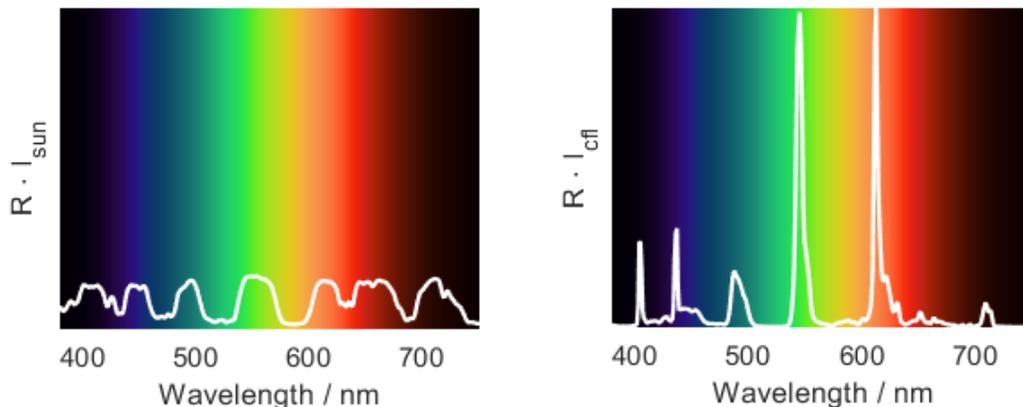


Figure 2.6 SPD of light diffusely reflected from anhydrous neodymium chloride under sunlight (left) and cool white fluorescent lights (CFL - right). Intensities are scaled as to maintain an even luminosity between the light sources. The background has been coloured as to approximately represent the perceived colour of each portion in the visible spectrum.

The results show that under sunlight the SPD is distributed across the spectrum in a series of bands. Under cool white fluorescent lighting, however, the distribution is completely different and dominated by two peaks in the green and orange regions. The contrast between these two spectra is particularly significant in the green region: under sunlight the purple appearance must be influenced by the absorbance bands at 525 nm (Nd: $^4I_{9/2} \rightarrow ^4G_{9/2}$, $^2K_{13/2}$ and $^4G_{7/2}$),⁹³ yet in fluorescent lighting the major green component comes from an atomic transition of mercury at 546 nm (Hg: $^3P_2 \rightarrow ^3S_1$),⁹⁴ which lies in between two absorption bands of neodymium chloride and is largely reflected.

The big discrepancy between the SPD of diffuse reflections under each light source shows that this is a case of illuminant metameric failure, and this can be confirmed by colourimetry. However, to exclude any contribution from colour inconstancy, the light sources must have the same colour temperature. The CCT of a light source may be approximated from CIE xy coordinates using the McCamy method (2.1).⁹⁵ It was found that the recorded sunlight spectrum approximates to 5310 K while the cool white fluorescent lights approximate to 4310 K.

$$CCT(x, y) = -449n^3 + 3525n^2 - 6823.3n + 5520.33$$

$$n = \frac{(x - 0.3320)}{(y - 0.1858)} \quad (2.1)$$

This is a significant difference and in order to compensate for it, the broad sunlight spectrum was replaced in calculations by a 4310 K black body radiator derived using eq. 1.3. The SPD of such a source is shown in figure 2.7 and the resulting CIE xy coordinates can be considered identical to those of the recorded cool white fluorescent lights (fluorescent lights: $x=0.369$ $y=0.375$ and black body radiator: $x=0.370$ $y=0.374$) within the accuracy of the instrumentation used in this study.

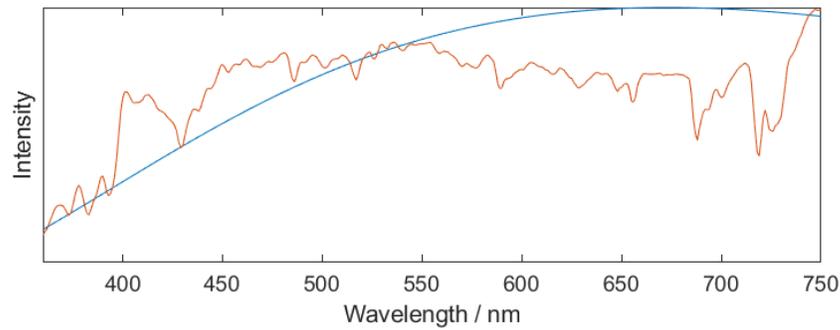


Figure 2.7 SPD of sunlight recorded in Durham, England at midday (orange) and a 4310 K black body radiator (blue) calculated using the plank equation (1.3), the same CCT as that of cool white fluorescent lighting used in this study.

The colour matching functions were applied on SPD's of light diffusely reflected from the neodymium chloride surface under the 4310 K black body radiator and cool white fluorescent lights. The von Kries chromatic adaptation functions were included to compensate for minor differences in the illuminant colour. The resulting CIE $L^*a^*b^*$ coordinates were found almost identical to those under measured sunlight (displayed in figure 2.4). The predictability of this colour change phenomenon confirms that this is a case of illuminant metameric failure as opposed to colour inconstancy, which has been excluded by keeping the colour temperature of the light sources equal.

These results can be verified by observers: both the human eye and digital cameras.

Photographs of the sample are provided in figure 2.3, and the Lab coordinates extracted from them are similar to those predicted with colour matching functions (provided in 5.2).

The colourimetry studies carried out have shown that the nature of the colour difference in anhydrous neodymium chloride is due to illuminant metamerism, however, with such a complex reflectance spectrum it is not possible to identify the specific features which are responsible for this effect from first principles. To investigate this matter further, a set of computational reflectance spectra has been constructed by sequentially removing sets of Nd absorption bands from the recorded diffuse reflectance spectrum. The resulting spectra were multiplied by the SPD of a 4310 K black body radiator and cool white fluorescent lights in turn to calculate the SPD of light diffusely reflected from the surface. Colour matching functions were then applied and the $L^*a^*b^*$ coordinates calculated under each light source (figure 2.8).

The results show that removing most of the absorption bands does not contribute to illuminant metamerism, does not have a significant effect on the degree of illuminant metameric failure observed, instead leading to a translation of appearance in the a^*b^* colour plane. The direction of the colour change shows little variation in the majority of the modified spectra, with the exception of **1** and **3** where the magnitude of the colour change is significantly reduced. Removing the bands at 430 nm (${}^4I_{9/2} \rightarrow {}^2P_{1/2}$),⁹³ **1**, changes the appearance under fluorescent lighting to pale cyan instead of green (figure 2.8 - **1**), whereas removing the transitions at 525 nm (${}^4I_{9/2} \rightarrow {}^4G_{7/2}, {}^2K_{13/2}, {}^4G_{9/2}$),⁹³ **3**, eliminates the purple appearance entirely (figure 2.8 - **3**). This suggests that the colour change observed in anhydrous neodymium chloride is largely due to these two absorption bands, which is to be expected given the overlap previously observed with the sharp mercury transitions utilised in fluorescent lighting.

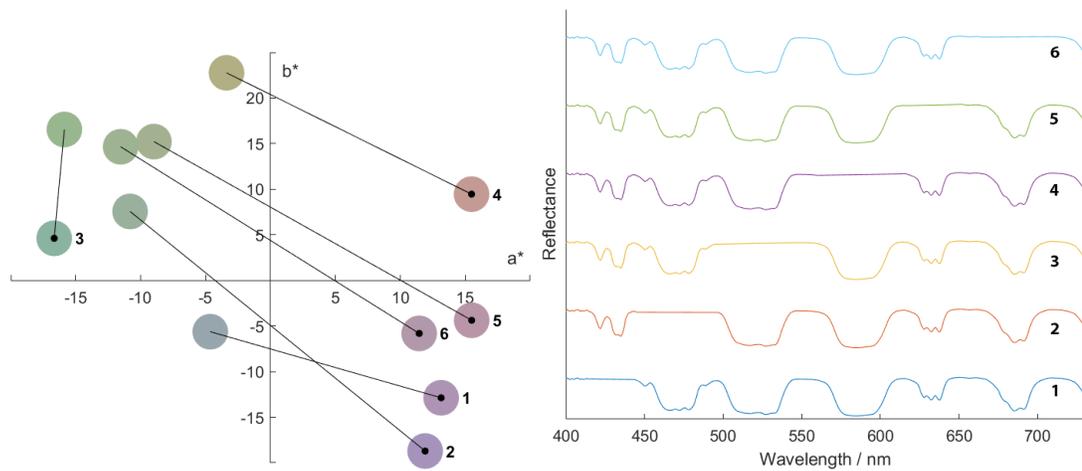


Figure 2.8 (Left) $L^*a^*b^*$ coordinates calculated from a set of modified reflectance spectra under a 4310 K black body radiator (central black spot) and cool white fluorescent lighting of the same CCT. Each set of points describes the appearance of a theoretical surface with a reflectance profile of corresponding number (shown to the right) under the two lighting conditions. The colour of the points is a conversion of the $L^*a^*b^*$ coordinates into the sRGB colour space. (Right) Set of reflectance spectra of anhydrous neodymium chloride which have been computationally modified by systematically removing groups of absorption bands. The spectra have been offset vertically for clarity.

The hypothesis can be investigated further by considering another set of modified reflectance spectra, which were simulated by including a single set of absorbance bands onto an otherwise 50 % reflective sample using the KM approximation of reflectance. The summation of reflectance spectra was carried out by: firstly applying a rectangular window function to isolate each set of absorbance bands in turn, then all sets of spectra were converted to KM space using equation 1.6 before being summed with the 50% reflective profile (also converted to KM space); applying the inverse KM function then yields the final simulated set of spectra (figure 2.9).

The same colorimetry procedure as that used on the previous set of modified spectra was carried out in order to study the illuminant metamerism contributed by each set absorption bands. The resulting $L^*a^*b^*$ coordinates show that the majority of modified spectra show very little illuminant metamerism with the exception of spectra **1** and **3**. **1** shows a grey to pale green change in hue, while **3** shows little change in hue, but a drastic change in saturation. This distinction is consistent with the previous set of simulated

spectra and suggests that a combination of the two bands could almost entirely describe the colour change observed in neodymium chloride. Such a spectrum, comprised of only two absorption bands was created (7) and the predicted colour is: purple under broad spectrum illumination, and yellow / green under fluorescent lighting, which is very close to that of the original sample. This confirms that observed colour change in neodymium chloride can be almost entirely accounted for by two sets of absorption bands, $^4I_{9/2} \rightarrow ^2P_{1/2}$ and $^4I_{9/2} \rightarrow ^4G_{7/2}$, $^2K_{13/2}$, $^4G_{9/2}$, in the blue and green portions of the visible spectrum.

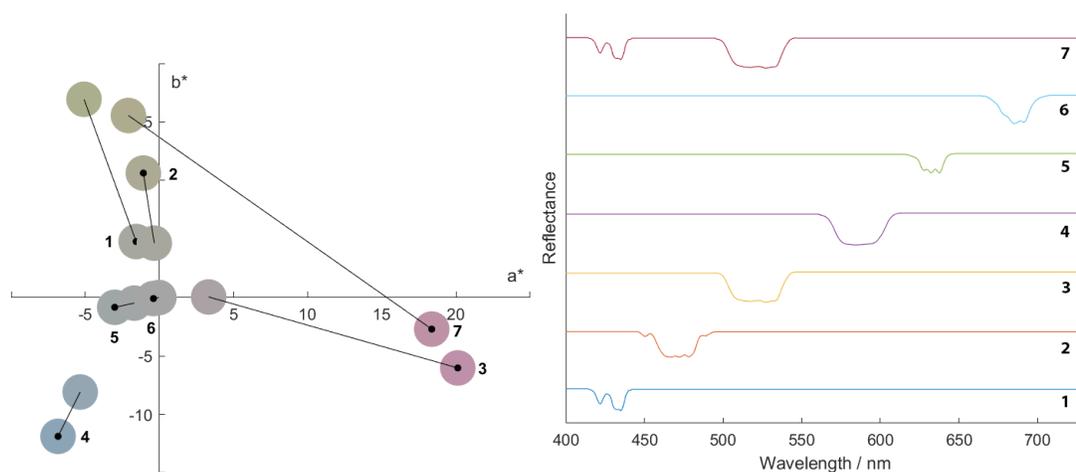


Figure 2.9 (Left) $L^*a^*b^*$ coordinates calculated from a set of theoretical reflectance spectra under a 4310 K black body radiator (central black spot) and cool white fluorescent lighting of the same CCT. Each set of points describes the appearance of a theoretical surface with a reflectance profile of corresponding number (shown to the right) under the two lighting conditions. The colour of the points is a conversion of the $L^*a^*b^*$ coordinates into the sRGB colour space. (Right) Set of theoretical reflectance spectra constructed by systematically including a single group of neodymium chloride absorbance onto an otherwise evenly reflective surface (with 7 as an exception which is a combination of 1 & 3). The spectra have been offset vertically for clarity.

Neodymium chloride is deliquescent, and hydrated neodymium chloride also exhibits an interesting susceptibility to illuminant metamerism. The hexahydrate sample used in this study consisted of larger crystals than those of the anhydrous sample. This leads to reduced scattering and a longer effective path length through the sample, resulting in stronger absorption of light. These effects are visible in the reflectance spectrum of the compound (figure 2.10), where the absolute reflectance is attenuated by a factor of

6. It is also seen that the addition of water to the crystal structure causes a variation to the shape the transitions at 525 nm (${}^4I_{9/2} \rightarrow {}^4G_{9/2}$, ${}^2K_{13/2}$ and ${}^4G_{7/2}$) and 580 nm (${}^4I_{9/2} \rightarrow {}^2G_{7/2}$, ${}^4G_{5/2}$) in the neodymium salt.

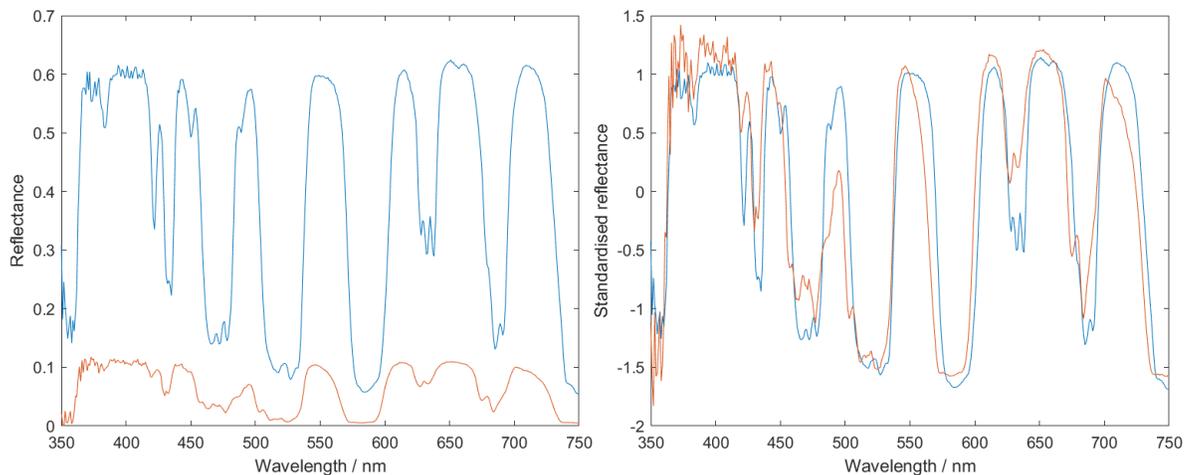


Figure 2.10 Diffuse reflectance spectrum recorded from two different samples of neodymium chloride (blue - anhydrous, orange - hydrated). Spectra displayed on the right are standardised to be mean-centred with a unity standard deviation.

The affected transitions are called hyper-sensitive transitions,⁹⁶ and the difference can be better visualised by comparing standardised reflectance spectra of both samples (figure 2.10 - right). The spectra appear only subtly different, however, the illuminant metamerism with fluorescent lighting becomes more complex. Colourimetric analysis (figure 2.11) shows that the colour of neodymium chloride hexahydrate under sunlight to be the same hue of purple as that of the anhydrous sample, only more saturated. Under cool-white fluorescent lighting, the reverse pattern is observed with the hydrate having a relatively less saturated green hue; and under warm-white fluorescent lighting, the green appearance is lost entirely and the sample appears a more-blue shade of purple.

These findings are difficult to rationalise from first principles due to the complex shape of the spectral profile of light reflected from the hydrated sample under both types of fluorescent lighting (displayed in figure 2.11 - right). The spectra appear similar with subtle differences, under warm white illumination the sample reflects more in the

blue region compared to cool-white illumination, while the green and red components appear more distributed through the visible spectrum. Aside from the influence on predicted colour, these results also suggest that under cool-white illumination, neodymium chloride is more susceptible to observer metamerism.

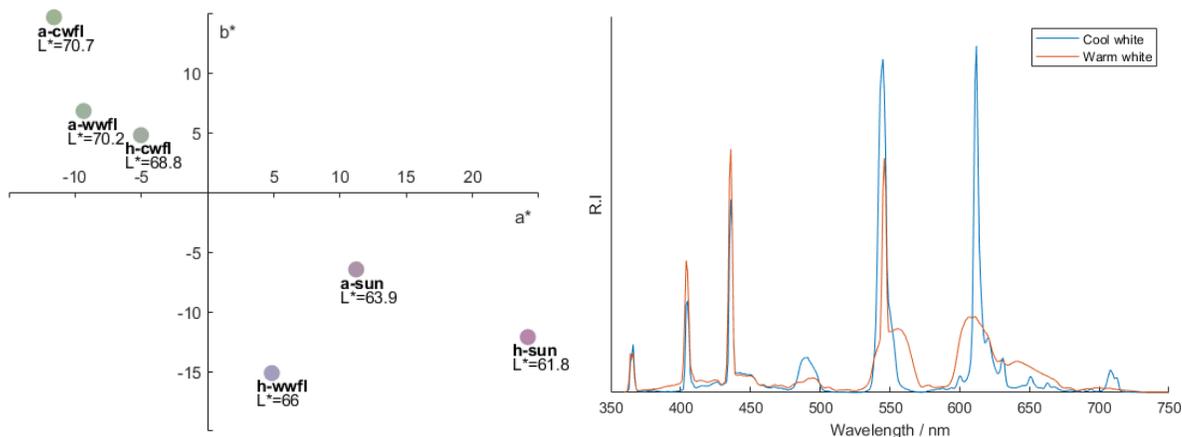


Figure 2.11 (Left) $L^*a^*b^*$ coordinates of neodymium chloride (a-anhydrous, h-hexahydrate) under sunlight and fluorescent lighting (cwfl-cool white, wwfl-warm white). The colour of the points is a conversion of the $L^*a^*b^*$ coordinates into the sRGB colour space. (Right) SPD of light reflected from neodymium chloride hexahydrate under warm-white and cool-white fluorescent lighting, calculated as the product of the light source SPD and the reflectance profile of the sample. Spectra have been scaled to maintain an equal luminosity.

Neodymium is not the only lanthanide that forms salts which are reported to show a significant colour change under fluorescent lighting. To study the colour change in other lanthanides, a series of compounds which represent a single sample from each lanthanide that has absorbance bands in the visible region was collected. A further six coloured lanthanide compounds were investigated for the alexandrite effect and metameric failure under fluorescent lighting. Reflectance profiles from hydrated salts of holmium sulphate, praseodymium chloride, samarium nitrate and erbium acetate were measured using the same DRS technique as that used for the neodymium chloride samples and the spectra are displayed in figure 2.12. Optical spectra from salts of promethium nitrate⁹⁷ and thulium sulphate⁷⁷ were investigated from previously reported studies where they were measured in aqueous solution, and after conversion of the raw spectra to transmittance, they will be used in place of reflectance spectra for an approximate comparative analysis.

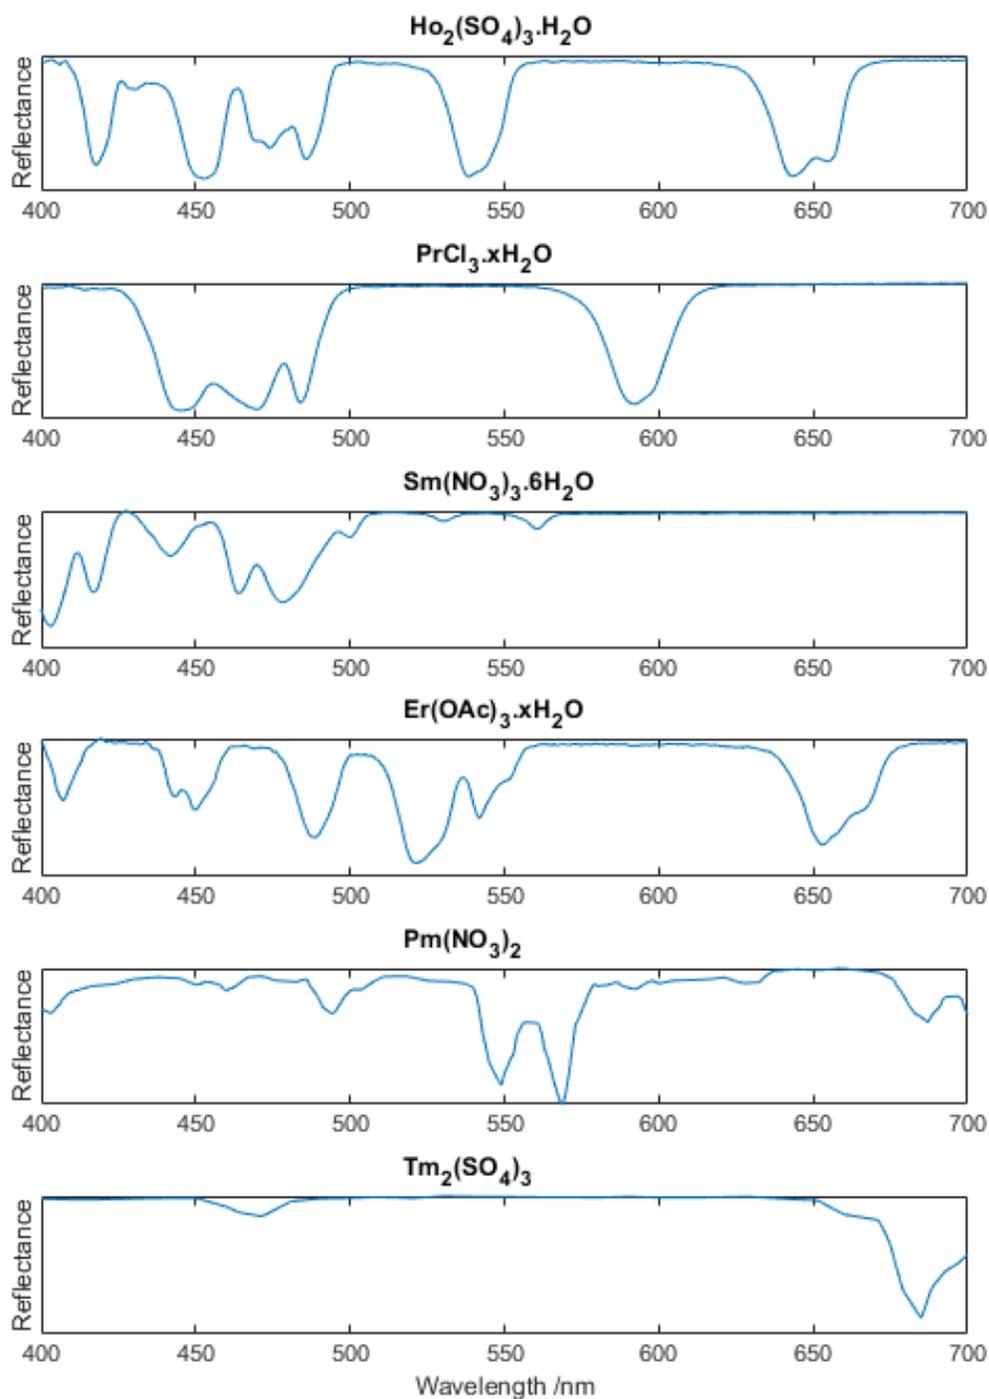


Figure 2.12 Standardised DRS spectra of hydrated lanthanide salts. Promethium⁹⁷ and thulium⁷⁷ samples were approximated to reflectance from aqueous transmission measurements which have been reported in literature.

The reflectance spectra again consist of a myriad of sharp $f - f$ transitions, similar to

those observed in neodymium chloride. The spectral profile of the light reflected from each sample was calculated for each compound under sunlight and an incandescent lamp, to test for colour inconstancy or the ‘Alexandrite effect’; as well as cool-white fluorescent lamp and a 4310K black body radiator (the same colour temperature as that of the cool-white fluorescent lamp) to test for illuminant metamerism. The results were analysed in the $L^*a^*b^*$ colour space by applying the CIE colour matching functions onto each product of the illuminant and the samples reflectance spectrum (figure 2.13).

The coordinates show that the tested samples display no significant degree of colour inconstancy, as the appearance of all samples is nearly indistinguishable when sunlight or an incandescent lamp is used at the illuminant. Further confirming that the colour change observed in lanthanide compounds is not due to the ‘Alexandrite effect’ (colour inconstancy) as previously reported.

Looking at the difference in appearance between fluorescent lighting (cool-white) and a black body radiator of the same colour temperature, shows that the holmium sample has a large susceptibility to illuminant metamerism under fluorescent lights, which is comparable in magnitude to that observed in neodymium compounds. The remainder of the tested samples show no significant change in hue between illumination from a broad-spectrum source and fluorescent lights.

It is surprising that the promethium salt appears colourless and yet it has distinct absorption feature in the green portion of the spectrum. This is somewhat displayed in the calculated $L^*a^*b^*$ coordinates as a purple hue but with a low saturation. Saturation of hue varies with the effective path length in the measurement, and the low value in this case likely resulted from the diluted solutions used in place of solid sample as sample could not be obtained and suitable reflectance data were not available in the literature. Regardless, it is expected that larger crystals of promethium salts would display a deep purple appearance, however the hue angle in the a^*b^* plane between promethium nitrate’s appearance under sunlight and fluorescent lighting is quite small,

therefore a large degree of illuminant metameric failure is not expected for such salts.

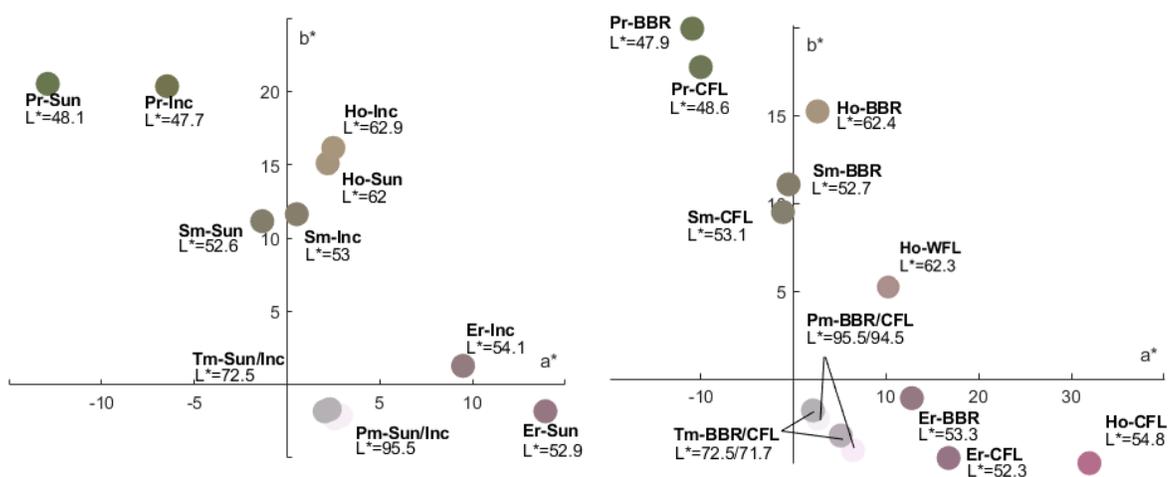


Figure 2.13 Colour appearance in the $L^*a^*b^*$ colour space of six hydrated lanthanide salts calculated from reflectance spectra shown in figure 2.12: holmium sulphate (**Ho**), praseodymium chloride (**Pr**), samarium nitrate (**Sm**), erbium acetate (**Er**), promethium nitrate (**Pm**), and thulium sulphate (**Tm**); illuminated by: a simulated 4310 K black body radiator (**BBR**), Incandescent lighting (**Inc**), cool-white fluorescent lights (**CFL**), and warm-white fluorescent lighting (**WFL**). The label next to each point corresponds to the lanthanide salt and light source used as the illuminant. The colour of the points is a conversion of the $L^*a^*b^*$ coordinates into the sRGB colour space.

The holmium sulphate hydrate sample is predicted to appear yellow when illuminated with sunlight or incandescent lighting and red-purple when illuminated with cool-white fluorescent lighting, yet under warm-white fluorescent lighting the appearance displayed a yellow-orange hue. These observations can be explained by comparing the emission spectra of fluorescent lighting with the absorbance of the hydrated holmium sulphate sample (which was approximated in figure 2.14 using the KM function). It can be seen that the main green component in cool white fluorescent lighting originates from the atomic transition of mercury at 546 nm ($(\text{Hg}: {}^3\text{P}_2 \rightarrow {}^3\text{S}_1)^{94}$), this overlaps significantly with the ${}^5\text{I}_8 \rightarrow {}^5\text{F}_4, {}^5\text{S}_2$ transition of holmium at 536 nm. The impact on colour appearance is that less green light is reflected under cool-white fluorescent lighting compared to sources in which the green component is distributed across more of the spectrum. This effect is further intensified by a similar overlap found between the emission of terbium at 488 nm ($\text{Tb}: {}^5\text{D}_4 \rightarrow {}^7\text{F}_6$),⁹⁸ which is used as a green-blue component in cool-white

fluorescent lighting, and the $^5I_8 \rightarrow ^5F_3$ transition of holmium at 483 nm.⁹⁹⁻¹⁰¹

The absorbance spectrum of holmium also displays two absorbance bands in the blue region at 447 nm (Ho: $^5I_8 \rightarrow ^5F_1, ^5G_6$) and 422 nm (Ho: $^5I_8 \rightarrow ^5G_5, ^3G_5$),⁹⁹⁻¹⁰¹ which in contrast, do not overlap with the atomic transitions of mercury at 404 nm (Hg: $^3S_1 \rightarrow ^3P_0$) and 436 nm (Hg: $^3S_1 \rightarrow ^3P_1$) which contribute to the blue component in both types of fluorescent lighting.^{94,102} The result is the holmium salt reflects less blue light under broad band sources compared to fluorescent lighting, which in combination with band of high reflectance in the green and red regions explains the yellow appearance under sunlight.

A combination of the afore mentioned features mean that the appearance of holmium sulphate hydrate under cool-white fluorescent lighting is dominated by large reflected red/yellow component, with a partially reflected blue component, and a little reflected green component. This is expressed in the predicted colour as a large positive a^* value and a small negative b^* value. In switching the illuminant to warm-white fluorescent lighting, the green component of the light source is now reflected more as it is distributed through the entire green portion of the spectrum. The combined increase in reflectance around the green - red portion of the spectrum is concordant with the orange appearance. The increase in total light reflected from the surface is also consistent with the increased L^* value observed for the holmium sample when illuminated with warm-white fluorescent lights.

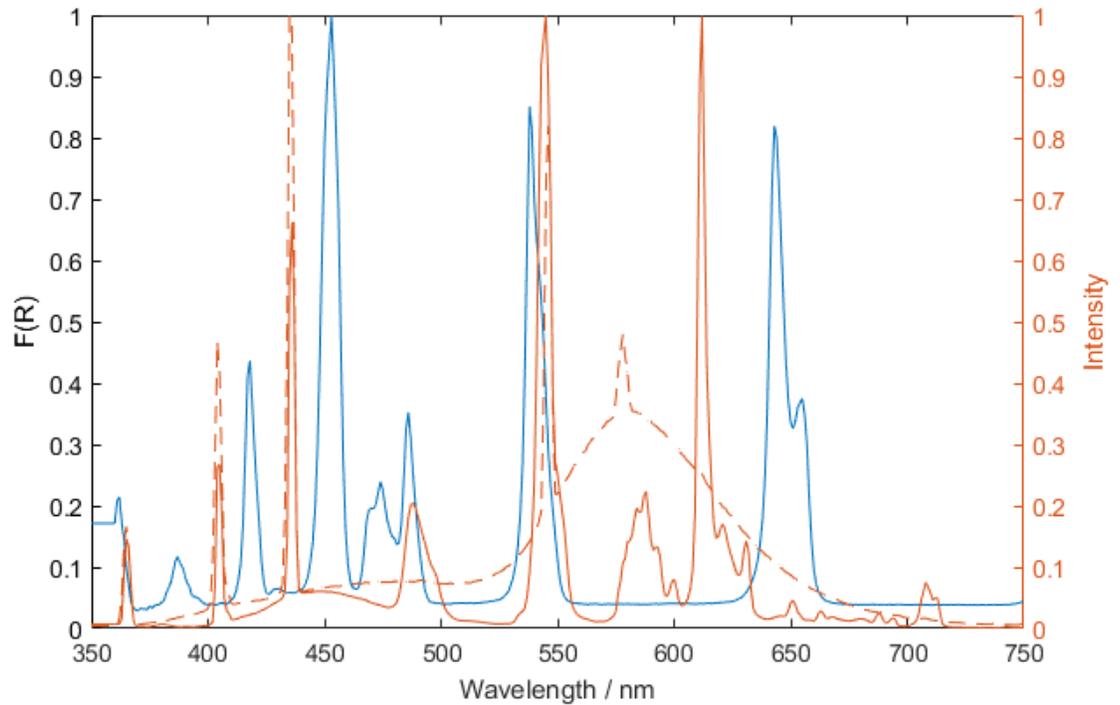


Figure 2.14 Comparative plot of the normalised KM function ($F(R)$) applied to the DRS spectrum holmium sulphate hydrate, including the normalised SPDs of cool (solid orange line) and warm-white (dashed orange line) fluorescent lighting for reference.

2.2.2 Tunable illuminant metamerism

Theory of operation

To achieve a truly colour tunable object, we must be able to tune the spectral composition of light reflected from its surface. In the previous section this was explored by computationally altering the reflectance profile of neodymium chloride. While this was proven effective, it is difficult to achieve experimentally. In this section, the idea that colour must be considered as an illuminant-object pair is exploited. The colour of an object can also be determined by the spectral composition of the light source used to illuminate it.

The illuminant is chosen such that it appears ‘white’ with a variable SPD. For a light source to appear white it must stimulate the photo receptors within the human visual system in the right ratios, however, the exact spectral distribution can vary significantly, exemplified by the case of sunlight and cool white fluorescent lighting (figure 2.5). The sample under illumination is chosen such that it can reflect a portion of each red, green and blue component, as well as absorb a portion of each component at a slightly different wavelength to that it reflects.

A ‘white light’ source can then be constructed by targeting each of the six components’ bands, two for each red, green and blue components required to generate white light. Anhydrous neodymium chloride was selected for initial attempts as its optical properties have already been well characterised. The property which makes it a suitable candidate is displayed in figure 2.15, where it can be seen that the narrow $f-f$ transitions allow the sample to absorb, as well as reflect red, green and blue components of the visible spectrum. A sample of hydrated neodymium chloride, which consists of larger crystals leading to a lower degree of scattering, will also be included in experiments for reference.

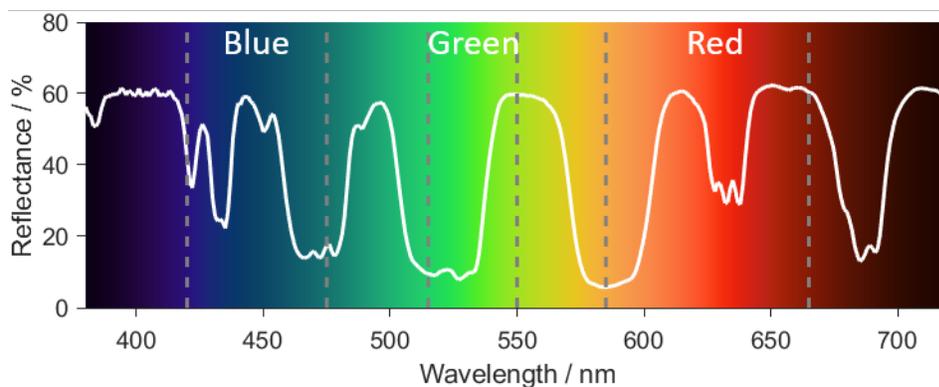


Figure 2.15 Diffuse reflectance spectrum of anhydrous neodymium chloride, highlighting the regions targeted by the spectrally tunable light source to create maximum illuminant metamerism. The background has been coloured as to approximately represent the perceived colour of each portion in the visible spectrum.

The blue and green components are well suited as the sample can reflect blue light at 450 nm while absorbing it at 475 nm, it can also absorb green light at 525 nm while reflecting it at 545 nm. The red component is non-ideal as the absorbance maximum occurs in the yellow part of the spectrum, around 590 nm, but the sample reflect red light at 620 or 650 nm. This may limit the resulting range of colours which the sample can be programmed to appear, especially in the high saturation blue and green hues which require low reflectance of red light.

Spectrally tunable light source

Tuning the spectral composition of a light source has previously been achieved using two main methods: a broad-spectrum light source may be dispersed using a prism or diffraction grating, then unwanted spectral components filtered out utilising devices such as spacial light modulators or digital-micromirror arrays, followed by recombination to yield an output beam;¹⁰³ alternatively, the output from a series of discrete, narrow spectrum emitters such as LED's, can be combined together and the spectral profile of the resulting light source can be controlled by adjusting the power of each emitter in the array.¹⁰⁴ The former typically offers a better spectral resolution, usually at the expense of output power; while the latter can offer higher output power due to

the absence of filters, although the availability of emitters at specified wavelengths and bandwidths may limit the operating range of such light sources.

The light source constructed for use in this study was based on a series of narrow spectrum emitters due to the required optical power output. Human colour vision is poor at low illumination levels, therefore to allow observation by humans with photopic vision, the output power of the LED's needs to be higher than 10 cd m^{-2} .¹⁰⁵ The light source consists of 8 LED's mounted onto a custom made printed circuit board (PCB), which is housed on an aluminium heat-sink facing into a custom integrating sphere. The integrating sphere is required to homogenise the output from the spatially offset emitters. The LED's are driven using a series of step-down converters on a main driver board mounted onto the side of the integrating sphere. The main board features an integrated microcontroller which is used to interfaced with software on a local computer, as well as generate power control signals feeding the LED drivers. The tunable light source unit is mounted on top of a photography booth that houses the samples as well as a colour standard card for reference. In parallel, the computer-side software interfaces with a spectrometer coupled to the output of the light source positioned next to the samples, allowing for feed-back of the optical spectrum. Further details about the light source are provided in section 2.4.2.



Figure 2.16 Experimental setup of spectrally tunable light source. A) tunable light source unit comprised of: integrating sphere shown as split halves in CAD model (B), LED array (D) and LED drivers & power control electronics (C). E) Photography enclosure. F) QP 203 colour standard card. G) Anhydrous neodymium chloride sample. H) PTFE diffusely reflective block. I) Collimator fibre-coupled to spectrometer.

The eight LEDs comprising the light source have emission maxima at: 420 nm, 440 nm, 470 nm, 515 nm, 545 nm, 590 nm, 630 nm, and 655 nm. These will be referred to as channels. The emission spectrum of each channel has been recorded and is presented along with the reflectance spectrum of anhydrous neodymium chloride (figure 2.17), for which the light source is expressly designed to work with. The 420 nm and 630 nm channels are only weakly absorbed by the neodymium chloride therefore were not utilised further. The 440 nm and 470 nm channels form the blue component of the light source. The 470 nm channel is designed to be reflected less than the 440 nm channel due to overlap with the $^4I_{9/2} \rightarrow ^4G_{11/2}$, $^2K_{15/2}$ and $^2G_{9/2}$ transitions. The 515 nm and 545 nm channels form the green component, with the later designed to have a higher reflectance than the former due to the $^4I_{9/2} \rightarrow ^4G_{9/2}$, $^2K_{13/2}$ and $^4G_{7/2}$ transitions in neodymium ions. The 590 nm and 655 nm channels make up the red portion of the light source, with the 655 nm channel designed to have a higher reflectance than the

590 nm channel which overlaps with the ${}^4I_{9/2} \rightarrow {}^2G_{7/2}, {}^4G_{5/2}$ transitions in neodymium ions.

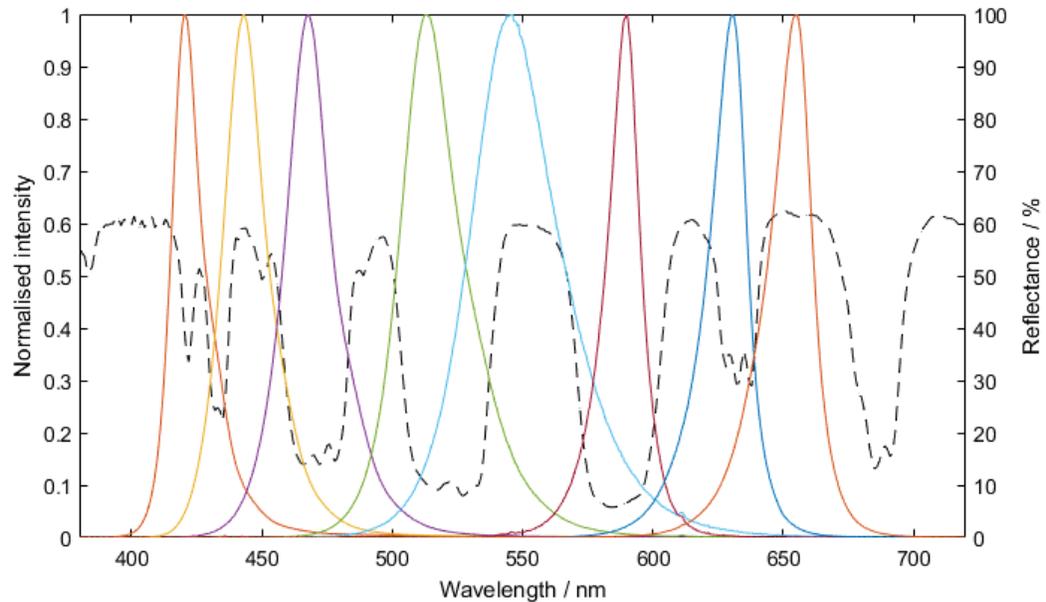


Figure 2.17 (Solid lines) Optical spectra recorded from each LED comprising the spectrally tunable light source. Spectra have been normalised for clarity. (Dashed line) Reflectance spectrum of anhydrous neodymium chloride for reference.

Tuning the spectral output of the light source is achieved by controlling the power supplied to each channel individually, this was done using pulse width modulation (PWM). The power output by each channel is regulated by continuously switching the LEDs on and off at a rate of 500 Hz, a frequency above that detectable by the human visual system, which is known as the *flicker fusion threshold* and has a maximum value of around 60 Hz.¹⁰⁶ The ratio of the time the LED spends turned on compared to off in a single cycle is called the *duty cycle* and it determines the power output by each channel. This technique was adopted to control the LED's brightness as opposed to limiting the current supplied to each emitter because the shape of the spectral output, as well as brightness, of LEDs is known to vary with operating current.¹⁰⁷ Previous studies have reported complications in colorimetry applications of LED tunable light sources utilising current control.¹⁰⁸ PWM dimming does not affect the spectral output of the LED, as the LEDs are always driven at a constant current during the on period

of each cycle.

The current LED control electronics allow the intensity of each channel to be set in 151 steps (0-150), and the power control of each channel was tested for linearity by recording a spectrum of the output as each channel is set to four power levels: 30/150, 60/150, 90/150, and 120/150. Each spectrum was then integrated to determine the total optical power output and the results for the first (420 nm) channel are shown in figure 2.18, where it can be seen that the emission peak does not vary in wavelength with varying power output, and the total optical power output responds linearly to PWM control. The other seven channels were also tested and show an identical trend.

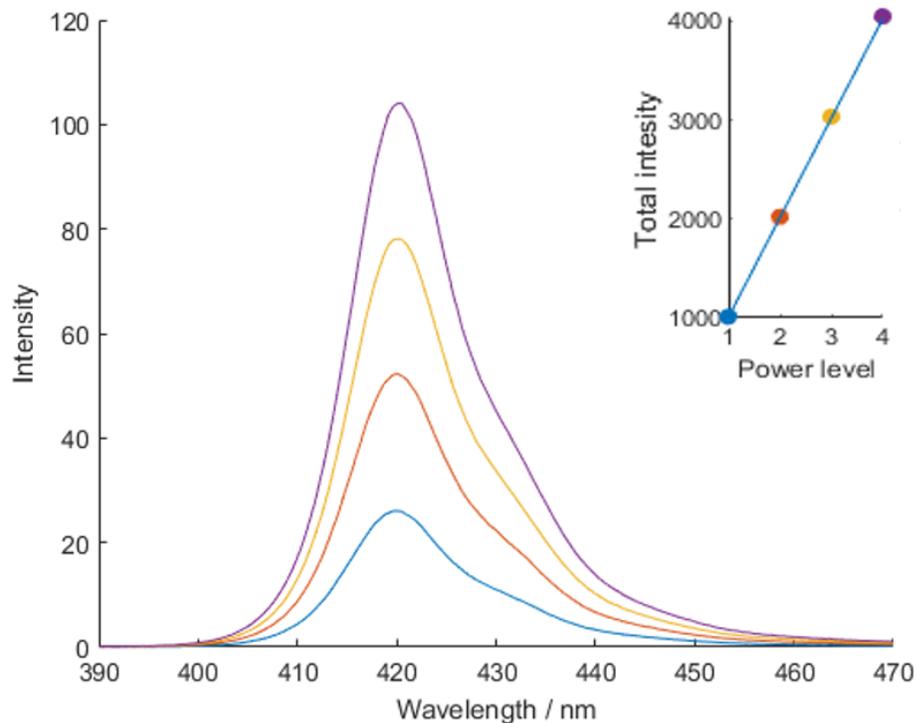


Figure 2.18 Optical power (arbitrary units) output from each channel of the spectrally tunable light source as it is set to four power levels: 30/150(1-blue), 60/150(2-red), 90/150(3-orange), and 120/150(4-purple).

Colourimetry studies

In working with spectra that contain sharp features, it is first important to consider the effects of observer metamerism. Biological differences between individual human

observers leads to slightly varied colour perception of identical stimuli, which is further exacerbated when the SPD of the stimuli contains sharp features.¹⁰⁹ For colourimetry, this usually results in a discrepancy between visual metamers and coordinates calculated using the CIE colour matching functions. A recent study using LED-based spectrally tunable light sources has concluded that even the most recent colour matching functions published by CIE were not good enough to render metameric light sources without clear visible differences in perceived colour by different observers.¹¹⁰ The perceived colour differences were attributed to the fact that the CIE colour matching functions were developed from an average of a set of observers, and for accurate colour matching each observer would effectively require their own set of colour matching functions. This effect can also be seen in digital photography where different image sensors have Bayer filters of varying transmission profiles, leading to differences in colourimetric response which cannot fully corrected for in software. When the neodymium samples used in this study were illuminated with a broad-spectrum white LED light source, the colour of the samples perceived by two different cameras (2.19) shows significant variation while the colour of the colour-standard card under identical conditions shows little difference, with the exception of the blue tones. Moreover, the intensity of light falling onto the eyes determines the ratio of which type of cells, rods or cones, dominate visual perception. These effects make it difficult to accurately describe the appearance of such samples, however, observer metamerism failure is typically minor when compared to the illuminant metamerism caused by sharp the SPD of the light source which was demonstrated in the previous section. Chromatic adaptation mechanisms also provide a margin for some unnatural illumination conditions to still appearing ‘white’.



Figure 2.19 Neodymium chloride samples illuminated by broad-spectrum white LED illumination photographed with different cameras: Sony a6000 (left), and Oneplus 5 featuring a Sony IMX 398 sensor (right). The in-camera white balance was set from the white tile on the QP203 colour standard card.

The suitability of LEDs used in the constructed light source was confirmed by looking at the appearance of the sample under illumination from individual channels of the light source, which was captured using digital photography and presented in figure 2.20. The exposure was normalised against the the brightness of the middle grey colour standard tile included above each image. It can be seen that for the red component, the samples appear significantly brighter under the 655 nm channel than the 590 nm channel; the same pattern is observed for the 545 nm / 515 nm channels forming the green component, as well as the 440 nm / 470 nm channels of the blue component. In the blue and green components, the colour of the light source shows little variation between the channels; while the red component shows an orange hue for the 590 nm channel, which is expected and will need to be compensated when combining the channels to form white light sources.

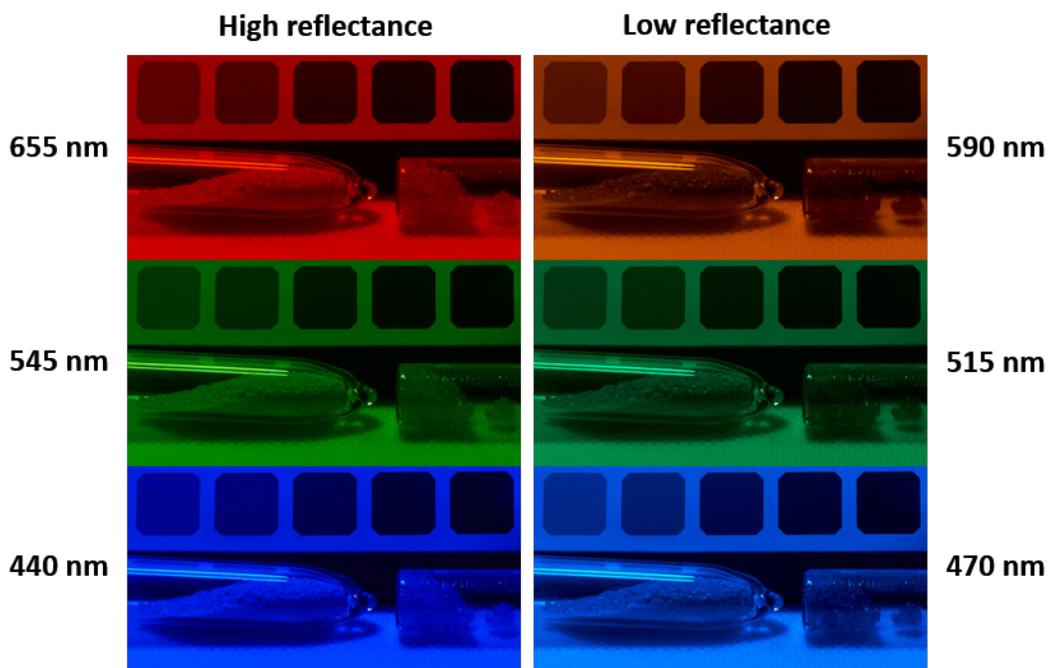


Figure 2.20 Neodymium chloride (anhydrous-left, hydrate-right) photographed while illuminated by different channels of the spectrally tunable light source. Images have been recorded with daylight white balance and exposure normalised to the 50 % grey reflectance standard (central grey tile above each image).

The tunable light source was used to create white light by mixing a single channel from each red, green and blue component. Feed-back from the spectrometer was used to simultaneously calculate the CIE xyz coordinates using the CIE 2° (2006) colour matching functions,¹¹¹ followed by a conversion to the $L^*a^*b^*$ colour space with D65 as the white point. The channel intensities were then manually tuned to get the a^* and b^* coordinates as close to zero as possible, while using the CIE Y parameter to measure the overall luminance and keep the total apparent brightness of the light source consistent between metameric combinations.

The appearance of neodymium chloride was tuned to primary colours of light (red, green and blue) by firstly selecting a single channel from the three components to have a high reflectance and the remaining two channels low reflectance, followed by the above procedure to balance the channels in the correct proportions to achieve white light of consistent luminosity. This is demonstrated in the case of the red appearance where

the channels chosen to form white light were: 655 nm, 515 nm and 470 nm. Secondary colour appearance (yellow, magenta and cyan) could also be made by selecting two channels to have high reflectance and the other low reflectance. For example, in the case of yellow the selected channels were: 655 nm, 545 nm and 470 nm. A further two combinations were also made where the channels were selected to all have high reflectance and all with low reflectance, causing the lightest and darkest possible appearance of the neodymium chloride sample. A total of eight illumination conditions were created and will be referred to as **1-8** in this section of work.

Images of the samples were captured under each type of illumination, with the in-camera white balance fixed to daylight, the same white point as that used in conversion to $L^*a^*b^*$ space when balancing the colour of the light source. The results can be observed in figures 2.21 & 2.22 where it can be seen that the appearance of neodymium samples is indeed altered to appear as both primary and secondary colours. These results confirm the possibility of tunable illuminant metamerism, and demonstrate that the perceived colour of an object can be entirely controlled by the spectral profile of the illuminant.

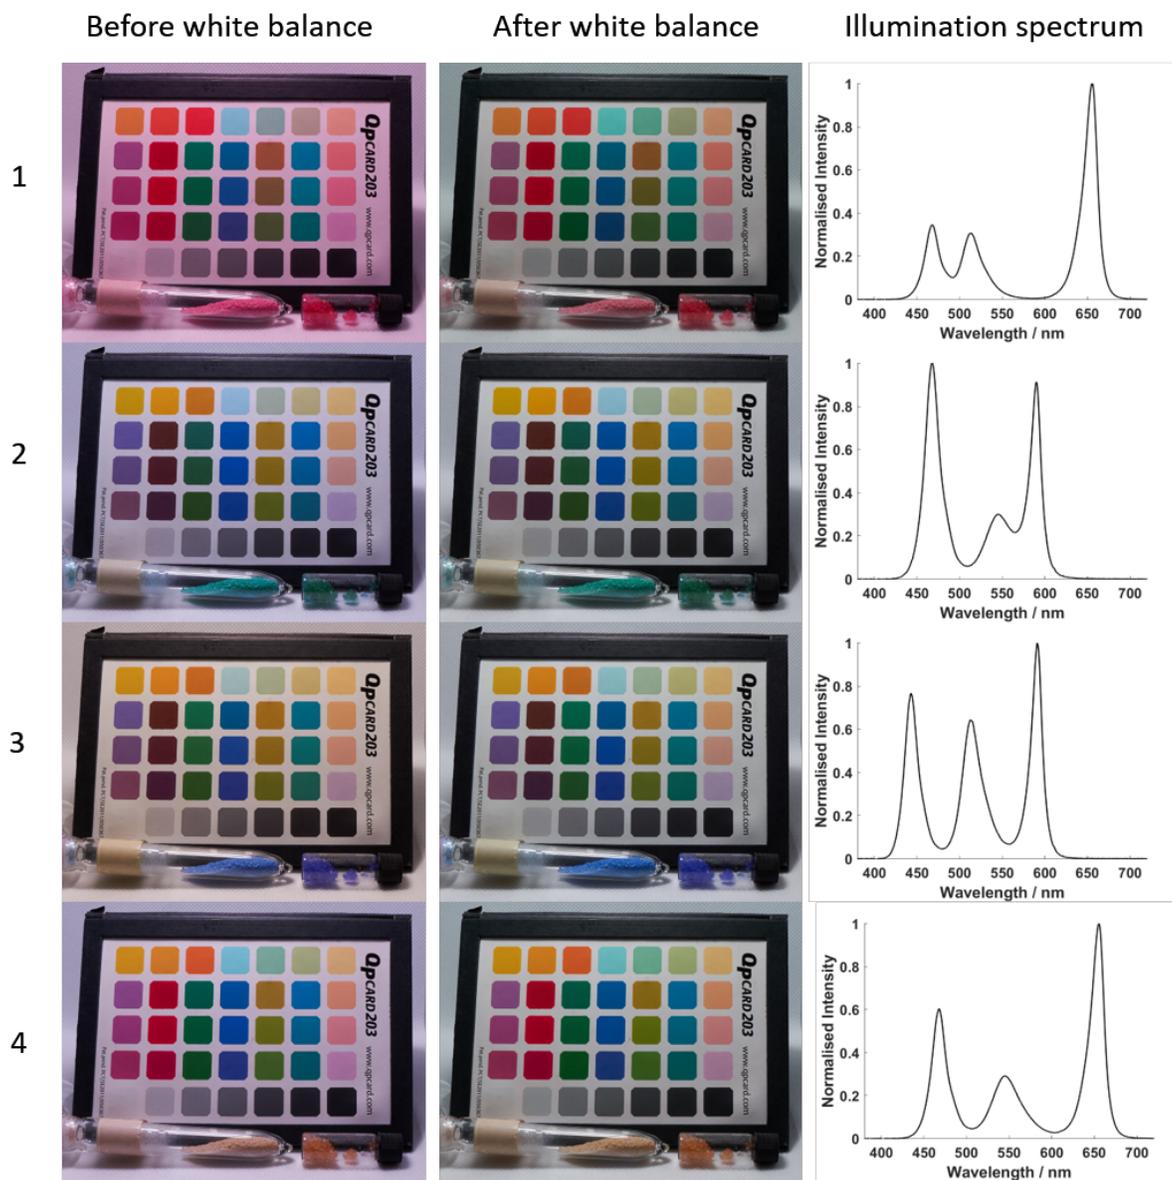


Figure 2.21 Neodymium chloride samples (anhydrous-left hydrate-right) photographed with illumination from spectrally tunable light source set to different spectral profiles shown in the right-hand column. Images have been recorded with daylight white balance and white balance corrected using the white tile of the QP 203 colour standard card.

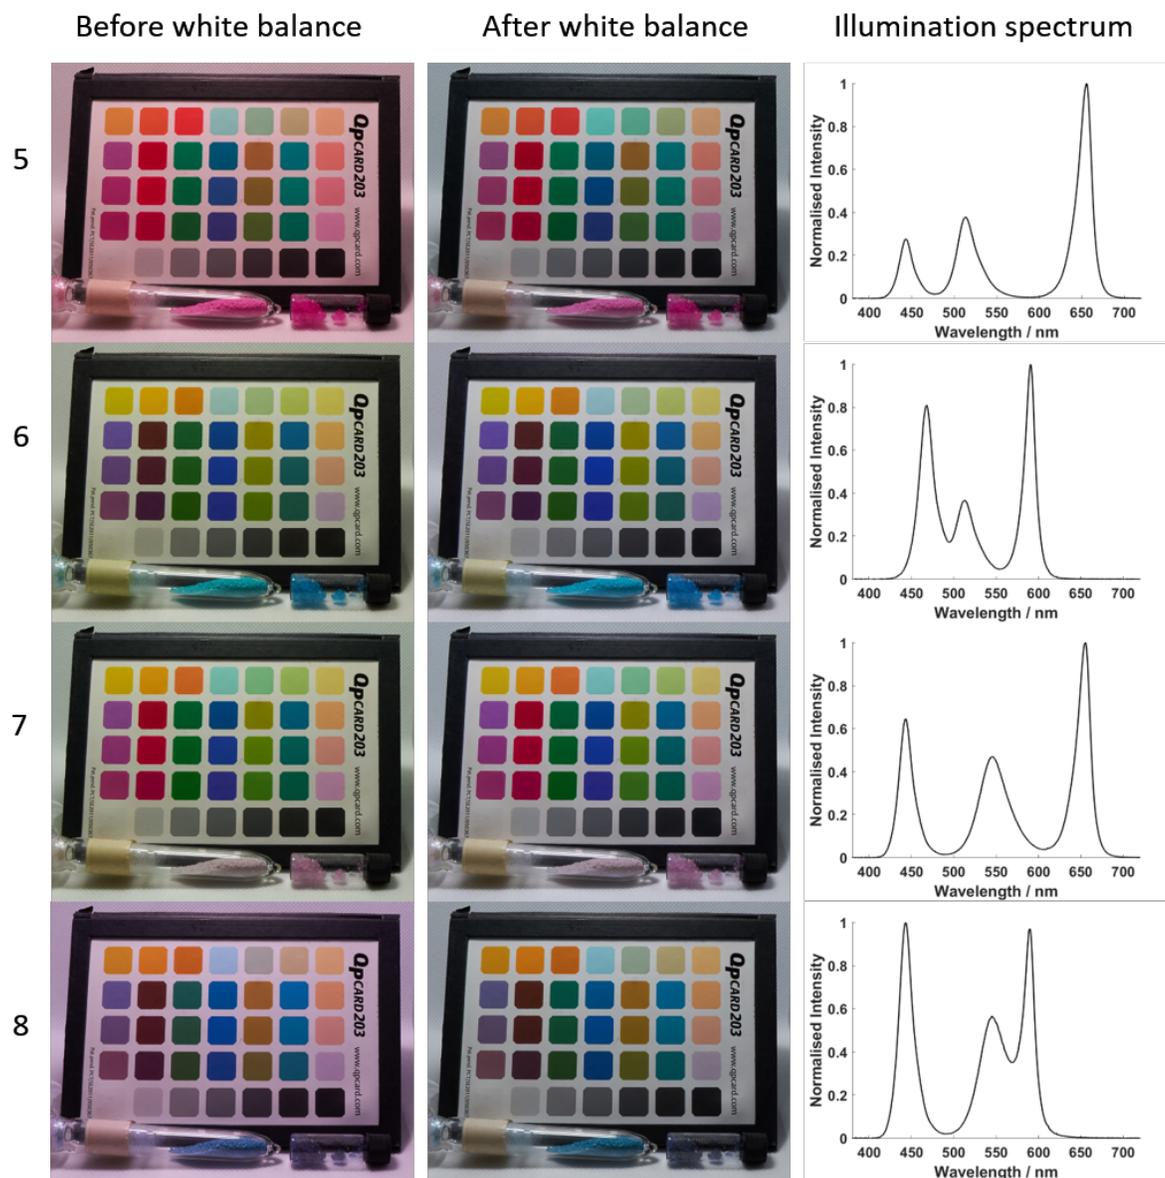


Figure 2.22 Neodymium chloride samples (anhydrous-left hydrate-right) photographed with illumination from spectrally tunable light source set to different spectral profiles shown in the right-hand column. Images have been recorded with daylight white balance and white balance corrected using the white tile of the QP 203 colour standard card.

The colour of the light source set to any single illumination condition was found to vary as described by different human observers or imaging sensors. In the case of **1** and **5** the light source appears to have a red hue, although the neodymium samples still adopts the desired red and magenta appearance respectively; scenarios **6** and **7** also show a slight green hue in the illuminant. These discrepancies are confirmed by human observation and likely arise due to the imperfect nature of the CIE colour matching

functions used in balancing the colour of the light source. Colorimetric analysis of the illumination spectra shows CIE xy coordinates (figure 2.23) scattered closed around the intended D65 white point and do not reflect the off-white hues captured in the images. In practice it was found impossible to balance the colour of the light source ‘by eye’ to achieve indistinguishable metamers due to chromatic adaptation effects, which suggests that the errors introduced by imperfect colour balancing may be un-noticeable if they are sufficiently subtle. To account for chromatic adaptation and estimate what the appearance would look like to an adapted observer, the images have been white-balanced using the colour of the white tile on the colour standard card used as the reference white point (column 2 in figures 2.21 & 2.22).

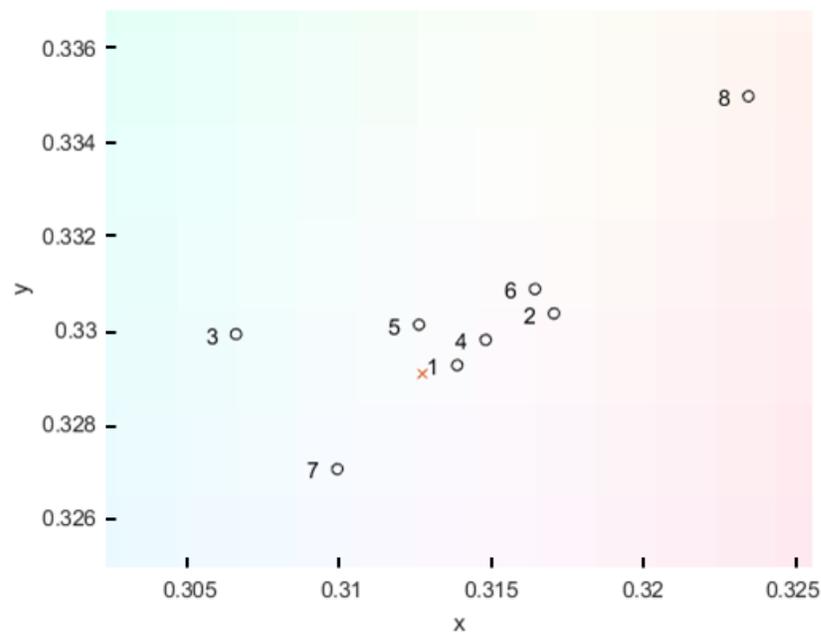


Figure 2.23 CIE 1931 xy coordinates of spectrally tunable light source set to illumination profiles 1-8 calculated from illumination spectra displayed in figures 2.21 & 2.22. The red cross pertains to the D65 white point. The colour of the background has been set to approximately represent the associated coordinates in the xy colour space.

The quasi-metameric illumination conditions do not exclusively influence the colour of neodymium chloride, the appearance of the colour standard card is also found to vary between different illumination conditions. To quantify these differences in appearance,

a set of 6 tiles representing the primary and secondary colour were selected and their colour analysed in the $L^*a^*b^*$ colour space (figure 2.24). The coordinates were extracted from the RGB values in the images and converted into the $L^*a^*b^*$ colour space using the colour of the white tile in each scenario as the white point. The results are displayed as projections on the a^*-b^* plane, with an image captured using high CRI LEDs as the illuminant for reference, which will be referred to as illumination condition **9**. While all of the coloured tiles show a degree of illuminant metameric failure, the difference in appearance is far more pronounced for the red, yellow and magenta tones. The red colour standard shows a particular pattern of two clusters dividing illumination conditions **1**, **5**, **4** and **7** with a lighter and more saturated appearance than **8**, **3**, **2** and **6**, with the reference condition **9** falling in-between the two clusters. This can be explained by the common factor of the red component in the tunable light source: when the 655 nm LED is used at the source of red, the red tones appear brighter and more saturated compared to broad spectrum illumination (**9**); whereas when the 590 nm LED is used as the source of red light, red tones appear marginally darker and less saturated relative to the reference. This is the same effect that causes the green appearance of neodymium chloride under fluorescent lighting, and occurs due to the narrow emission spectra of the LEDs; the effect is particularly prevalent in the red tones as the L and M cone cells in human eyes have peak sensitivity weighted towards lower wavelengths (figure 1.7), meaning the visual system can better resolve colours at the red and green portions of the spectrum.¹¹² The high degree of metameric failure seen in the yellow and magenta tones can also be attributed to the same cause, as they all require high reflectance in the red portion of the spectrum.

Poor rendering of red tones by the tunable light source constructed in this study is not unusual in lighting. Good colour rendering of high saturation red tones is especially challenging even with broad-spectrum artificial illumination, for this reason deep-red colour standards are sometimes omitted from CRI calculations that describe the quality of white light sources to yield inflated values.¹¹³ This could be reduced by exchanging

the 655 nm LED (chosen to have high reflectance by neodymium chloride) with a red LED lower wavelength, the high reflectance of the sample at 620 nm suggests that this could be a good alternative.

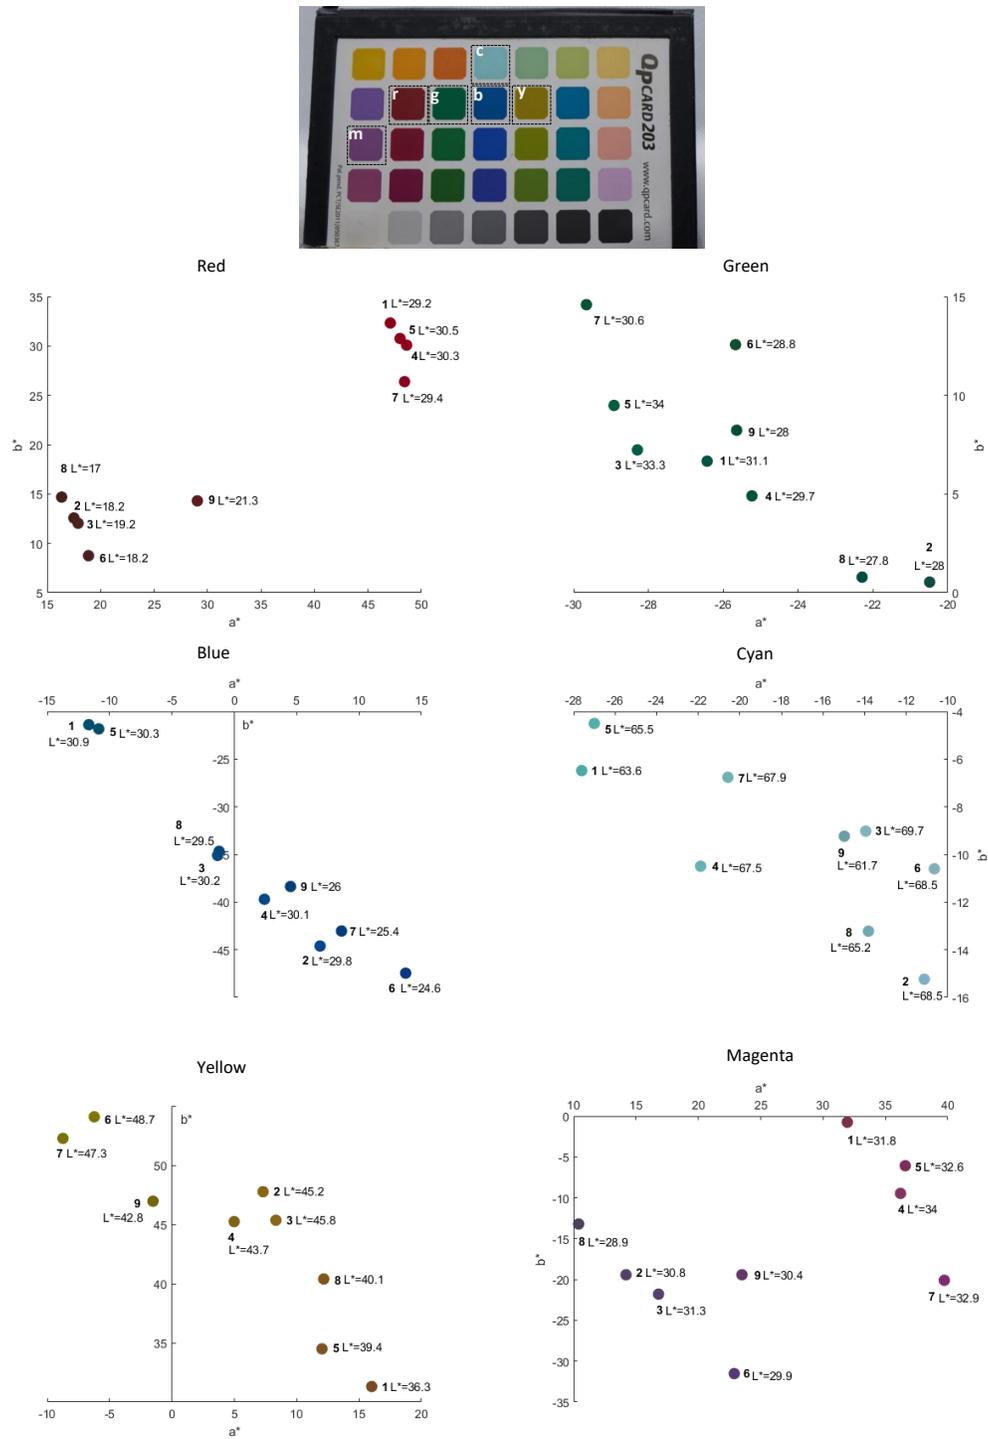


Figure 2.24 $L^*a^*b^*$ coordinates calculated by image analysis of tiles on a QP 203 colour standard card under illumination from led tunable light source set to illumination conditions 1-8 in figures 2.21 & 2.22 (9 pertains to broad-spectrum LED illumination). Each set of plots describes the appearance of a single tile of the QP 203 colour standard card, shown at the top with the letter denoting: red, green, blue, cyan, yellow and magenta, under different metamers of white light. The label next to each point describes the appearance of the corresponding colour tile under illumination conditions 1-9. The white point has been taken from the colour of the white tile in each scenario. The colour of each point is a conversion of the $L^*a^*b^*$ coordinate to the sRGB colour space.

The tuned colour of the anhydrous neodymium chloride sample can be further examined by looking at the spectral profile of the light reflected from its surface under each lighting condition. This can be calculated as the product of the spectral power distribution of the light source and the reflectance profile of the sample. The resulting spectra are shown in figure 2.25. From the primary colours: **1** and **3** are mainly composed of a single dominant peak, while **2** shows two broad bands in the blue-green and green-yellow regions of the spectrum. This is likely a result of the incomplete overlap of the 545 nm and 470 nm LEDs with the high and low reflectance bands of the neodymium chloride sample at 550 nm and 475 nm respectively, which limits the intensity or maximum saturation of the green appearance. From the secondary colours: **4** and **5** are dominated by the reflectance of the 650 nm LED, which must be the most intense due to the low sensitivity of the human visual system in this region of the spectrum. **4** is designed to reflect little blue light yet shows a larger than expected reflectance in the blue region, which was also observed in condition **2**; this has the effect of lowering the saturation of the yellow appearance and suggests that the 470 nm led used in this case is not ideally suited to the application. **6** shows a good spectrum consisting of two strong peaks as designed, which is congruent with the bright cyan appearance of the sample.

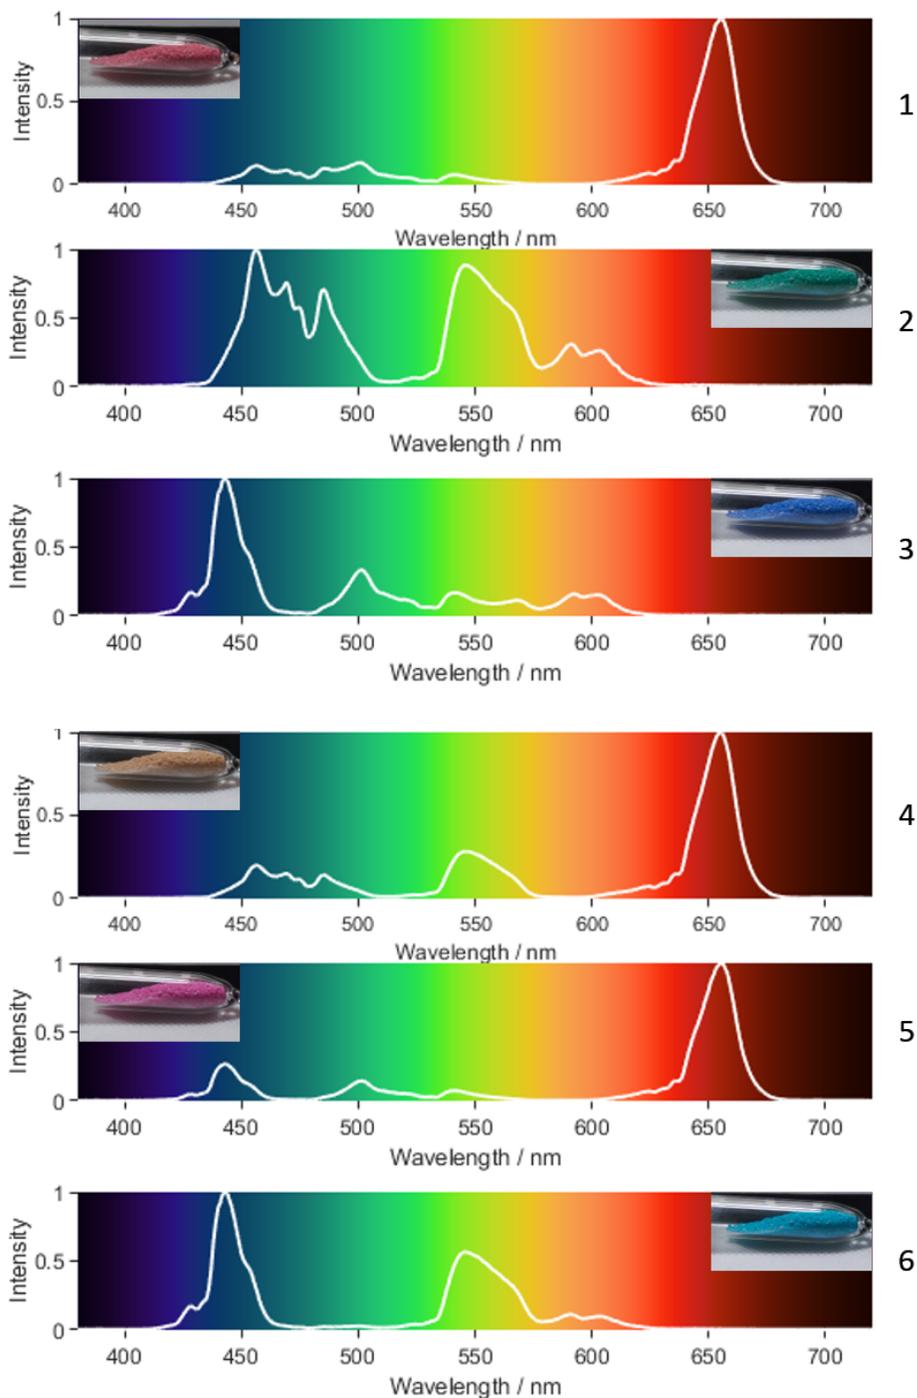


Figure 2.25 SPD of light reflected from the surface of anhydrous neodymium chloride calculated by the product of its reflectance spectrum and the SPD of the light source tuned to illumination conditions 1-6. Picture inset on each spectrum shows the appearance of the anhydrous neodymium chloride sample under each illumination condition. The background has been coloured as to approximately represent the perceived colour of each portion in the visible spectrum.

Colorimetric analysis can be extended further by applying colour matching functions to the spectral profiles reflecting from the surface of the anhydrous neodymium chloride sample under each illumination condition. The resulting $L^*a^*b^*$ coordinates have been plotted as projections onto the $a^* b^*$ plane in figure 2.26. The calculated coordinated closely match the colours observed with digital photography, and the appearance of the samples covers a large area in $L^*a^*b^*$ colour space. The maximum saturation of the blue and red hues which constitute a^* values almost as high as 50, and b^* values nearly as low as -40, is found to be far greater than that of the yellow and green hues; which only have a magnitude around 20 in a^*b^* space. The maximum saturation of the yellow and green LEDs is likely limited by the high reflectance of the 470 nm LED as well as the yellow colour of the 590 nm LED used as the red component in **2**, **3**, **6** and **7**.

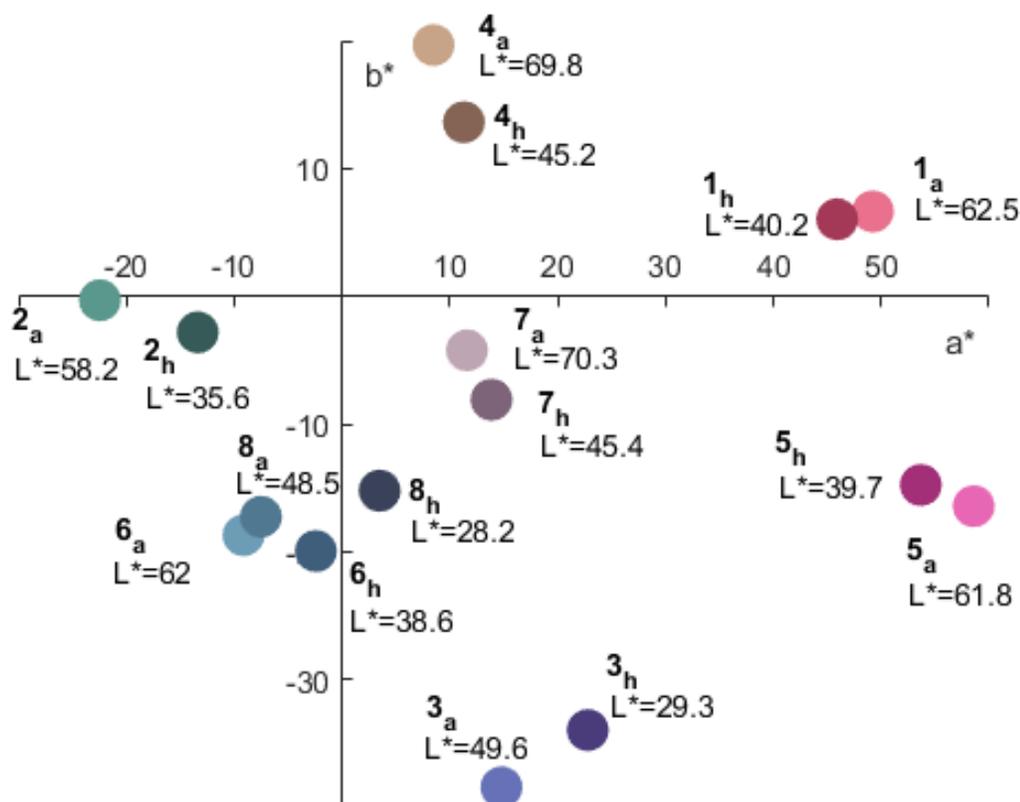


Figure 2.26 $L^*a^*b^*$ plot showing calculated colour appearance of two samples of neodymium chloride, anhydrous and hydrated, under illumination from led tunable light source set to illumination conditions 1-8 demonstrated in figures 2.21 & 2.22. The label next to each point describes the illumination condition with subscript 'a' denoting anhydrous neodymium chloride and 'h' denoting neodymium chloride hexahydrate. The colour of the points is a conversion of the coordinates to the sRGB colour space.

Illumination conditions 7 and 8 were designed to control the lightness of the sample. This is described by the L^* coordinate, which shows that a single sample of neodymium chloride has its lightest appearance under illumination condition 7, and the darkest appearance when illuminated by 8 as expected.

In condition 7, the appearance of the sample is grey with a pale purple hue, while 8 appears dark blue. Ideally, these would be white and black respectively, and the chromatic shift shows that: the LED's chosen to have a high reflectance with neodymium chloride are reflected in similar proportions, with the 650 nm channel being reflected marginally more than the 545 and 440 nm channels; and from the LEDs chosen to have low reflectance, the 470 nm channel is reflected the most. The intensity of visible light

reflected from the surface of the sample can be further defined using a different CIE metric called *light reflectance value* (LRV), which describes the percentage of visible light reflected from a surface accounting for the sensitivity of the human visual system.¹¹⁴ LRV for anhydrous neodymium chloride under illumination conditions **1** - **8** was calculated (table 2.1) using the formula:

$$LRV = \frac{\int_{380nm}^{750nm} I(\lambda)R(\lambda) d\lambda}{\int_{380nm}^{750nm} I(\lambda) d\lambda} 100\% \quad (2.2)$$

where I is the spectrum of the light source and R denotes the reflectance of the sample. The LRV values for the sample under **7** and **8** are the highest and lowest from the entire set of illumination profiles respectively, with the sample reflecting more than twice the amount of light under **7** compared to **8**. It is also found that **1**, **4** and **5** also show a relatively high LRV, and this is likely due to the 655 nm red LED which is used in each case and is largely reflected by the sample.

Table 2.1 Light reflectance values (LRV) of anhydrous neodymium chloride under illumination from spectrally tunable light source set to illumination profiles **1-8**.

	Illumination profile							
	1	2	3	4	5	6	7	8
LRV (%)	41.8	25.4	26.6	42.5	44.5	34.1	46.6	20.9

The established illumination conditions create an appearance of the neodymium chloride samples which is the most saturated hue of the primary and secondary colours possible using the constructed light source. The idea of mixing illumination conditions to fine-tune the appearance of the sample within these bounds is next investigated. Interpolated illumination profiles have been made by combining two illumination settings with a scaling factor, herein defined as ξ , which determines the proportion of each illumination profile present in the combined set. Profiles **1** ($\xi = 1$) and **3** ($\xi = 0$) were selected for an initial test, and four interpolated illumination profiles ($\xi = 0.2, 0.4, 0.6$ and 0.8) were tested by setting the intensity of each channel (I_i) in the light source

using a weighted sum:

$$I_i(\xi) = \xi I_1 + (1 - \xi) I_3 \quad (2.3)$$

where I_1 and I_3 represent the power level of each channel in illumination profiles **1** and **3** respectively. The resulting appearance of the neodymium chloride samples, as well as the illumination spectra, are displayed in figure 2.27. The colour of the samples gradually changes from red to blue, passing through shades of purple as the scaling factor ξ is decreased. The colour of the illuminant remains near-white as determined by the spectrometer (figure 2.28), although there is a red hue in **1** which is not reflected in colorimetric analysis. This issue has been previously addressed, and the red hue appears to decrease with ξ , suggesting that this is a systematic issue which can be resolved with better colour-balancing of the tunable light source.

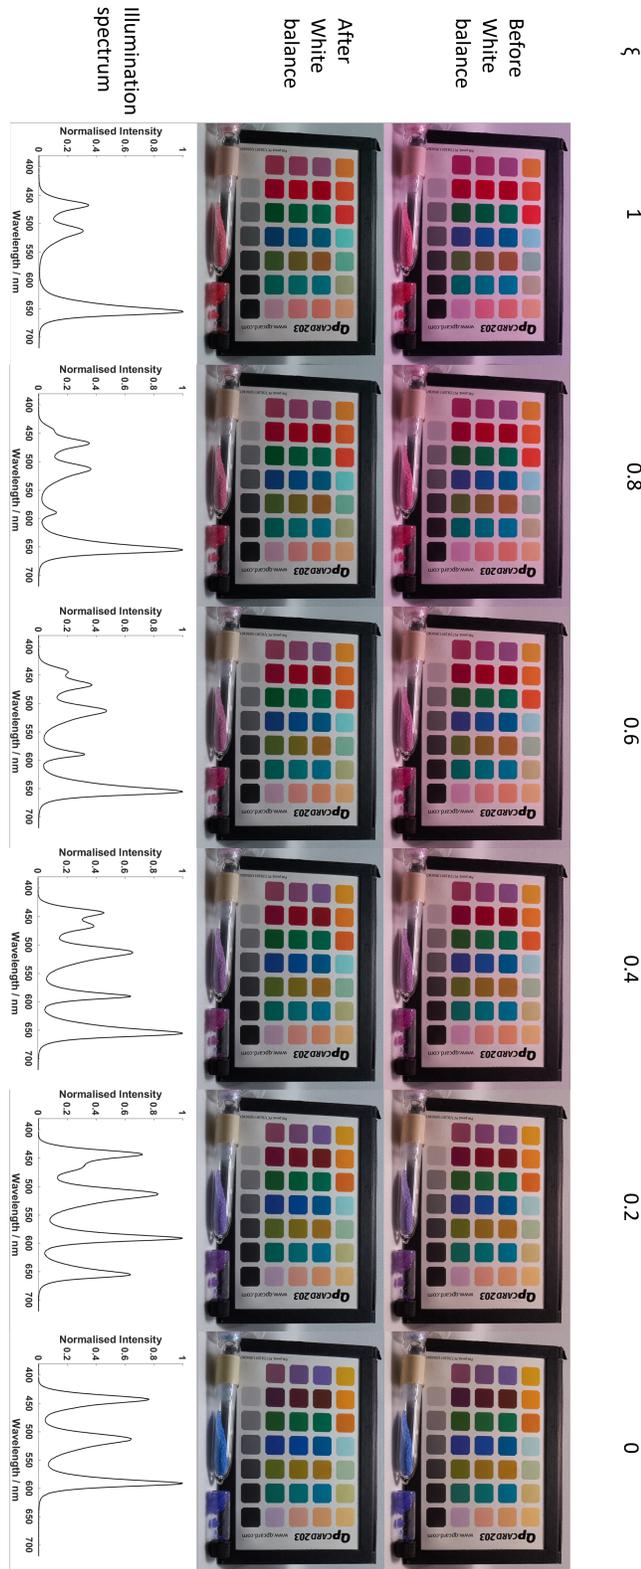


Figure 2.27 Appearance of neodymium chloride samples (anhydrous-left, hydrate-right) illuminated by spectrally tunable light source with spectral profiles interpolated from **1** and **3**, with scaling factor ξ determining the proportion of each profile in the total. White balance is set using the white tile on the QP203 colour standard card.

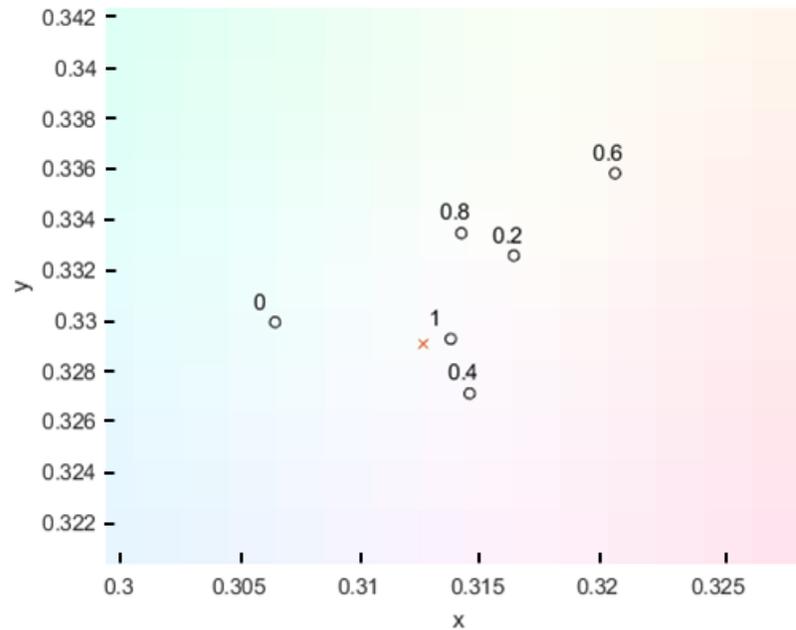


Figure 2.28 CIE 1931 xy coordinates of spectrally tunable light source calculated from interpolated illumination profiles between **1-3** demonstrated in figure 2.27, with number above each point showing the scaling factor ξ , which determines the proportion of each profile in the total. The red cross marks the D65 white point. The colour of the background has been set to approximately represent the associated coordinates in the xy colour space.

Colorimetric analysis of the neodymium chlorides appearance in the $L^*a^*b^*$ colour space, which was carried out from the spectrum of light source measured by the coupled spectrometer, shows a linear relationship between ξ and the apparent colour of the sample in $L^*a^*b^*$ space between the most saturated red and blue appearance. Similar behaviour is observed between any two illumination profiles, showing that this method of mixing illumination conditions via a weighted sum is valid for interpolating the appearance of neodymium chloride between the most saturated primary and secondary colours while maintaining near-white illumination.

The appearance of neodymium chloride illuminated by this spectrally tunable light source can therefore be considered as a polygonal gamut in the $L^*a^*b^*$ colour space bound by coordinates of the pure primary and secondary colour illumination conditions **1-6**. The resolution with which the appearance of neodymium chloride can be tuned to,

within this colour appearance gamut, is foremost limited by the resolution of the power control to each channel of the tunable light source; an appropriate level of stability of the set power level is also required to maintain a ‘white’ colour of the light source as the programmed appearance of the sample. In the current setup, the resolution of the power control to each channel is limited to 151 steps, which is likely a major source of error in mixing illumination profiles as calculations using equation 2.3 result in fractional values that must be rounded to integers to function with PWM power control. The resolution of the PWM power control can be increased with readily available integrated circuits, therefore there is no reason that this resolution cannot be increased until the parasitic thermal and chromatic effects of the LEDs, which are difficult to account for,¹⁰⁸ become the limiting factor.

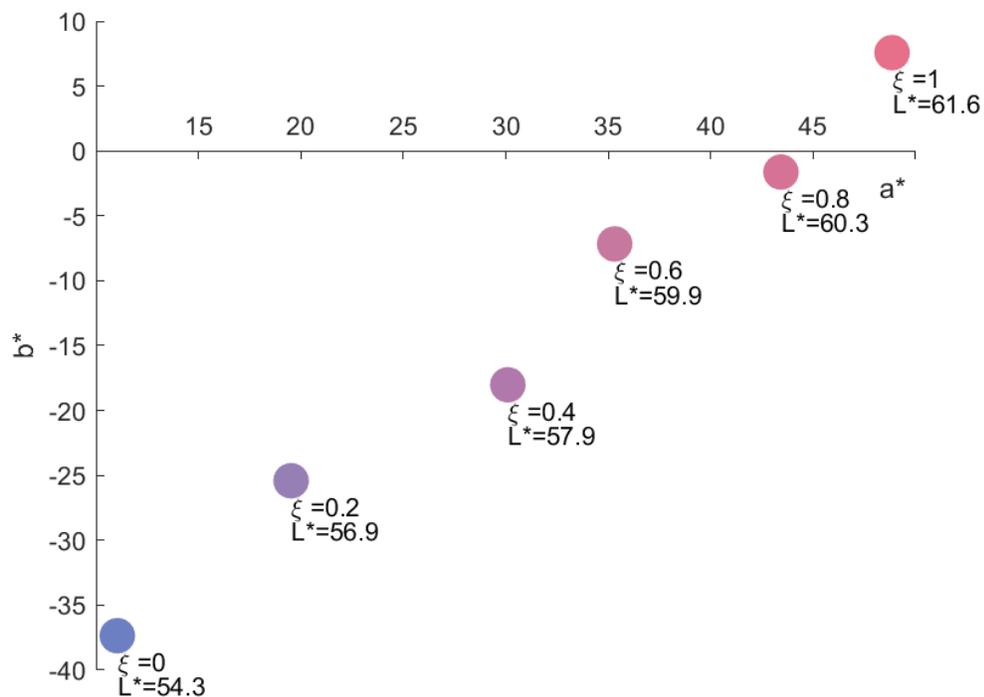


Figure 2.29 $L^*a^*b^*$ coordinates of anhydrous neodymium chloride under illumination from spectrally tunable light source set to interpolated illumination profiles between **1-3**, demonstrated in figure 2.27. The scaling factor ξ which determines the proportion of each profile in the illumination profile of the tunable light source is labelled next to each point. The colour of each point is a conversion to the sRGB colour space.

2.2.3 Modelling the effects of fluorescent whitening agents and hueing dyes on the appearance of fabric

Scope

The CIE colour appearance modelling that was rigorously tested in the previous two sections can also be applied to model the appearance of fabric, accounting for the natural degradation and subsequent colour correction during the laundry process (described in section 1.4). This section of work will focus on using semi-empirical computer modelling of the effects that yellowing, FWAs, HDs and lighting conditions have on the appearance of fabric.

Clean fabric model

To quantitatively describe the appearance of fabric free from degradation or additives introduced during washing, a reflectance spectrum of new cotton fabric ($R_{fab}(\lambda)$) was measured using DRS (figure 2.30). The fabric was stained with FWAs from the manufacturer, and to exclude this contribution from the reflectance spectrum a 400 nm long-pass filter was added into the probe beam, effectively blocking the excitation of the FWAs. The absorption of the FWA is the major feature in the reflectance spectrum, but this occurs in the UV-deep blue part of the spectrum (below 400 nm), therefore has very little influence on colour appearance. The product of $R_{fab}(\lambda)$ and the SPD of the illuminant ($I_{ill}(\lambda)$) chosen to assess the fabric's appearance under gives the SPD of light reflected from the fabric and reaching the observer ($I_{obs}(\lambda)$).

$$R_{fab}(\lambda)I_{ill}(\lambda) = I_{obs}(\lambda) \quad (2.4)$$

The spectral function $I_{obs}(\lambda)$ can then be input into the CIE colourimetry system to calculate the $L^*a^*b^*$ coordinates of the fabric under a light source of choice, with the

white point set as the colour of the illuminant. For initial studies, the measured spectrum of mid-day sunlight (figure 1.11) will be used. This forms the basis of the fabric appearance model, and further properties of real fabrics can be introduced algebraically.

The output of this fabric appearance model is a set of $L^*a^*b^*$ coordinates which can be used in combination to assess trends, although each 3-dimensional set of coordinates is difficult to visualise. One solution is to translate the coordinates into the sRGB colour space, which can then be displayed via a screen, which usually have some level of colour-calibration. However, this would still result in the output of a single colour, and real fabrics are typically highly textured which instead leads to a small distribution of colours. In order to aid visualisation, a simple method was devised based on image-colour manipulation. First, an image of a textured fabric was taken next to a white block for reference (figure 2.30-bottom right), making sure that all pixels covering the fabric do not exceed the dynamic range of the camera. A mask of the textured fabric in this image was then created by a contrast threshold against the black background. The RGB coordinates of all pixels within the mask are then converted to $L^*a^*b^*$ coordinates, forming a point cloud in $L^*a^*b^*$ space. The centre of this point cloud is evaluated and the difference between the centre coordinate and the model output is used to calculate a translation vector, which can then be used to shift the entire point cloud in $L^*a^*b^*$ space. The result is that the $L^*a^*b^*$ coordinates output by the model set the centre of the image point cloud. Finally, a reverse conversion of each point to the sRGB colour space yields an image where the texture of the fabric is preserved, but the overall colour appearance is determined by the colourimetry model (figure 2.30-bottom left). In the simple case of clean fabric, the results appear very similar to the original appearance of the fabric sample.

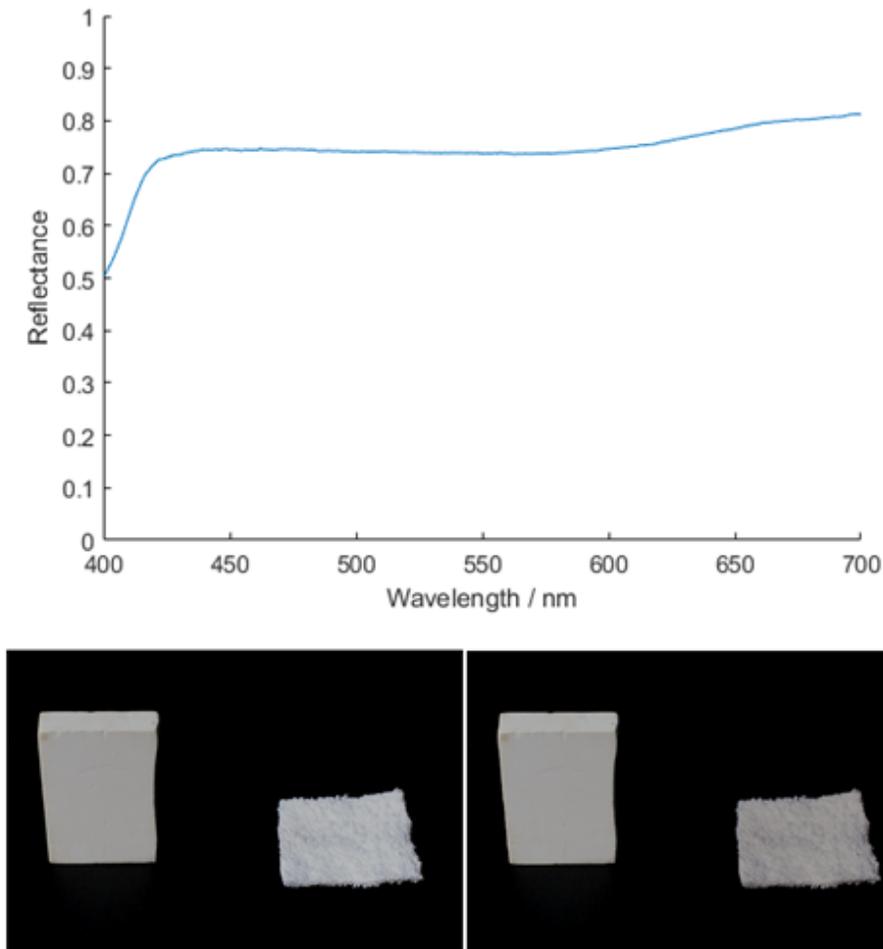


Figure 2.30 (Top) DRS spectrum of white cotton fabric, measured with 400 nm low-pass filter in probe beam. (Bottom-right) Raw image of cotton fabric swatch used to make aid for visualisation of fabric model results. (Bottom-left) colour manipulated image with colour of fabric swatch set by model output.

Yellowing model

Yellowing can be considered as an increasing concentration of blue-absorbing chromophores. The absorbance spectrum of these chromophores is likely to vary slightly from sample to sample due to the nature of yellowing. An empirical approach was therefore taken in modelling the effects of yellowing. A piece of paper free from FWA's which had been left in front of the office window and showed a significant degree of yellowing (primarily caused by months of exposure to UV light) was measured using DRS ($R_{yell}(\lambda)$). Paper has a reflectance spectrum which strongly resembles that of

cotton fabric, due to its cellulose composition.

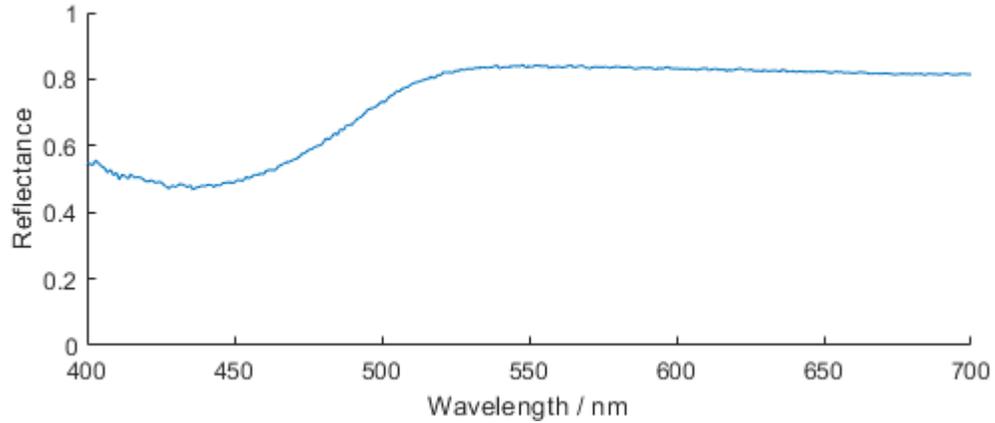


Figure 2.31 DRS spectrum of naturally yellowed paper used to model yellowing in fabric.

The main feature in the reflectance spectra of this yellowed sheet is a strong absorbance band in the blue region (figure 2.31), which peaks around 430 nm and is responsible for the yellow colour. This is introduced into the fabric model by attenuating the reflectance spectrum of clean fabric. The attenuation is carried out as a weighted sum of the clean fabric's reflectance spectrum, and that of the yellowed fabric in KM space, followed by the inverse KM function ($f'(R)$) on the sum:

$$f'(f(R_{fab}(\lambda)) + wf(R_{yell}(\lambda))) = R_w(\lambda) \quad (2.5)$$

where the coefficient w determines what proportion of the yellowed spectrum is included in the output, which is termed as the working reflectance spectrum ($R_w(\lambda)$) and can be used in place of $R_{fab}(\lambda)$ in equation 2.4 to calculate the colour appearance of yellowed fabric.

The effect yellowing has on the spectrum of light reflected from the fabric can be seen below for a practical range of w .

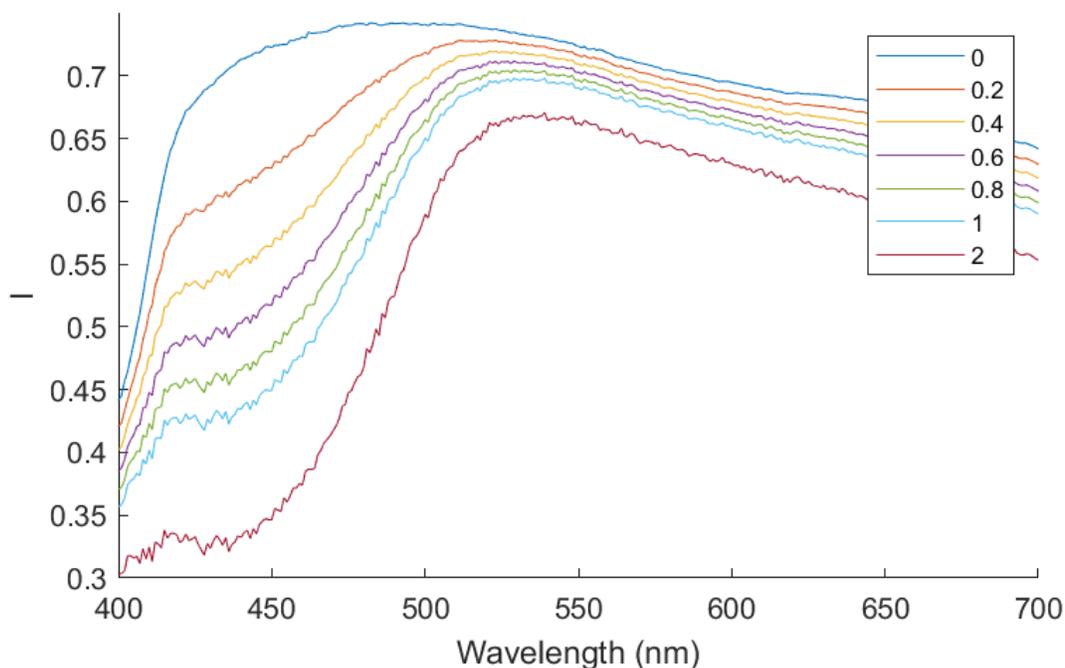


Figure 2.32 Simulated spectra of sunlight reflected from the surface of yellowed cotton fabric (I_{obs}) modelled using equation 2.5 with arbitrary intensity units for comparison. The w values used to construct each spectrum shown in the legend.

As expected, the spectra show the largest variation at 430 nm, where the absorbance of yellowed cellulose is greatest. The effects of yellowing in the green and red region are, however, very minor. Since the KM function is designed to scale linearly with concentration of absorbing species, this modelling approach should be valid within the range where the KM approximation is valid. In particular, KM models tend to perform poorly on strongly absorbing material, which is not the case in yellowed fabric.

The resulting simulated DRS spectra can then be used in conjunction with a spectrum of the illuminant (sunlight for initial trials) to calculate the colour appearance of aged cotton in $L^*a^*b^*$ space. Initial tests were carried out with value of $w = 0, 0.2, 0.4, 0.6, 0.8, 1$ and 2 (figure 2.33). It can be seen that increasing the yellowing coefficient causes a uni-directional shift in appearance towards a subjectively unappealing yellow hue, although the $L^*a^*b^*$ coordinates also contain a significant b^* contribution which increases with ‘yellowing’.

It is also observed that the L^* value decreases with yellowing, which results from the

decreased total light reflected from the sample.

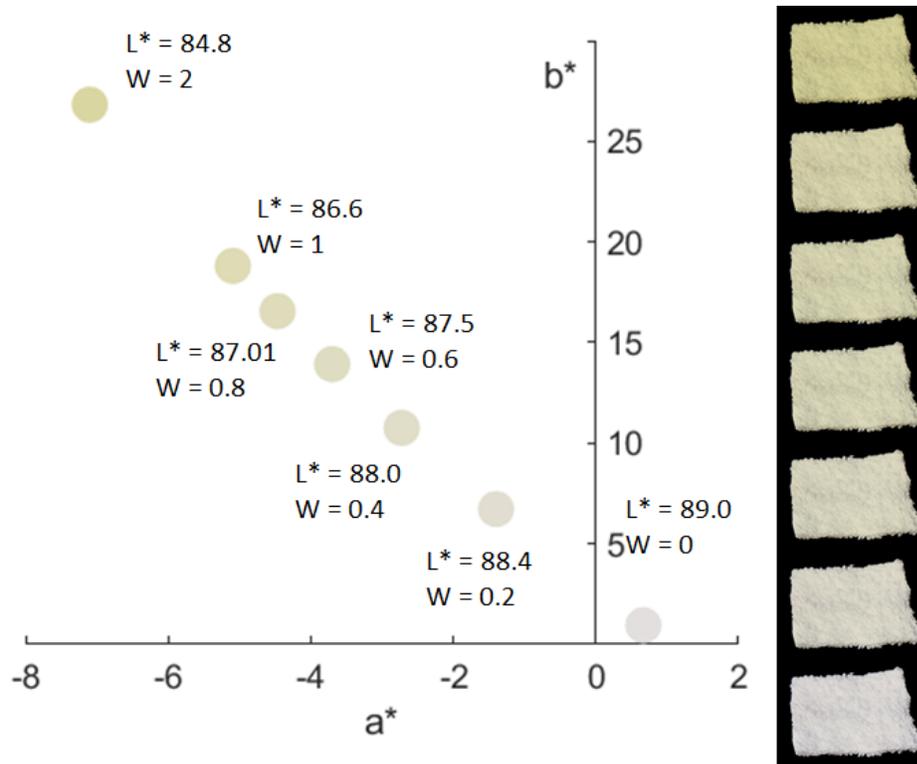


Figure 2.33 (Left) $L^*a^*b^*$ coordinates calculated from modelling yellowed fabric. w refers to the yellowing coefficient in equation 2.5. The colour of the points is a conversion to the sRGB colour space. (Right) Simulated appearance of corresponding yellowed fabric in descending w top-bottom.

The fabric visualisation model performs well in aiding the visualising of calculated $L^*a^*b^*$ coordinates, from the results it can be seen that values of w beyond 1 are rarely encountered in real-world laundry.

FWA model

To model the effect of FWAs on fabric, a fluorescence spectrum (I_{fl}) was recorded from the clean cotton swatch used as the model base, which was pre-stained with FWAs from the manufacturer (figure 2.34). A 365 nm LED was used as the excitation source, and as this is not a relative measurement, the recorded emission spectrum was correctly to ensure a linear response across the visible spectrum.

The measured fluorescence spectrum shows just how well-suited the proprietary FWA

technology used in fabrics is to combat yellowing when compared to the absorbance of yellowed paper (approximated with KM function in figure 2.34). The two spectra match almost perfectly and should combine to restore the appearance of yellowed fabric.

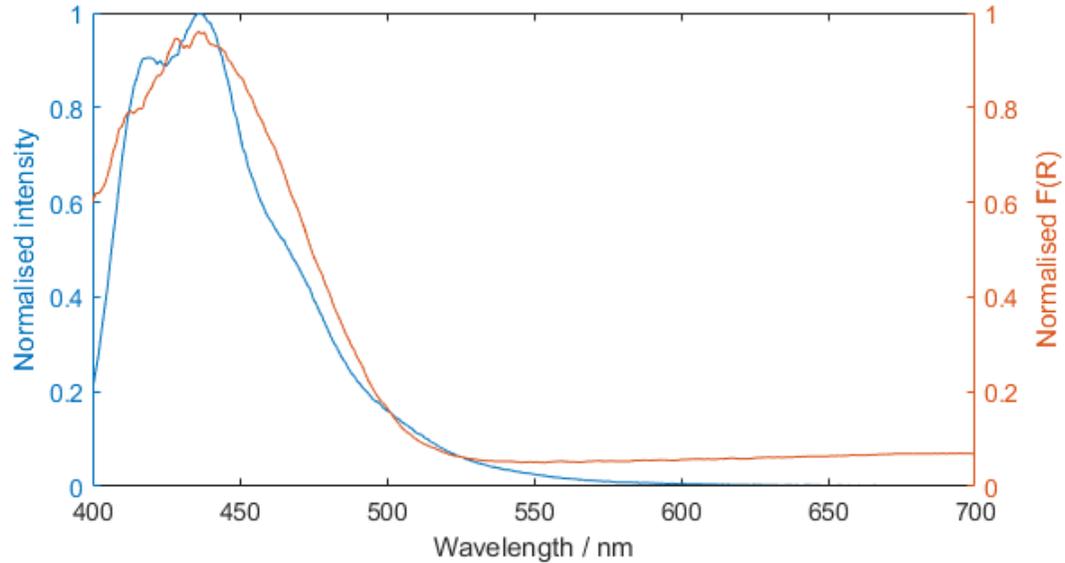


Figure 2.34 (Blue) Corrected fluorescence spectrum recorded from clean cotton fabric with excitation at 365 nm. (Orange) Absorbance of naturally yellowed paper, approximated using the KM function on DRS spectrum.

An arbitrary fluorescence coefficient (F) will be necessary when incorporating FWAs into the fabric model as the amount of fluorescence emitting from the sample is dependant on many variables. The wavelength dependant excitation product of the FWA and the light source, as well as the concentration of FWA at the point under observation are just a few examples.

This approach, however, does not account for the complex relationship between the spectral output of the FWA fluorescence and the vertical distribution of the FWA through the fabric matrix. The complex relationship arises due to the diminishing power of the excitation light as it is scattered upon penetrating further into the fabric, as well as the wavelength-dependant and depth-dependant attenuation that the fluorescence is subject to before leaving the fabric and reaching the observer. These complexities have been somewhat reduced by using thin fabric with a black backing to

reduce the back-reflections reaching the observer.

Another consideration is that in reality FWA concentration cannot be increased continuously, eventually the molecules begin to aggregate, and the spectral profile of the emission as well as the quantum yield changes. In this model we assume that due to the high quantum yield of modern FWAs, they are in low enough concentration for practical values of F that aggregation is negligible.

A fluorescence coefficient allows us to track the appearance of the fabric with increasing FWA concentration, and in future can potentially be transcribed to a physical property. This has previously been the subject of research which concluded that the relationship is complex but prediction is possible.¹¹⁵

With these approximations in mind, the measured fluorescence spectrum can be incorporated into the fabric appearance model by adding it into the light reflected from the surface of the fabric (equation 2.4) via a weighted sum:

$$R_{fab}(\lambda)I_{ill}(\lambda) + FI_{fl} = I_{obs}(\lambda) \quad (2.6)$$

The magnitude of F is arbitrary and depends on the relative intensity of the illuminant. In this case, the illuminant has been scaled to have a unity luminosity, or perceived brightness, and the fluorescence spectrum has been scaled relative to the illuminant such that 0-1 forms a sensible scale for real-world fabrics. In the fabric model, FWAs can be added to fabric with any degree of yellowing.

For an initial test, the condition $w = 0.5$ was examined. Simulated optical brightener was added in units $F = 0, 0.2, 0.4, 0.6, 0.8, 1$, and the effect on $I_{obs}(\lambda)$ can be seen below:

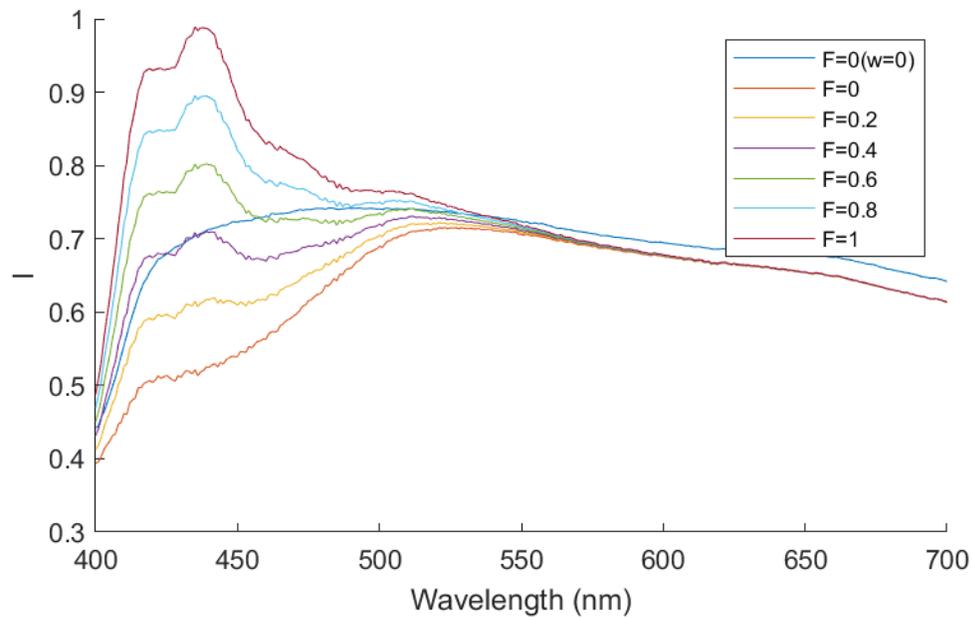


Figure 2.35 Simulated SPD $I(\lambda)$ of sunlight reflected from the surface of yellowed fabric (simulated with $w = 0.5$) stained with an increasing concentration of FWA. The units of intensity are arbitrary for comparison between spectra. The legend indicates the values of F introduced in equation 2.6.

The results clearly show the mechanism of FWAs operation, where the fluorescence makes up for the blue absorbance of yellow chromophores and restores a 'flat' colourless reflectance profile when the concentration of FWA is matched to the degree of yellowing. In this example, a near-matched condition is observed for values $w = 0.5$ and $F = 0.4$, with the resulting spectral density of light reflected from the surface resembling that of un-yellowed fabric ($w = 0$, $F = 0$). The effects of FWA can be quantified by further analysing the spectra in terms of their predicted appearance in $L^*a^*b^*$ space (figure 2.36).

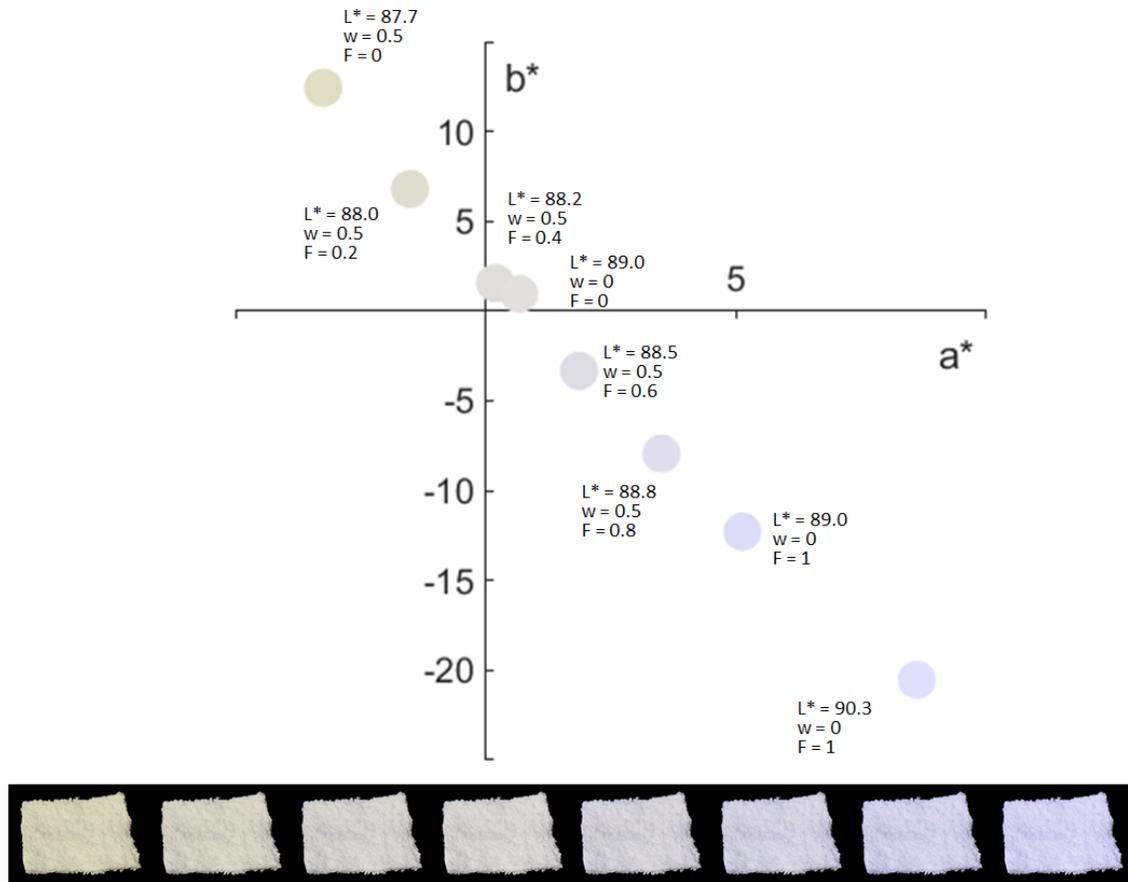


Figure 2.36 (Top) $L^*a^*b^*$ coordinates calculated from modelling yellowed fabric stained with a varying concentration of FWA. w refers to the yellowing coefficient in equation 2.5 and F to the fluorescence coefficient introduced in equation 2.6. The colour of the points is a conversion to the sRGB colour space. (Bottom) Simulated appearance of each corresponding modelled fabric condition, arranged left-right in order of increasing a^* coordinate.

The coordinates display a unidirectional blue shift in appearance on increasing the amount of fluorescence F added to the model, which is a near-perfect polar opposite to the appearance shift caused by yellowing. This originates from the overlap of the two spectra displayed in figure 2.34. The matched condition ($w = 0.5$, $F = 0.4$) appears visually indistinguishable (ΔE below 2) from unstained fabric, even matching in L^* value. The model also shows the effects of adding too much FWA. In the example cases, values of F greater than 0.4 result in a light blue appearance of the fabric. Moreover, addition of FWA to un-yellowed fabric causes a strong and un-natural blue appearance of the fabric. A useful property of the human visual system aids in extending the usable range of FWAs here. In perception of whiteness, small appearance shifts in

the blue direction are still perceived as ‘white’, in some cases even more ‘white’ than neutral tones as found in psychology experiments.⁸⁷ This blue-preference, combined with FWAs ability to increase lightness forms the basis for the ‘whiter than white’ claims often advertised.

Since the yellowing vector (w) was found to be a near-negative of the FWA concentration vector, it can be assumed that in this model, any practical values of w can be corrected for by a corresponding value of F . Modelling how FWAs correct fabrics appearance at different stages of yellowing shows that this is the case. A matching F was found for a practical range of w (0-1) where the resulting match is visually indistinguishable from un-yellowed fabric, and the resulting relationship is displayed in figure 2.37. The amount of fluorescence needed to correct for an increasing concentration of yellowing chromophores is found to decrease at higher levels of yellowing leading to a non-linear relationship. This is due to the imperfect spectral overlap of the yellowing chromophore absorbance and FWA emission shown in figure 2.34, and can be considered beneficial to FWAs function as small levels of fluorescence make the largest impact in correcting for the appearance of yellowing.

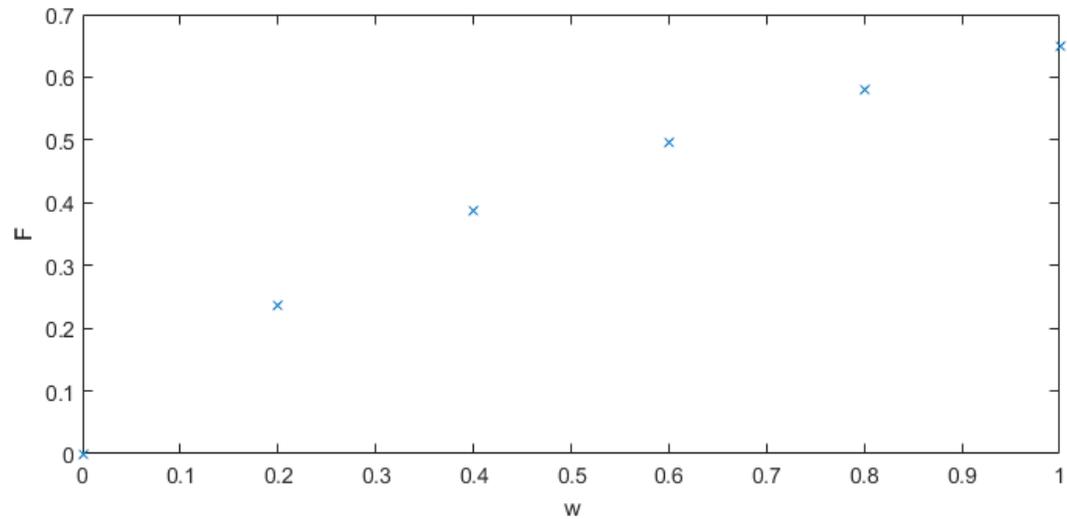


Figure 2.37 Simulated conditions where the appearance of fabric at various stages of yellowing is perfectly correct for by addition of FWA. Yellowing has been modelled using equation 2.5 with coefficient w plotted on the x -axis, and FWA action has been modelled using equation 2.6 with coefficient F forming the y -axis. Each point represents a modelled fabric condition where the correction by FWA of concentration F restores the appearance of fabric yellowed by w to the point where it is visually indistinguishable from unyellowed fabric.

Hueing dye model

To model the effect of HDs on appearance of fabric, a similar approach to that used in modelling of yellowing can be applied. A representative HD was found in literature as a reflectance spectrum ($R_{HD}(\lambda)$) of cotton fabric stained with monomeric v200 dye (a commercial HD).⁴⁶ The spectrum features a high reflectance peak coinciding with the absorbance caused by yellowing with heavier absorbance in the green and red regions of the spectrum, as demonstrated in figure 2.38.

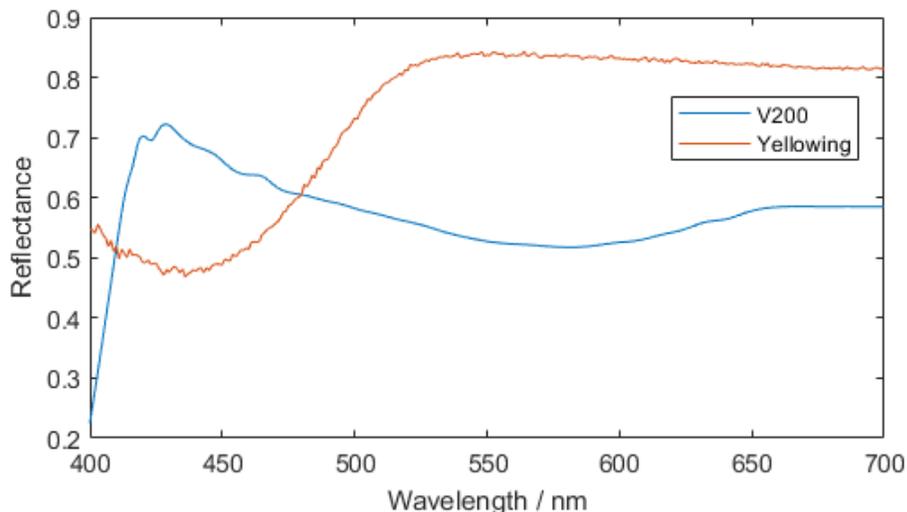


Figure 2.38 DRS spectrum of: cotton fabric stained with v200 HD (Blue),⁴⁶ and naturally yellowed paper for reference.

Since a HD can be considered as an increasing concentration of green/red absorbing species, the absorbance can be added to that of paper and yellowing via a weighted sum in KM space. This can be achieved by modifying equation 2.5 to incorporate the measured spectrum via a HD coefficient H :

$$f'(f(R_{fab}(\lambda)) + wf(R_{yell}(\lambda)) + Hf(R_{HD})) = R_w(\lambda) \quad (2.7)$$

This approach to modelling should again be valid within the approximations of KM theory, and in this case at low concentrations of dye. The HD coefficient H should scale linearly with HD concentration, due to the proportional nature of the KM function. HDs are prone to aggregation at high concentrations, which is not accounted for using this technique. Therefore H values of less than 1 will be used in modelling as the dye in the reference HD stained fabric ($H = 1$) is known to be free from aggregates,⁴⁶ making the approximations feasible.

The effects of HD on light reflected from yellowed fabric ($w = 0.5$) can be seen in figure 2.39 where light in the green/red regions, that are little affected by yellowing,

are the absorbed by the HD; which does not absorb in the blue region, restoring a ‘flat’ reflectance profile with around 20% lower reflectance than un-yellowed fabric.

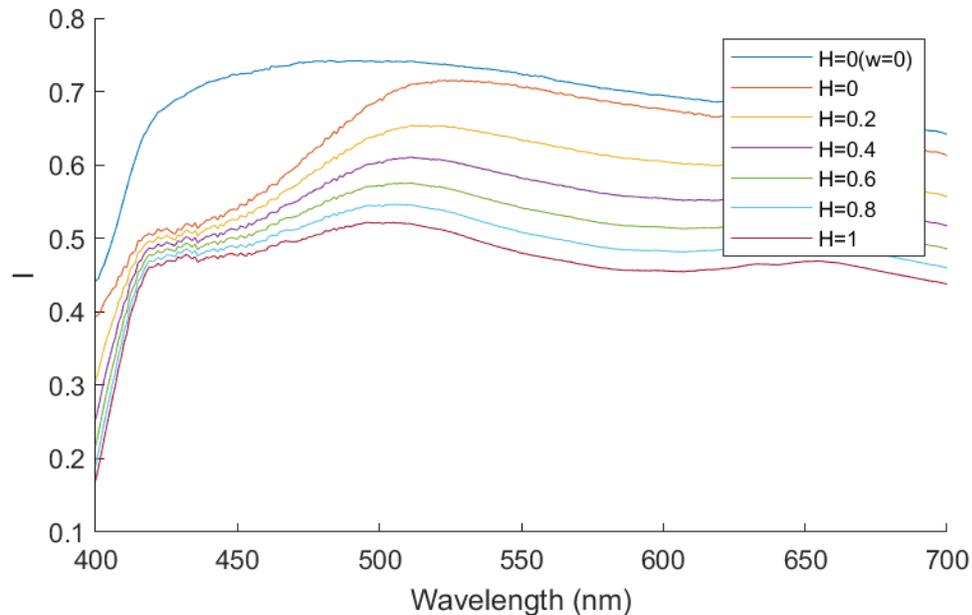


Figure 2.39 Simulated spectra of sunlight reflected from the surface of fabric showing a varying degree of yellowing. The spectra are modelled using equation 2.7 with $w=0.5$ unless otherwise stated and the corresponding HD coefficient shown the legend. Intensity is presented in arbitrary units for comparison between spectra.

A flat reflectance profile yields a neutral appearance, but the lower reflectance manifests as a dull-grey appearance of the fabric. This can be seen by further analysis in the $L^*a^*b^*$ space (figure 2.40), where the action of increasing HD concentration on yellowed fabric increasingly shifts the appearance towards neutral grey at the expense of decreasing L^* value. In condition $w = 0.5$ and $H = 1$, the fabric has no discernable colour, yet the appearance compared to un-yellowed fabric is visually darker and 13% less in lightness. This is of course less favourable than the correction from optical brighteners, however, HDs do not require the light source to contain an invisible UV or even deep violet component.

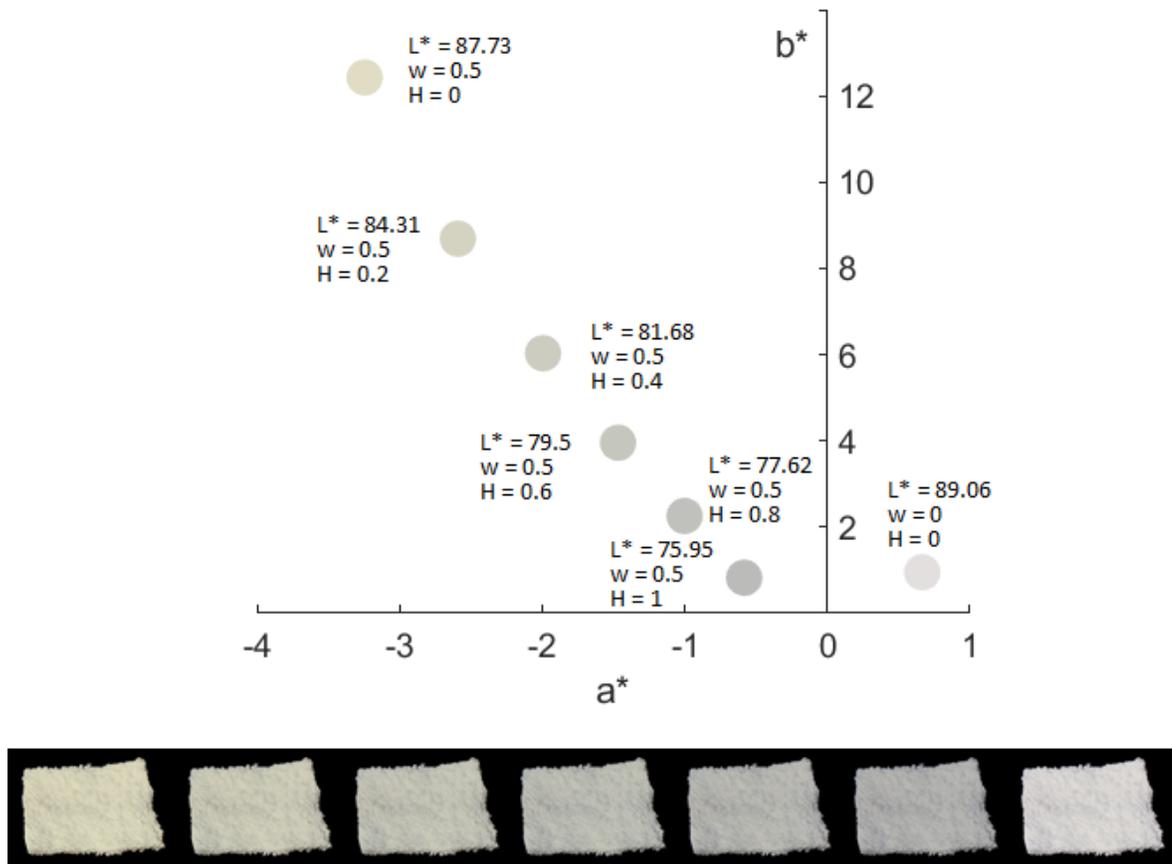


Figure 2.40 (Top) $L^*a^*b^*$ coordinates calculated from modelling yellowed fabric stained with a varying concentration of HD. H refers to the yellowing coefficient in equation 2.5 and H to the HD coefficient introduced in equation 2.7. The colour of the points is a conversion to the sRGB colour space. (Bottom) Simulated appearance of each corresponding modelled fabric condition, arranged left-right in order of increasing a^* coordinate.

A matched condition, like the one observed for the case $w = 0.5$ and $H = 1$, can be found for any value of w , provided that HD is added in enough concentration. An analysis into the concentration of HD needed to neutralise the yellow hue introduced from a varying degree of w was carried out. This was done by minimising the b^* value using least squares fitting, and the results are displayed in figure 2.41. It is found that the concentration of HD needed to neutralise yellow chromophores scales linearly with the simulated yellowing factor w .

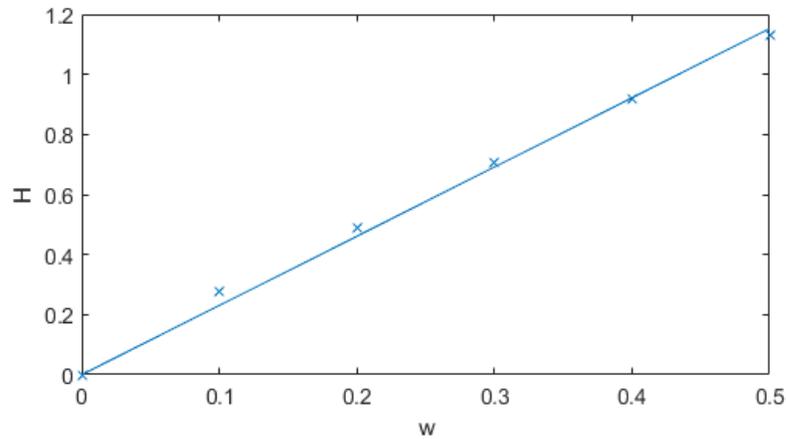


Figure 2.41 Simulated conditions where the appearance of fabric at various stages of yellowing is perfectly correct for by addition of HD factor H . Each point represents a modelled fabric condition which has a colourless appearance.

Effects of lighting

So far in this section we've looked at the ability of FWA's and HD to correct for yellowing under sunlight, and the impact that different light sources could have on the function FWAs and HDs correcting yellowing fabric forms the focus of this section.

Firstly we look at the reason why FWAs are becoming redundant under LED lighting from a spectroscopy perspective. The excitation of the FWA used in this study was roughly measured by recording the fluorescence spectrum at 10 different excitations wavelengths (figure 2.42). This typically resembles a mirror image of the fluorescence spectrum, and in this case peaks around 380 nm. The shape of the fluorescence spectrum does not vary with excitation wavelength, as fluorescence occurs in appreciable yield only from the lowest excited state, according to Kasha's rule.¹

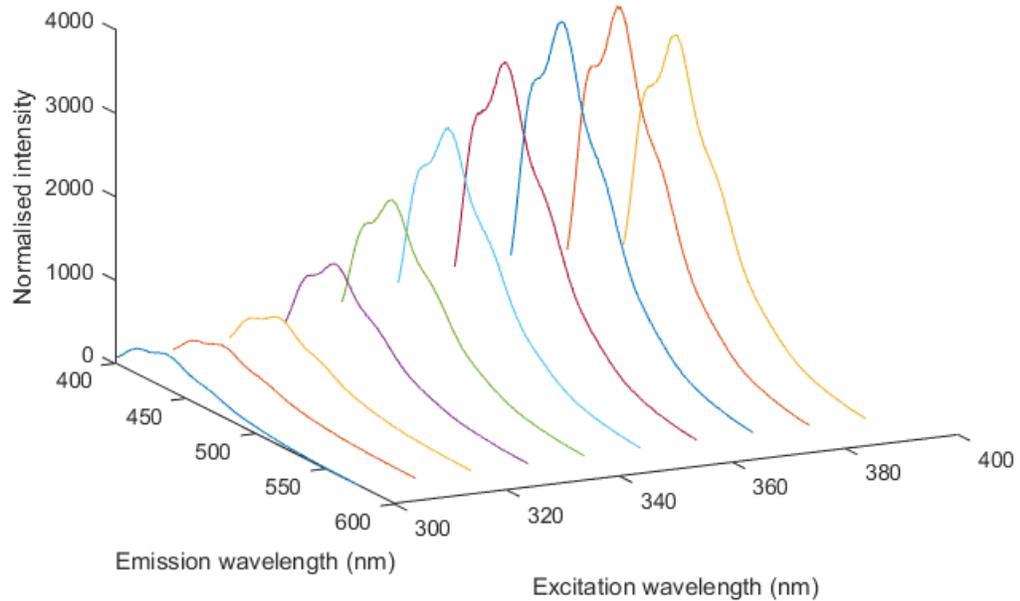


Figure 2.42 Fluorescence of cotton fabric stained with FWA as a function of excitation wavelength, demonstrating the wavelength dependant quantum yield of the FWA.

The emission from a series of light sources was recorded and when compared to the excitation of the FWA used in the model fabric it becomes evident that sunlight, fluorescent lighting and even incandescent lighting (to a small degree) all contain a significant emission in the UV region, which is sufficient to stimulate emission from the FWA used in our model fabric. However, the blue-pumped white LED shows no discernable output that overlaps with the excitation of the FWA, meaning that the mechanism for correcting the appearance for yellowing does not function. This is concordant with the findings of Houser's team in their approach to studying this issue with psychology experiments on a group of forty human participants.⁸⁷ In their study, participants appeared to show a distorted perception of 'whiteness' for white standards containing FWAs under blue-pumped LEDs, and concluded that this must be caused by the inability of this light source to stimulate emission from FWAs.

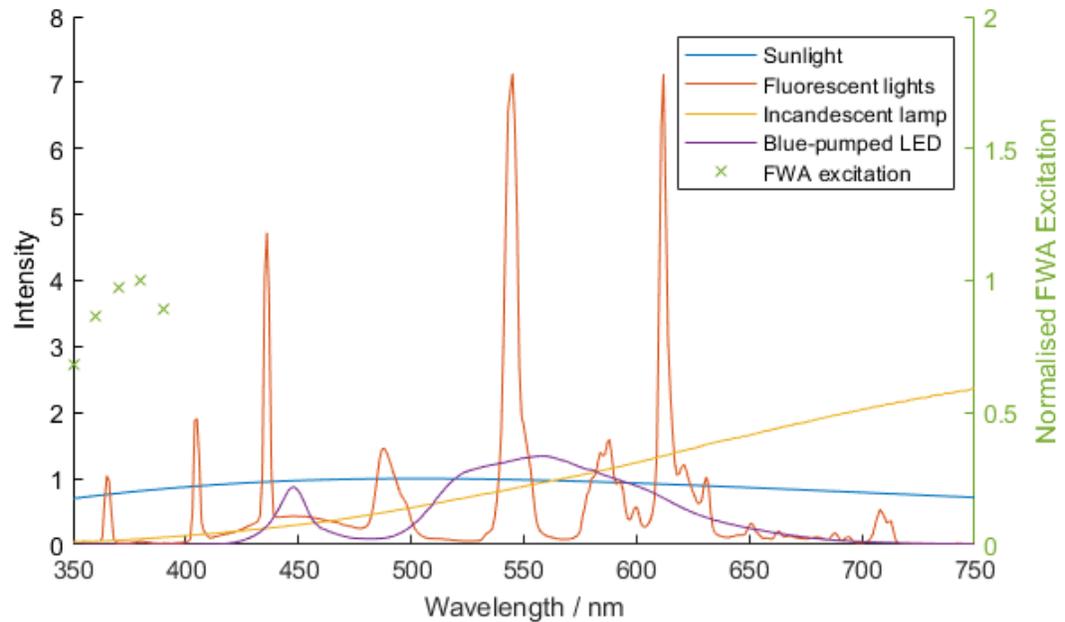


Figure 2.43 Emission spectra from a collection of light sources, intensities have been scaled to represent a uniform luminosity. The excitation spectrum of the FWA used in this study is included for reference.

HDs do not suffer from the same issue as they do not rely on fluorescence to function. The workings of HDs were examined under different light sources, a series of yellowed fabric ($w = 0.3$) was simulated for a range of HD concentrations and the resulting conditions were studied in the $L^*a^*b^*$ colour space (figure 2.44).

The results show that, firstly, yellowed fabric shows susceptibility to colour inconstancy which is displayed as a difference in appearance between sunlight and incandescent lighting, that have a high and low CCT respectively. It can also be seen that with increasing concentration of HD, the susceptibility decreased which results from HDs restoring a flat reflectance profile after yellowing. Comparing the appearance of samples between sunlight and fluorescent lighting, which is visually indistinguishable, shows that yellowed fabrics are not susceptible to illuminant metamerism that is observed in rear-earth elements. It should be noted, however, that even the largest differences in appearance (ΔE) between different light sources is still an order of magnitude smaller lower than that observed in Alexandrite or the rare-earth elements. Un-yellowed fabric shows no major susceptibility to either metamerism or colour inconstancy, which is to

be expected given its featureless reflectance profile.

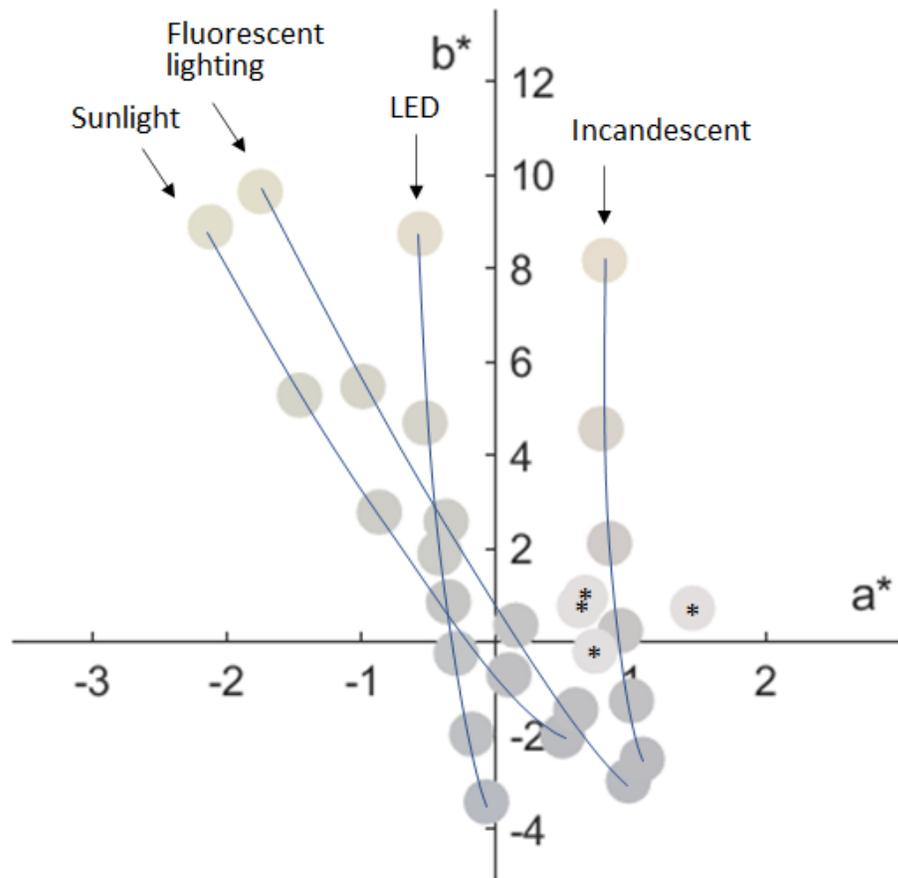


Figure 2.44 $L^*a^*b^*$ coordinates calculated from modelling yellowed fabric ($w = 0.3$) stained with a varying concentration of HD ($H = 0, 0.2, 0.4, 0.6, 0.8, 1$ and trend is displayed with blue line). The colour of the points is a conversion to the sRGB colour space, and points with a central asterisk show appearance of un-yellowed ($w = 0$) fabric. All $L^*a^*b^*$ coordinates are available in section 5.2.

2.3 Summary and conclusions

The colour change phenomenon reported in compounds which contain rare earth elements has been shown not to be the Alexandrite effect. The Alexandrite effect was investigated using the CIE system of colourimetry and found to be an extreme case of colour inconstancy, which is concordant with recent literature in the field.⁵⁴

Neodymium chloride was selected for preliminary studies to probe the colour change observed in rare-earth compounds, and colourimetry studies have shown that colour inconstancy could not explain the observed colour change in the examined samples. Thus, colour inconstancy was eliminated by using light sources of the same colour temperature, cool white fluorescent lighting and a 4310 K black body radiator; Alexandrite showed little change in hue, but the colour change in neodymium chloride was comparable to observations. This colour change was attributed to illuminant metameric failure due to coincidental overlap of the sharp emission lines in fluorescent lighting with the narrow absorption bands caused by the $f-f$ transitions in the lanthanide salt. Further investigation showed that a set of absorption bands centred at 430 nm and 525 nm are almost entirely responsible for the observed difference in colour.

The perceived colour change in six other coloured lanthanide salts was studied for colour inconstancy and illuminant metamerism under fluorescent lighting, together this set constitute one example from each coloured lanthanide in the whole series. The results showed that only the holmium sample displayed a large degree of illuminant metamerism, while no significant sign of colour inconstancy was found in any of the lanthanide salts. The yellow-red colour change of holmium sulphate hydrate was found to be far more pronounced under cool-white fluorescent lighting, and this was again attributed to a similar set of interactions between the narrow absorption bands of the salt and the sharp emission lines of mercury and other elements emitting light in fluorescent tubes.

Therefore, the colour change reported when compounds containing lanthanides are

viewed under sunlight and fluorescent lighting is caused by illuminant metamerism as opposed to the ‘Alexandrite effect’, which has been commonly attributed to this phenomenon.

The same colour change phenomenon has also been reported in actinide compounds by Albrecht-Schmitt and co-workers,^{80,81} where they demonstrated that the appearance of neptunium and americium salts exhibit a yellow-brown colour change when illuminated by broad-spectrum or fluorescent lighting. Although the experimental study of these compounds was not feasible, we may speculate as to the reported difference in colour of these heavier elements. Simply considering the illuminants involved, mainly fluorescent lighting, are more likely to facilitate illuminant metamerism as opposed to the ‘Alexandrite effect’. Moreover, the narrow absorption bands of the reported compounds in the visible region are also highly suggestive of this idea.

Having studied the origin of the colour change in lanthanides under fluorescent lighting, the object of the studies turned to harnessing the effect in a tunable manner. A custom spectrally-tunable light source was constructed with LEDs that correspond to absorption and reflectance bands of the neodymium chloride, which was chosen as the target sample due to its suitable and well-characterised optical properties. Using a spectrometer for *in-situ* spectral feed-back, it was possible to create quasi-metameric near-white illumination conditions which force the appearance of neodymium chloride to any primary or secondary colour. Furthermore, by combining illumination conditions through linear interpolation of individual channels in the tunable light source, it was possible to fine-tune the appearance of the target samples between the primary and secondary colours while maintaining near-white illumination. It was even possible to control the luminosity of the sample while maintaining a uniform luminosity of the light source.

To the best of the authors knowledge, this is the first example of entirely controlling the apparent hue of an object (exclusively materials containing neodymium ions) by

spectral control of the illuminant while maintaining near-white illumination, an effect which was herein termed ‘tunable illuminant metamerism’. This effect provides a novel approach for controlling the spectral profile of light reflected from the surface of an object and may well find application in commercial lighting where colouration cannot be controlled by the altering the object.

The colourimetry models which passed rigorous testing in examining the colour of Alexandrite and the rare-earth elements were then applied to construct a unified semi-empirical model of fabric appearance. The model was used to study the natural degradation of fabric, which is a large contributor to yellowing, and demonstrated the workings of FWAs and HDs in countering the resulting yellow appearance and restoring a ‘flat’ reflectance profile.

The output of the model requires further validation with real samples, however this work forms a platform which is a convenient companion to psychological trials, which are typically used to study such effects at a far greater expense to conduct. The platform can easily be modified to accept any reflectance profile for the fabric, and any SPD for the light source to give a near-spontaneous prediction of colour appearance.

The excitation of FWAs was also investigated to validate the result of a psychology experiment run by Houser and co-workers, which raised the issue of FWAs not functioning under a certain type of LED lighting. The results could be explained by showing that blue-pumped white LEDs do not emit radiation of sufficient energy to excite FWAs into fluorescence.

In combination, these studies conclude that colour must be primarily considered as an illuminant-object pair if accurate results are to be expected. This is a growing consensus in literature and has implications for transitional colour science as well as human colour perception more generally, where colour is often considered as an intrinsic property of an object.²⁶

Studying such extreme cases of colour inconstancy and illuminant metamerism is important to test our pre-existing models of colour perception. In studying compounds with high susceptibility to colour inconstancy and illuminant metamerism, it was found that the CIE system of colorimetry could accurately predict the observed phenomena. However, balancing the spectrally tunable light source and creating visually indistinguishable metamers was not possible even with the latest colour matching functions published by the CIE. This conclusion has also been achieved by other studies¹¹⁰ and suggests that further advances to the field of colour matching are required for delicate colorimetry problems, such as those encountered in tunable illuminant metamerism.

These studies are of great importance to industry. In the lighting industry for example, the quality of colour rendering by white light sources is typically assessed using CRI, with a set of samples showing a wide range of hues at various saturations.¹¹⁶ However, such a set of samples is presented in images alongside the anhydrous neodymium chloride examined in this study (figure 2.3), and the standard set shows very little change in hue between sunlight and fluorescent lighting while the neodymium chloride sample changes hue entirely. These considerations should be applied when designing the next generation of lighting technology, as well as the metrics for its evaluation. Additionally, a more rigorous set of samples may be required to assess colour rendering, and samples based on lanthanide salts are shown to be promising candidates by virtue of their susceptibility to illuminant metamerism.

2.4 Experimental

Reflectance measurements were carried out using a custom fibre-optic reflectance probe: a bifurcated optical fiber (low-OH) couples an ocean optics HL-2000 incandescent light source into two collimators mounted on the probe. The collimators are arranged in a ‘v’ shaped geometry, and focus the channelled probe beam onto a single spot (around 2 mm in diameter) 30 mm from the detector. A third collimator is then used to guide the diffusely-reflected light into a fiber-coupled Ocean Optics Maya 2000 Pro spectrometer (100 μm slit variant). More information about the design is available elsewhere.¹¹⁷ All spectra were referenced to a Spectralon⁹ standard unless otherwise noted. SPD of illuminants was recorded using the same spectrometer and intensity corrected using the procedure described in section 1.2.2. Unless otherwise stated, all images were captured using a Sony a6000 digital camera equipped with a Sigma 16mm f/1.4 DC DN lens. All calculations were carried out in MATLAB, and the raw data is available in section 5.2. All graphical user interfaces were implemented in MATLAB using the ‘app designer’¹¹⁸ add-on package.

2.4.1 Rare-earth elements study

All chemicals were purchased from Sigma-Aldrich and transferred to a sealed glass vessel before use. Fluorescent lights used in colourimetry studies were Phillips brand fluorescent tube lights, with phosphor type confirmed by spectrometry.

2.4.2 Tunable LED Light source

The tunable LED light source used for investigation of tunable illuminant metamerism was constructed from eight individual three watt LED’s purchased from Future Eden Ltd.¹¹⁹ Mounting and heat dissipation was achieved via a custom 1-layer aluminium printed circuit board (PCB), the spacing between LED’s was minimised in order to

achieve the best possible diffusion and even output intensity across the spectral bandwidth of the light source. The LED's were housed in a 3D printed integrating sphere which was painted internally with titanium dioxide based matt-white paint in order to maximise scattering.

LED's were driven by individual current limiting drivers based on the Diodes Incorporated AL8805 IC. The operating current of the LED's was set to 600 mA, which is well below the maximum 900 mA specified by the manufacturer; this was done in order to minimise localised heating of the LED's and maintain a more consistent light intensity output. Intensity was controlled by pulse width modulation at 500 Hz. The signal, which resembles a 500 Hz square-wave of varying duty-cycle, was injected into the driver via a MOSFET forcing the IC's feedback voltage to ground on the high cycle of the PWM signal, resulting in smooth switching of the LED's to guarantee an accurate duty cycle to set current conversion.

PWM signals were generated by an STM32F303K8 microcontroller clocked by 8 MHz and 25 KHz temperature compensated oscillators, which were used as the clock source in order to maintain consistent timings. The drivers and microcontroller were mounted on a custom 4-layer PCB, which was mounted separate from the LED PCB for thermal isolation.

The graphical user interface for tuning spectral intensity is implemented using MATLAB software (control panel shown in figure 2.45) which communicates with the microcontroller via a RS232 virtual serial bus. Feedback of the generated spectrum is achieved by a fiber-coupled spectrometer (same as that used for reflectance measurements), with intensity corrected using the method described in section 1.2.2.

Further information and software source code is available in section 5.2).

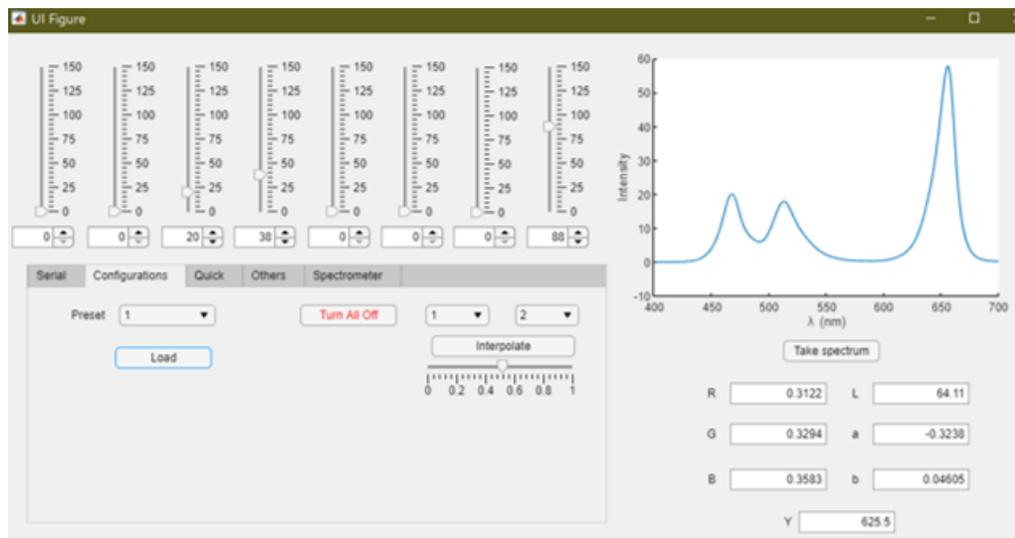


Figure 2.45 User interface for spectrally tunable light source implemented in MATLAB software, allowing for: brightness control of each LED channel, feedback from coupled spectrometer, integrated colourimetry analysis, storage of preset illumination profiles and interpolation between saved profiles.

2.4.3 Modelling fabric appearance

The user interface created to model fabric appearance allows facile colourimetric evaluation of simulated fabric under a range of conditions as well as export of data for further analysis, and is provided in full along with all of the $L^*a^*b^*$ coordinated derived from it in section 5.2. The front-end of the interface is presented in figure 2.46. The variables w , F , and H allow for control of the yellowing, fluorescence, and HD coefficients; which describe the conditions possible to simulate using this model when combined with the choice of illuminant used in CIE calculations.

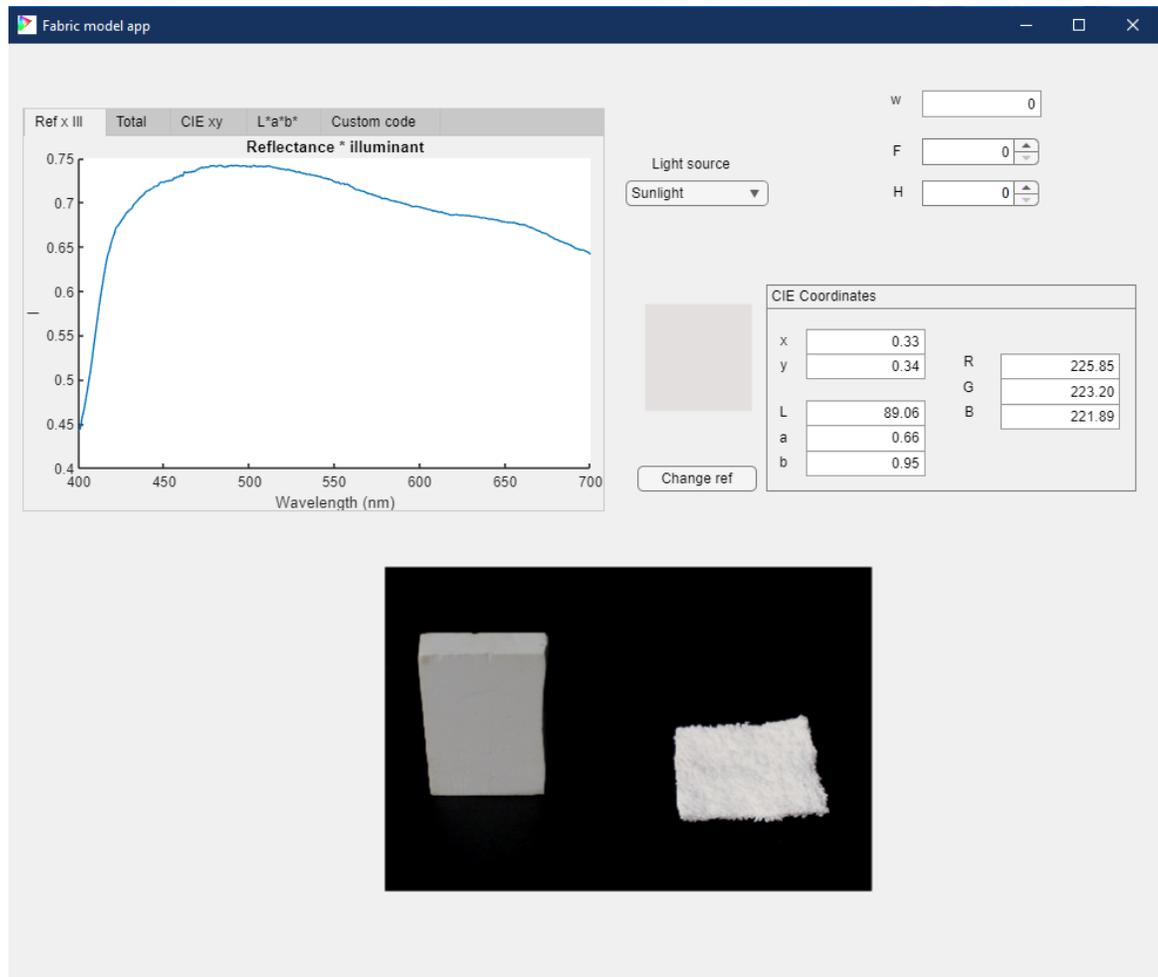


Figure 2.46 User interface of app designed to model fabric appearance including simulating the effects of yellowing, HDs, and FWAs. The app also displays: the SPD of light reflected from the surface, a CIE xy plot, and a CIE $L^*a^*b^*$ slice plot which are not visible in this menu.

Chapter 3

Towards quantification of lipid residues on fabric

The development of a technique that can accurately and consistently quantify the level of lipid present on fabric forms the main topic of study in this chapter. Quantifying the concentration of otherwise invisible lipid trapped in the fabric matrix of a laundry stain is an essential tool to facilitate the development of better laundry formulations.

Initial work was focused on the development of a method to manufacture calibration samples with a set mass of lipid on a fabric sample of known dimensions. These standard samples were then used to define a unit system which could be used to evaluate the level of lipid staining on fabric.

The standards were then employed to develop a novel approach to lipid measurement on fabric, based on DRS in the NIR-SWIR region of the spectrum. The key to this new approach lies in combining DRS with multivariate data analysis and machine learning to intuitively extract quantitative information out of a measured reflectance profile. Further studies were carried out in employing the developed lipid quantification method to imaging, allowing for quantification of absolute lipid mass in a given area of fabric.

3.1 Foreword

The research presented in this chapter have been affected by the outbreak of the COVID-19 pandemic. Due to national restrictions imposed by the British Government, it was also not possible to carry out further work on: validating the developed model on all types of external datasets, and producing reliable calibration stains using the application of inkjet printing.

3.2 Introduction

A laundry stain is typically defined as any undesirable matter which has an adverse effect on the appearance of a fabric.¹²⁰ In the context of cleaning, however, it is important to consider all of the components, including lipids which may not directly affect the appearance. Lipids found in laundry stains may originate from: food stuffs, which can be coated in oils used for cooking and flavouring; sebum, an oily substance excreted by human skin; or from the environment, where lubricants and other oil based products are abundant. Lipid residues are often transparent in the visible portion of the spectrum, and therefore almost impossible to detect by eye. An example is shown in figure 3.1, where a fabric swatch has been heavily stained with vegetable oil yet shows no visible difference in colour across the surface. Certain fabric weaves may allow lipid stains to be visualised due to the change in refractive index of oil-stained fibres leading to a visible difference in colour, but this is only discernable at relatively high concentrations of lipid in a stain. This property is likely the reason that lipid removal on fabric has previously been difficult to quantify, and the amount of lipids present in typical laundry stains is yet to be reported in the literature.



Figure 3.1 Segment of fabric contaminated with lipid stain (outlined with red circle) showing no visible sign of contamination.

Fabrics contaminated with lipid stains unnoticeable to the human eye may not seem

problematic, however, lipids on fabric surfaces are more prone to secondary contamination. Furthermore, lipid deposits comprising unsaturated carbon chains, such as those present from sebum, undergo autoxidation inside the fabrics they contaminate.¹²¹ This leads to polymerisation or stiffening of the fabric, as well as formation of chromophores.³⁸ The oxidation of unsaturated lipid deposits can also liberate short chain fatty acids that cause malodour when metabolised by bacteria present on human skin.¹²² These processes are sometimes grouped with *yellowing*, which as previously described (section 1.4.1) ultimately leads to degradation in appearance and mechanical properties of the textile.

The removal of lipid contaminants is challenging as they have very low solubility in water. Even with the use of surfactants, it has been shown that present day laundry detergents under normal wash conditions cannot remove all the lipid from typical stains. Novel laundry detergents are therefore formulated to contain additives which breakdown the large hydrophobic molecules into more soluble components that can then be removed using surfactants.¹²³ This technology is still in the early stages of development, but the key to further optimisation is a method that allows for the measurement of the level of lipid present on the fabric. Such a technique would allow quantitative assessment of the effectiveness of new laundry detergent formulations.

3.2.1 Chemistry of lipid stains

The type of stains focussed upon in this project are ones which are most commonly encountered on clothing garments, effectively limiting the scope to food stuffs and sebum, which can be transferred onto garments that make physical contact with the skin.

The major lipid components in common food stuffs are chemically quite similar, these are a type of compound called triglycerides.

Glycerides comprise a glycerol backbone esterified with a range of saturated and unsaturated long fatty acids to form a triglyceride (figure 3.2). The nomenclature of these compounds is usually defined by 3 letters which refer to the fatty acids constituting R_1, R_2 and R_3 , for example the common tri-oleic compound is referred to as OOO. The fatty acid esters are abbreviated to: C-caprate, La-Laurate, M-Myristate, P-palmitate, S-stearate, O-oleate and L-linoleate.

The other constituents include: diglycerides in which only two of the glycerol alcohol functionalities are esterified, free fatty acids, and trace compounds present from the plant or animal that the lipid was extracted from.

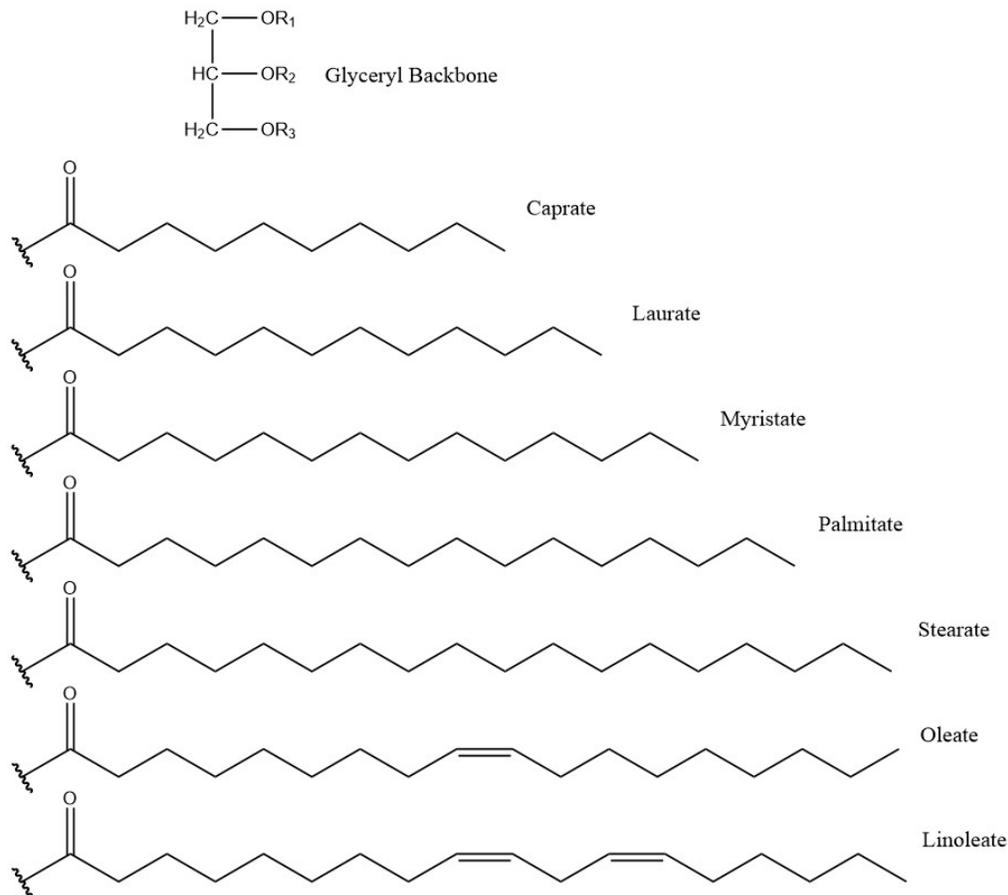


Figure 3.2 Structure of triglyceride backbone along with fatty acid ester groups commonly found in nature which take positions R_1, R_2 and R_3 .

The main difference between the oils and lipids found in food stuffs is therefore the

composition and ratio of the fatty acid esters making up the triglycerides, these can be summarised for the most common oils in food stuffs and sebum:

Olive oil The composition of olive oil tends to differ across different regions in the world, but the most common triglycerides contain the oleic group, with OOO accounting for 40 % - 60 % of the total mixture. POO and OOL account for another 12 % - 20 % and finally the unsaturated POL accounts for around 5 %.¹²⁴

Sunflower oil Sunflower oil is also mainly composed of OOO. Different genetic lines of sunflower have been bred which have higher palmitic, stearic or linoleic content, therefore the exact composition is usually source-dependent. However, OOL, SOO and POO are present in oils from all common mutations.¹²⁵

Coconut oil Coconut oil has a high laurate content, with LaLaLa, LaLaM and CLaLa accounting for more than 60 % of the total. At room temperature it is a soft waxy solid which begins to melt at around 24 °C.¹²⁶

Lard Unlike the previously mentioned oils, lard is an animal product but its chemistry however is not dissimilar. The main constituents are triglycerides rich in palmitic and oleic acids: POO , POS and POP. The ratios of which vary between sources. Lard is a solid at room temperature and begins to melt at 30°C - 40 °C.¹²⁷

Sebum Sebum is an oily substance excreted by human skin to protect it against friction and makes it more resistant to moisture.¹²⁸ It is a complex mixture which varies

in composition from person to person, part of the body and between reported studies. The main components can be summarised as: Squalene, triglycerides, cholesterol, cholesterol esters, wax esters and free fatty acids.^{129–131}

3.2.2 Quantitative assessment of laundry detergent performance

Assessing the level cleaning achieved by laundry formulations is typically carried out using: visual assessment, gravimetry, colourimetry, or exotic methods such as radioactive labelling.^{38,120,123} However, all of the afore mentioned results are relative measurements, meaning that they can be used to compare levels of staining, but not give an absolute value as to the amount present.

For colour stains, colourimetry is a popular approach and the technique can be traced back to a study published in 1980.¹²⁰ Formulations are assessed using the following method: a imitation stain (model stain) is applied onto standard fabric swatches in a controlled quantity. Samples are then placed into a box with homogeneous illumination and a standard colour camera is then used to image the samples. The fabric swatches are then washed, dried and imaged again using the same equipment. The RGB values obtained from the images are converted into the CIE L*a*b* colour space, and the Euclidean distance between the L*a*b* coordinates of the clean fabric and the center of the stain are measured using a parameter called ΔE^* , which is defined:

$$\Delta E^* = \sqrt{(L_2^* - L_1^*)^2 + (a_2^* - a_1^*)^2 + (b_2^* - b_1^*)^2} \quad (3.1)$$

where $(L_1^*, a_1^*$ and $b_1^*)$ represents the L*a*b* coordinates of the clean fabric and $(L_2^*, a_2^*$ and $b_2^*)$ the L*a*b* coordinates of the stain centre. Performance of laundry detergents is then assessed using a metric called stain removal index (SRI) which is thought to indicate the percentage of stain removed after a wash cycle. It is calculated by taking the difference in ΔE^* between washed and unwashed fabrics as a ratio of the ΔE^* in

the unwashed stain (equation 3.2).

$$SRI = \frac{\Delta E_{unwashed}^* - \Delta E_{washed}^*}{\Delta E_{unwashed}^*} \times 100\% \quad (3.2)$$

Measuring the performance of formulations on colour-less lipid stains is, however, more challenging. Some studies have used solvent extraction, followed by weight measurements as a substitute for ΔE ,¹²³ although there are many sources of error in using this approach as any soluble components are assumed to be lipid. Another approach, which was reported in the original study,¹²⁰ has been to add a hydrophobic dye to the model stain in low concentrations (demonstrated in figure 3.3), allowing the colorimetric method to approximately compare the amount of lipid removed between samples in model stain wash tests.

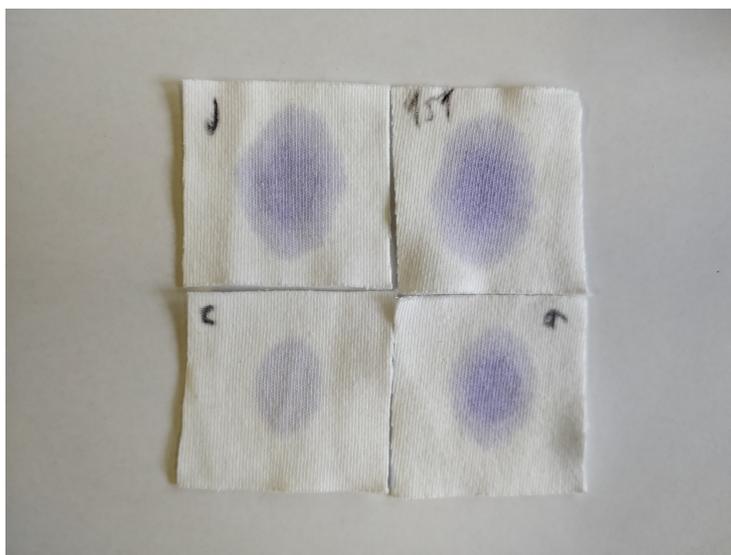


Figure 3.3 Example of dyed lipid stains on standard fabric swatches used for colourimetric analysis.

This empirical metric of observing the wash-off of the lipid-bound dye has been developed as it allows the otherwise colourless lipid stains to be visualised with basic instrumentation. It stems from measurement of laundry detergent performance where

appearance is the key factor, however, the authors of the original study noted ‘a number of problems in using these methods of calculation’, in particular ‘percentage stain removal may not describe the relative performance of products’. The errors can be traced back to several assumptions which are key to the accuracy of the approach, the validity and implications of which are discussed and tested experimentally further in this chapter.

In a similar field, the quantification of lipid present on hard surfaces, such as those found in dish washers, has been previously explored. A widely used method in the field is called ‘bath substrate flow’ and was originally published in 2003.¹³² The technique is based on having a large confined surface, such as glass spheres packed into a column, called the substrate. This container is subject to model stain and washing solution is then continuously exchanged between a reservoir and the substrate to simulate washing. Following the ‘wash’ cycle, the mass of the substrate and washing solution can be measured to determine the amount transferred. Alternatively, if the model stain is appropriately chosen then the pH of the wash solution can be continuously monitored. This method is also prone to several issues, mainly that the mass of the substrate or bath is usually far greater than the mass of the model stain making accurate measurements difficult. Measurement of the pH circumvents this issue and even allows real time analysis during a wash. However this is limited to stains which cause a change in pH of the usually alkaline wash solution, which excludes the majority of greasy food stain components.

In an attempt to resolve these issues, an improvement to the method has been published. The improved technique involves addition of a hydrophobic dye to the model stain, followed by extraction of the dye into an organic solvent post-wash and examination of the extract by absorption spectroscopy or colourimetry.^{133–135} Such a method however, inherits the problem associated with dye-staining which have been discussed earlier in this section.

3.3 Calibration sample set

In order to develop a technique for quantifying lipid stains on fabric, a set of samples which have a known concentration of model stain is required. Such protocols for making standards does not currently exist and therefore needs to be developed prior to the development of new analytical methods. This section describes the approach taken to develop a method for sample preparation which yields uniformly stained fabric samples of known lipid concentration.

3.3.1 Unit system

There exists no standard definition for the units of lipid concentration on fabric. In this application, the lipid being measured can have varying compositions which will be measured on several different fabric substrates. For this work, a system of measurement suitable for this purpose has been devised and will be used throughout the course of this study. The quantity found most suitable was determined to be the mass of stain per mass of fabric as a universal scale in describing the concentration of lipids on fabric. This quantity will be denoted c_l , and as a ratio of two masses it can be used as a unit-less value to describe the amount of lipid present on a fabric surface as a fraction of the fabric substrates mass (eq 3.3). In order to simplify the interpretation of results, the quantity will be presented as percentage and referred to throughout this work as ‘wt%’:

$$c_l = \frac{m_l}{m_f} \quad \text{wt}\% = 100 c_l \quad (3.3)$$

where m_l is the mass of lipid in a given volume of fabric and m_f is the mass of the fabric substrate occupied by the lipid. Interpretation of the wt% scale should also be done with care. While values of the wt% scale are expected to lie within the range of 0-100 wt% based upon measurement of real garments and laboratory made replicates,

the scale can also extend above these values; a sample may comprise more mass of stain than the fabric substrate to yield a measured c_l above 100%.

3.3.2 Model compounds

To study the level of lipid soiling on fabrics, at least in initial studies, a reproducible and readily available compound is needed to represent realistic lipid stains. This will be referred to throughout as a *model stain* and needs to be representative of real stains encountered on fabric. Due to the low selectivity of the techniques which will be developed in this work (Colourimetry and SWIR-DRS) as well as the chemical similarity of lipid stains, the choice of lipid used as a model stain will likely have little effect on the resulting measurements. Sunflower oil was therefore chosen as the model stain due to its availability and ease of handling.

The substrate used as the *model fabric* to represent typical items of laundry will be cotton in initial studies, and further work will expand the technique to include polyester and poly-cotton (a blend of polyester and cotton typically in around a 3:7 ratio). Together these fabrics account for a large proportion of the laundry industry. They were selected as together they may show how a reflectance based technique may be suitable for quantification of lipids on textiles of different fabrics as well as fabrics consisting of blended materials.

3.3.3 Sample preparation

In order to have samples with a known concentration of model stain, a known mass of fabric needs to have a known mass of model stain distributed evenly across its volume, meaning measurements can be carried out at any location on the sample. Directly staining fabrics with varying amounts of model stain results in a central region of high concentration with the edges remaining completely unstained, although capillary action

in between the fibers does help in homogenisation of the oil across the fabric.

A better approach to creating uniform stains was established by diluting the model stain with a volatile-organic solvent and applying a larger volume to the fabric, with the aim to completely saturate the swatch and allow the model stain to distribute evenly before evaporation of the solvent. This methodology was found effective in producing homogeneous samples with stain concentrations above 10 wt%. Below this limit, samples exhibit a ‘coffee-ring’ like effect, with more stain deposited on the edges as opposed to the center of the fabric. The accuracy of this method, as well as the issue in making low-concentration samples will be addressed further in this section.

Cotton model sample set

A sample set designed to create a prediction model exclusively for cotton fabrics was made using 2x2 cm square (cut using a surgical scalpel) 100 % knitted cotton fabric swatches of area density 33 mg cm⁻². The concentration range of the sample set was limited to a maximum at where the fabric can no longer absorb any more oil in between its fibers, which was found to be 97 wt%. 10 Stain solutions of 1 ml volume were made up in individual vials using an auto pipette to add the right proportions of sunflower oil and ethyl acetate. In order to allow colorimetric analysis of the stains, a trace amount of lipophilic dye (0.1 % solvent violet 13) was added to the sunflower oil prior to making the solutions. In such low concentrations, it was found that the dye has no detectable contribution to the NIR-SWIR reflectance spectrum. Staining was replicated by applying 200 μ l of model stain solution, diluted to varying degrees with ethyl acetate, directly to the center of fabric swatches and leaving the samples at room temperature for two hours to allow the solvent to evaporate. In total, 10 pairs of stained samples with lipid concentrations increasing in steps of 9.7 wt% up to the maximum were made, with a further 14 unstained sample yielding the complete set of 34 samples.

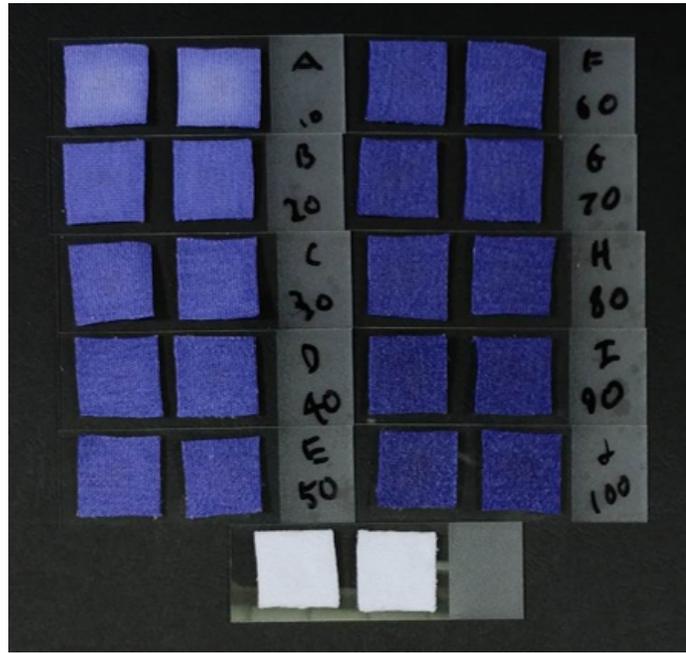


Figure 3.4 Calibration set of samples with set lipid concentration. Only two clean samples are shown for clarity. A(10): 9.7 wt% - J(100) 97 wt%.

Multi-fabric model sample set

To expand the scope of fabrics used in modelling, the set of calibration samples was expanded to include polyester, poly cotton, as well as a composite of the two materials used in collars. Samples were prepared in the same manner as the cotton set with a few modifications to address issues which were identified in the process: the relatively high viscosity of the sunflower oil leads to difficulties in accurately dispensing the liquid. This was mitigated by first diluting the oil to 70% (by volume) before preparing stain solutions. The uncertainty in model stain concentration of the stain solution was also minimised by preparing the solutions in batches of 100 ml. Fabric swatches were cut to 2x2 cm using a CNC laser cutter, which minimised the error in the dimensions of the swatches, however, the edges were slightly degraded from the cutting process. The concentration range of the samples was limited to a maximum of 60 wt% as it was found that some fabric weaves could not bear a greater proportion of model stain. The amount of model stain solution deposited onto the swatches was scaled to account for differences in fabric density. In total, the sample set used in modelling consists of

84 stained samples: two different types of cotton, poly-cotton and polyester; as well a batch of collar material, with 12 samples for each fabric type, at 6 concentrations increasing in steps of 9.7 wt% and two samples at each concentration. A further 120 unstained samples were included to the set. The properties of the fabrics have been summarised in table 3.1.

Table 3.1 Area-density of fabrics used in multi-fabric calibration set

Fabric type	Density (mg cm ⁻²)
Knitted cotton	33.0
Flat cotton	13.1
Poly cotton(1)	11.8
Poly cotton(2)	11.1
Polyester(1)	14.3
Polyester(2)	18.7
Collar material	49.4

Error analysis

It is important to assess the accuracy associated to the process of making the calibration set of samples, which will aid in quantifying the error in the subsequent reflectance dataset used in development of machine learning algorithms.

The experimental factors which are likely to contribute most to the error in the set lipid concentration have been identified as: the variability in dispensing model stain solution onto fabric swatches, and the uncertainty in dimensions of the fabric swatches. A series of experiments were designed and carried out to measure the certainty with which the set lipid on fabric concentration in the training data-set can be known. This practical approach allows for reliable estimation of the errors associated with the tools which are used in the preparation of stained fabric calibration standards.

The accuracy and precision of model stain deposition was measured by repeatedly transferring 200 μL aliquots of model stain solution into a vessel. The increase in mass of the vessel was subsequently monitored for each repeat using an analytical balance. The consistency of fabric swatches was measured using a vernier caliper to gauge the horizontal dimensions and consequently the area. An analytical balance was employed in combination to allow for an area density measurement. The dataset acquired from these experiments is available in section 5.2 and the results have been summarised in table 3.2.

To aid the propagation of errors, the set lipid concentration of samples in the training set can be expressed mathematically using the definition of lipid on fabric concentration in equation 3.3. Furthermore the relationship can be rewritten in terms which have been experimentally probed:

$$c_l = \frac{m_l}{m_f} \propto \frac{c_{ms} m_d}{\rho_d \rho_f A_s T_f} \quad (3.4)$$

where c_{ms} is the concentration of model stain solution, m_d is the mass of model stain dispensed onto the sample, ρ_d is the density of model stain solution, ρ_f is the area density of model fabric, A_s is the area of the model fabric swatch, and T_f is the thickness of the model fabric swatch. This equation shows that the probed parameters are linearly related to the lipid concentration present in the sample. The errors in the measured parameter can therefore be consolidated into an error in the set lipid concentration using a root sum of squares.

The thickness of model fabric could not be accurately measured due to the softness of the material causing ambiguity in the reading. The error contributed by the fluctuations in the thickness of model fabric is, however, reflected in the error measured for ρ_f , which is accounted for in error propagation. The error in making solutions of a set concentration, c_{ms} , was initially a dominant factor due to the high viscosity of the model stain. The

procedure was however improved by working on a large scale (100 mL) compared to preparation of samples (0.2 mL). The error associated to c_{ms} is estimated to be less than 0.2% and therefore insignificant in these calculations.

The errors from experimentally tested parameters were calculated as the standard error of the recorded data, otherwise called the standard deviation of the mean, and takes into account the offset from the set value for each quantity. The calculated error in each quantity along with propagation to error in lipid concentration are tabulated in their relative form in table 3.2.

Table 3.2 Parameters used to probe the magnitude of the error in the set lipid concentration of the calibration sample set. ^A The value in brackets denotes the error in the last significant digit.

Quantity	Mean	Error % ^A
m_d	183.8 mg	0.4(1)
A_s	404.1 mm ²	2.4(6)
ρ_f	0.33 mg mm ⁻²	0.9(3)
$C_{l-training}$		3(1)

The error in the error was calculated based on the number of repeat measurements performed in the experiments for each quantity.¹³⁶ This quantity needs to be interpreted with care as it does not account for any experimental factors which were not controlled. Alternatively, it provides an estimate of how much the calculated errors are likely to vary if the number of repeat readings was increased. The results show that the largest contribution to the error in set lipid concentration is the size variation of the fabric swatches. The 2-4% error calculated in $C_{l-training}$ sets the upper limit on the accuracy of prediction models training using this dataset. This value will be used in detecting over-fitting present in prediction models.

3.3.4 Spectroscopy of lipid and fabric substrates

The optical spectrum of lipids in the SWIR region of the spectrum is dominated by five features: two major features at 1700 nm (5880 cm^{-1}) and 2300 nm (4350 cm^{-1}), and three minor features at 1200 nm (8330 cm^{-1}), 1400 nm (7140 cm^{-1}) and 2150 nm (4650 cm^{-1}). The major features are attributed to the first overtone of C-H stretching modes (1700 nm) and combinations of C-H stretches (2300 nm), while the minor features are difficult to assign.¹³⁷ These can be observed in the transmittance spectrum recorded from the model stain in figure 3.5. Due to the similar chemical nature of lipids, specifically regarding the long alkane chains constituting many C-H bonds, these features are likely to be present in all lipid stains encountered within the scope of this project. In some oils such as olive, the trace presence of pigments, mainly carotenoids, anthocyanins and chlorophyll is responsible for a yellow and green colouring, resulting from two absorbance bands around 420–460 nm and 668 nm.¹³⁷

The composition of cotton fabric largely consists of two components, cellulose and water, and these account for 91% and 8% of the total mass respectively, with the remaining components consisting of protoplasm, pectins, fatty substances, and salts.¹³⁸ As such, the optical spectrum of cotton in the SWIR region is dominated by the absorbance bands of cellulose. This is seen as broad absorption bands at 1500 nm (6670 cm^{-1}), 1900 nm (5260 cm^{-1}), 2100 nm (4760 cm^{-1}) and 2500 nm (4000 cm^{-1}). The band at 1900 nm is due to the presence of water in the fabric, and the band at 2100 nm is a combination band of O-H / C-H bending modes as well as an O-H stretch.¹³⁹ These features can be observed in the measured DRS of the model cotton fabric in figure 3.5.

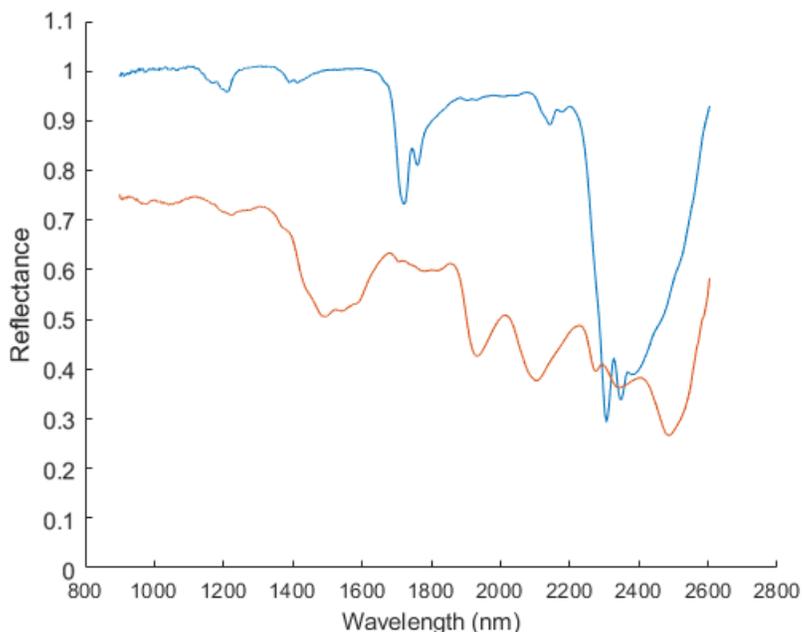


Figure 3.5 DRS spectrum of kitted cotton fabric (Orange) and transfectance spectrum of sunflower oil (Blue) recorded as a thin layer deposited on a teflon backing and reference corrected.

The DRS spectra of polyester and ploy cotton model fabrics from the calibration set are also presented here for reference (figure 3.6). Polyester refers to a class of polymers that contain the ester group in the backbone, but most commonly is used to describe the material polyethylene terephthalate (PET). The optical spectrum of PET in the SWIR region shows many distinct features. A distinct peak at 1660 nm originates from the first overtone of the aromatic C-H stretching mode. The band around 1900 nm can be attributed to the first overtone of the O-H stretch, and overlaps with the second overtone of the C=O stretch. The feature around 2170 nm is assigned to combination bands of aromatic C-H and combination band of the methylene group. Finally the features beyond 2200 nm are difficult to distinguish, although they are know to involve the methylene groups, aromatic C-H and COO moieties.¹⁴⁰ The spectrum of polycotton shows absorbance bands from both of it's constituent components.

As fabrics are stained with an increasing concentration of oil, the absorption bands from the oil around 1200 nm, 1700 nm and 2300 nm become increasingly prevalent. The

general trend seen with increasing stain concentration also shows an overall decrease in reflectance, which is non-linear with wavelength and attributed to changes in scattering by the oil saturated fibres.

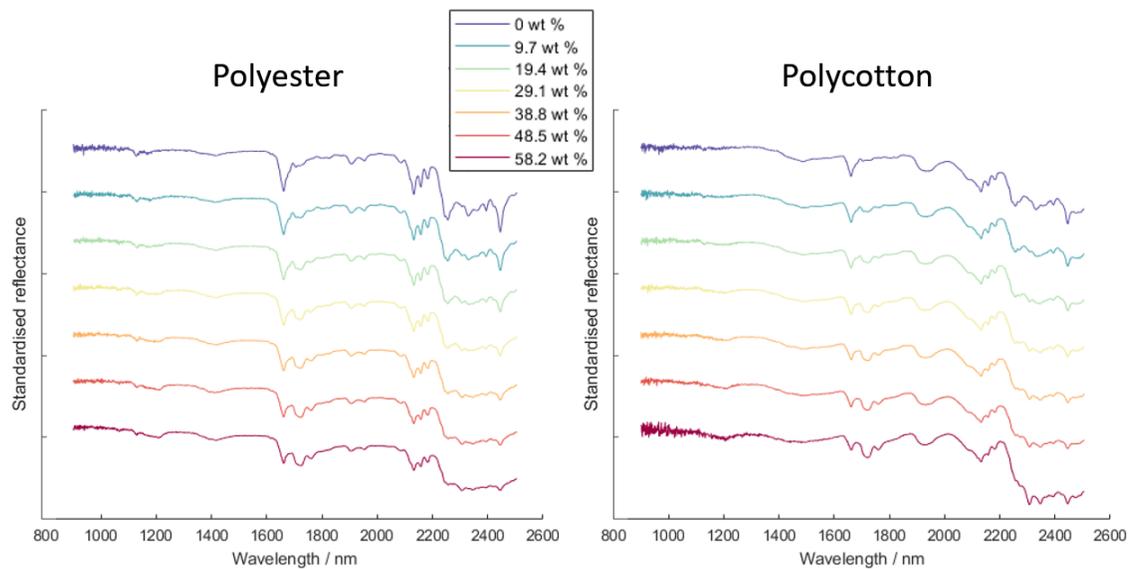


Figure 3.6 DRS spectra sets of fabrics stained with an increasing concentration of model stain (displayed in the legend). (Left) Polyester training set. (Right) Polycotton training set. Spectra have been standardised by subtracting the mean and dividing by the standard deviation, and are presented vertically offset for clarity.

3.4 Modelling

3.4.1 Cotton model

Cotton fabric is widely used in textiles and serves as a good starting point for developing a chemometrics model. A system including a single type of fabric and a single type of model stain lends itself well to multivariate analysis due to the reduced number of variables. The interpretation of the inner workings of such models is also usually possible as opposed to more complex models. The use for such a model is limited to cotton garments, which restricts the application in real-world stains that may be based on a variety of other fabric, or even blended fabric substrates. While gathering all real-world fabrics into a single model is beyond the scope of this project, the techniques which would be used to do so if such a data-set become available would have a firm grounding in this work. Furthermore, even a limited cotton model has application to assessment of laundry detergent performance in stain-removal experiments.

Preprocessing

Prior to construction of a mathematical model based on reflectance data, it is necessary to apply some initial processing of the raw reflectance spectrum in order to remove device dependant artefacts, minimise experimental artefacts, maximise analyte features, or linearise the response with respect to analyte concentration.

Selecting a pre-processing method is usually limited to trial-and-error, and is therefore considered somewhat subjective.¹⁴¹ In this study, the performance of pre-processing methods is assessed by using unsupervised machine learning algorithms to reduce the dimensionality of the dataset.

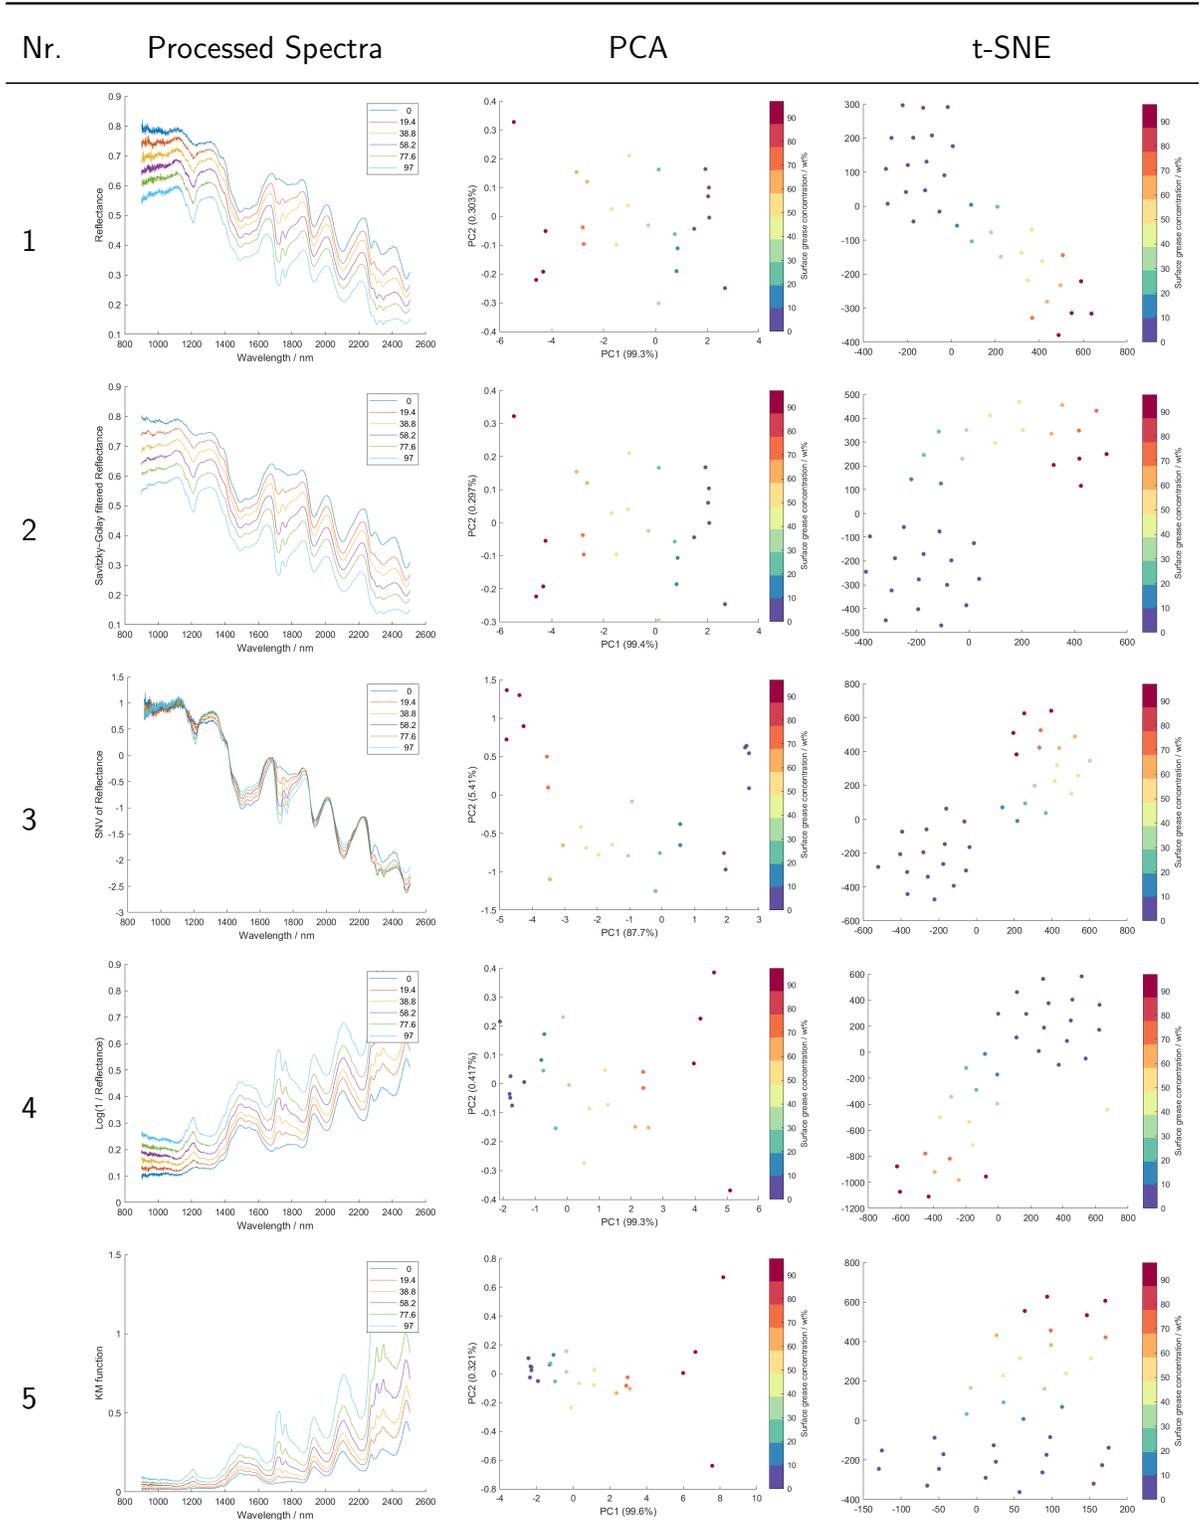
As the main source of inter-sample variance should be caused by changing the concentration of analyte, reducing the dimensionality of a pre-processed dataset should define a pattern in the higher rank components. The clarity of the pattern and similarity of

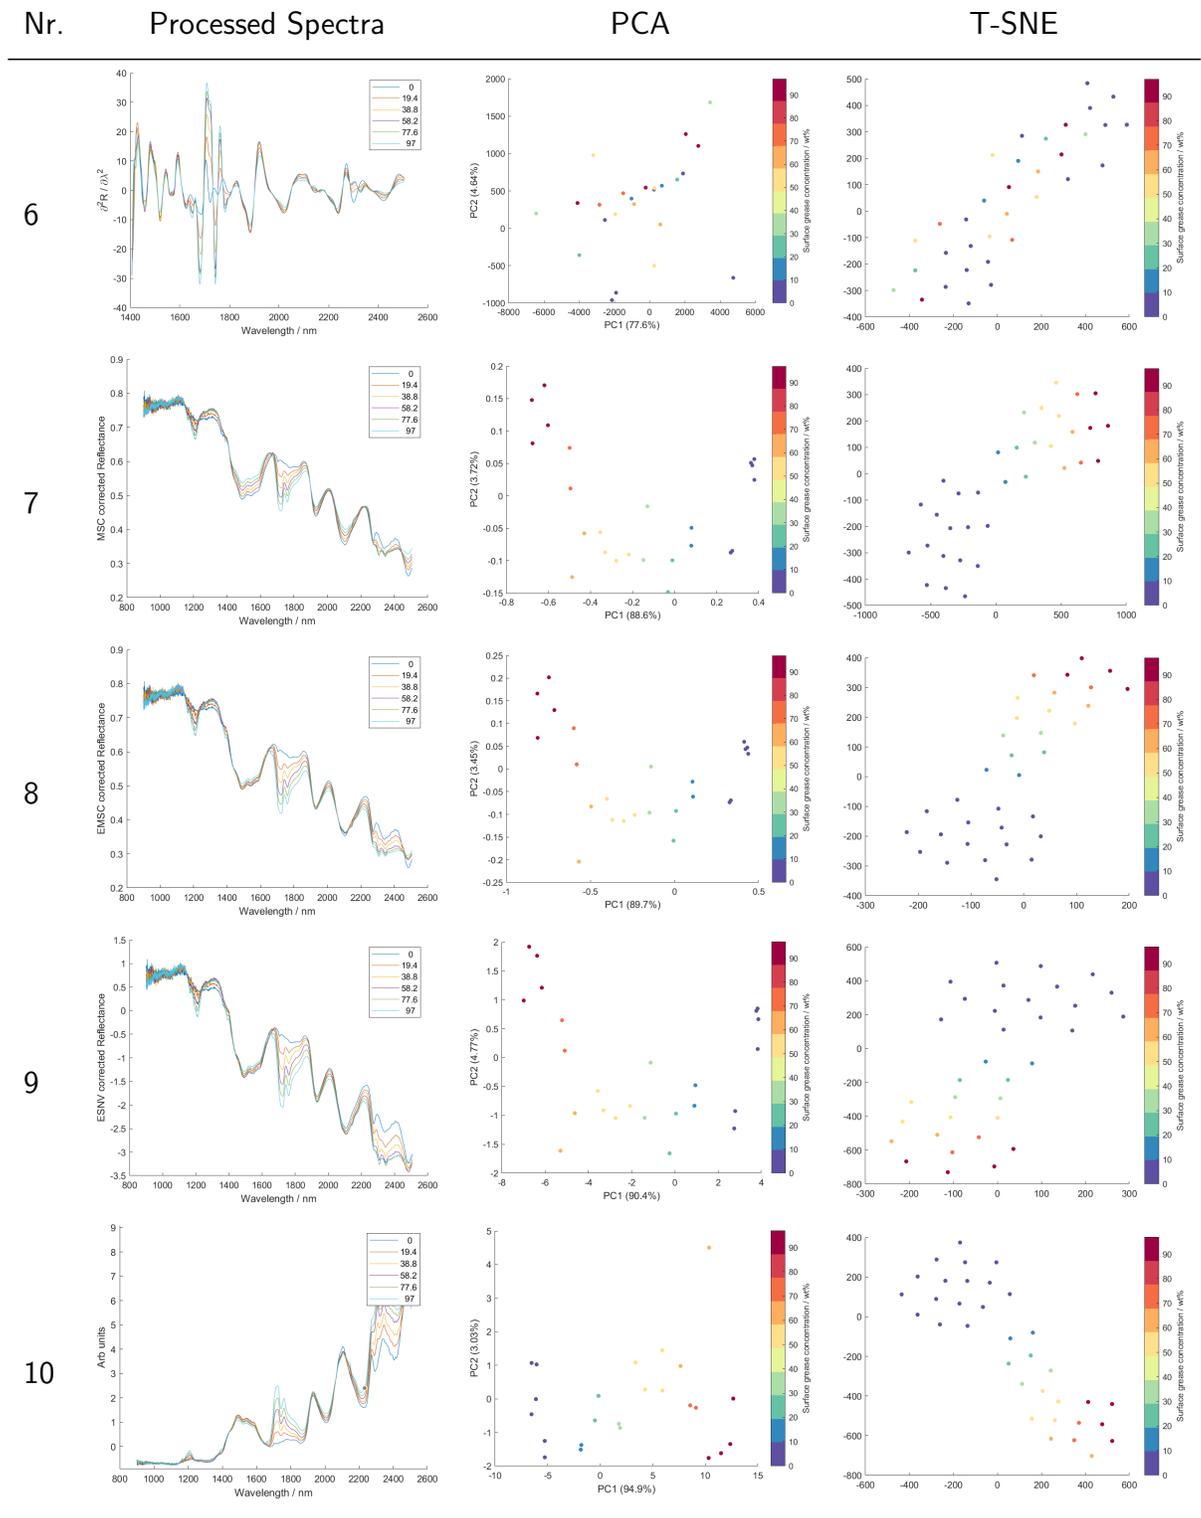
readings repeated on identical samples can then be used to justify the most suitable pre-processing method. Furthermore, if a mathematical relationship between the new variables and analyte concentration present in the samples can be found, it can be used as a basis for a chemometric model.

The unsupervised ML techniques used in this case are PCA and t-SNE; PCA is a tangible technique which can be later used to build and model. Unlike PCA, which can only show linear projections, t-SNE is a highly non-linear method which provides another point of comparison if PCA shows similar results or the relationship between analyte concentration and the obtained reflectance data cannot be described by a linear projection.

The relatively small size of the dataset used in this initial model makes computation on an average computer easily assessable, therefore a variety of preprocessing methods are explored including combinations which were found to give promising results. The result of each pre-processing method, along with the outputs of the unsupervised machine learning algorithms is shown in table 3.3 and the corresponding transformations are explained thereafter.

Table 3.3 preprocessed DRS spectra recorded from the stained cotton calibration set. The result of pre-processing is shown on a selection of spectra covering the analyte concentration range, with concentration shown in each figure legend with units of wt%. The scores plot of the first two principal components is shown with the amount of variance they explain in axis labels. The pre-processing method are denoted (Nr): (1) Raw spectra. (2) Savitzky-Golay filtering. (3) SNV. (4) Log(1 / R). (5) Kubelka-Munk function. (6) second derivative. (7) MSC. (8) EMSC. (9) E-SNV. (10) ESNV of Kubelka-Munk transformation of EMSC corrected R.





Raw spectra First the condition of unprocessed spectra is investigated. Using Unprocessed data is undesirable as it contains device dependant artefacts which may be incorporated into the later developed machine learning models, making them unreliable on spectrometers other than the one used to collect the calibration dataset. For example in the SWIR spectrometer used to collect the dataset, the sample was manually brought into focus of the probe beam. The alignment is subject to the operators judgment and will vary between operators. Differences in probe alignment will result in a non-linear distortion of the recorded reflectance spectrum, which is difficult to reject in ML models.

Looking at the ML output in table 3.3 (1), the first PC can account for 99.3% of the variance in the dataset. This is a strong indication that linear projections will be suitable in this case. The t-SNE projection shows uneven clustering through out the concentration range of the analyte, although this may be due to the uneven number of samples taken at different surface lipid concentrations.

Noise reduction Any measurement system is inherently subject to random noise. In SWIR spectra this is usually present as Johnson–Nyquist noise from the detector which has been amplified. The noise is unevenly present throughout the spectrum due to the uneven output of the probe light source and subsequently the SNR across instruments bandwidth. Noise can be reduced at the expense of spectral resolution. A common approach is to approximate the spectrum by fitting weighted piecewise polynomials to adjacent points, otherwise called a Savitzky–Golay filter. The outcome is reduced noise with minimised distortion to the signal shape. In this context the result of Savitzky–Golay filtering can be seen in table 3.3 (2). Spectra appear less noisy, and PCA shows more covariance described by a single PC. These techniques however, generally lower the amount of information present in the data set, therefore models should instead be trained with tolerance for common instrumental noise.

Magnitude and offset correction. Errors which cause spectra to be scaled or offset are difficult to remove without external information, for example the output power of the probe beam optical alignment on the sample. However, spectra can be standardised via a linear transformation to a form where such errors are excluded, which allows better comparison between spectra. Although the original intensity information is lost and only the shape of the spectrum considered. It is applied by subtracting the mean of the entire spectrum (\bar{x}) from each point (x) and dividing by the standard deviation (S) of the spectrum (equation 3.5).

$$z(\lambda) = \frac{x(\lambda) - \bar{x}}{S} \quad (3.5)$$

In chemometrics this is called the Standard Normal Variate (SNV) of a spectrum, and the result is a standard

The outcome of SNV on the calibration set can be seen in table 3.3 (3). Spectra become overlapping in regions which are little affected by the analyte and the C-H bands introduced to the spectrum become prominent. PCA shows points of the same analyte concentration appearing closer together, although the amount of variance explained by a single principal component is reduced.

Response linearity In construction of a model using regression, it is beneficial that the predictor variables (reflectance spectra) are linearly related to the response variables (concentrations). In reflectance spectroscopy this is inherently not the case, and there exists no exact transformation to relate the reflectance intensity to the concentration of absorbing species. To address this issue, several transformations have been proposed which input a reflectance spectrum and output a function which is approximately proportional to the concentration of absorbing species. A popular transformation is the Kubelka-Munk function (F(R)):

$$F(R(\lambda)) = \frac{(1 - R(\lambda))^2}{2R(\lambda)} \quad (3.6)$$

Kubelka-Munk theory was developed using several approximations and is based on the probability of a ray escaping the surface depending on how deeply it penetrates the sample.¹⁴² It has since been shown that a more accurate approximation should use $4R(\lambda)$ in the denominator,¹⁴³ and this will be referred to as the Kubelka-Munk equation.

Another commonly used approach in reflectance based chemometrics is to use the logarithm of the reciprocal of the reflectance:

$$f(R(\lambda)) = \log_{10}(1/R(\lambda)) \quad (3.7)$$

There is no theory which suggests that this transformation should be linearly dependent on the concentration of absorbing species, however, in practice it has been found to give good correlation in many systems.¹⁴⁴

The use of Kubelka-Munk theory and $\log(1/R)$ on the calibration dataset can be seen in table 3.3 (4) and (5) respectively. The transforms both show better clustering of similar samples in PCA compared to unprocessed data. Kubelka-Munk theory performs marginally better in that respect as well as showing more variance explained by a single principal component, and more even clustering throughout the analyte concentration range in t-SNE.

Derivatives When building machine learning models to work on features which are difficult to distinguish, it may be useful to take derivatives of reflectance with respect to wavelength. This causes peaks to become more prominent compared to the relatively small gradients of the background. In chemometrics, the second derivative is

found to give good results. it is also important to remove instrumental noise from the spectrum using the aforementioned Savitky-Golay filtering, otherwise the derivative becomes dominated by random fluctuations.

In this case, the result of second derivative pre-processing seemed to show poor correlation with the analyte concentration in PCA and t-SNE as can be seen in table 3.3 (6).

Scatter correction In reflectance spectroscopy the relationship between recorded spectra and concentration of absorbing species is complicated as the introduction of such species changes the refractive index, and introduces non-linear wavelength-dependant distortions. This is extremely difficult to correct for and a popular empirical approach involves correcting the spectrum ($R_s(\lambda)$) by fitting it to a reference one ($R_r(\lambda)$) via a polynomial function in a technique called Multiplicative Scatter Correction (MSC).¹⁴⁵ In this case the reference spectrum is that of clean cotton fabric and a second order function is found to give sufficient results, therefore the scatter correction function take the form:

$$R_r(\lambda) = \alpha R_s^2(\lambda) + \beta R_s(\lambda) + \gamma + \epsilon(\lambda) \quad (3.8)$$

where α , β and γ are constants found by least squares fitting in order to minimise ϵ , the vector of residuals. The MSC function is then applied by using the determined coefficients on the reflectance spectrum of the sample in question ($R_s(\lambda)$):

$$R_c(\lambda) = \alpha R_s^2(\lambda) + \beta R_s(\lambda) + \gamma \quad (3.9)$$

where $R_c(\lambda)$ is corrected spectrum to be used in modelling. This correction works to remove non-linear distortions which hinder models based on linear projections. The application of this method to the training set shows results which are comparable to

those from SNV and can be seen in table 3.3 (7).

Despite the promising results, MSC is fundamentally flawed as the absorbance bands from the analyte cannot be fitted to the reference spectrum, and therefore reduce the overall quality of the fit. For this reason MSC was later improved to only fit the regions where the analyte absorbance has little effect on the spectrum, and the scatter correction function can include intensity dependant terms, as well as wavelength dependant terms.¹⁴⁶ In this case, the wavelength dependant terms were found unnecessary and the scatter correction function was built by excluding the absorption bands of the model stain (described in section 3.3.4) when finding the α , β and γ coefficients. The results of EMSC are shown in table 3.3 (8) where the fit in areas least affected by the analyte become significantly improved, along with the variance described by a single PC.

E-SNV In this study, a novel pre-processing method was devised and found to give excellent results. Based on the method used to extend MSC, SNV can be improved by excluding spectral regions which are heavily affected by the analyte. This is carried out by applying a window function on the spectrum in question, to exclude absorbance bands of the analyte. This is followed by calculation of the mean and standard deviation of the resulting truncated spectrum. SNV can then be carried out in the same manner using equation 3.5 with the previously calculated mean and standard deviation values. This approach remains unsupervised and as shown in figure 3.3 (9) has an excellent capability to remove the differences between sample in spectral regions which are little affected by the analyte. This improvement is also observed using PCA where the variance described by a single principal component is increased compared to SNV. This approach is however limited in the range of analytes it can be used with.

Spacing The aforementioned methods aim at improving model performance by fundamentally distorting the input space to maximise linearity in the output space. But it is also important to consider how the input space is sampled, as model input vectors contains no information on units, the only requirement that they are consistent. In DRS, the input space may either be sampled linearly in units of wavelength, or units of frequency. The conversion between these two quantities may be possible via piecewise interpolation functions. Scaling of the input space is not required as the reflectance spectrum is a ratio of intensities, and therefore the scaling factors described in section 1.4 ratio to unity.

In this application, the input space will be linear in frequency units as this is the raw output from the FT spectrometers used in this study. Interpolation to linear spacing in wavelength units was found to cause little difference to model performance, and therefore only serves to reduce the amount of information present in the dataset. It is still however an important consideration when applying a model to the output of a new reflectance spectroscopy instruments.

Combinations It is also possible to use the aforementioned techniques in succession to combine the advantages of each individual method. The combination of pre-processing steps which will be used in this initial model will need to be as effective as possible at removing device dependent artefacts, as the dataset is of a relatively small size and collected on a single type of reflectance spectrometer.

The combination which was found to give the best results is a Kubelka Munk transformation, followed by SNV and finally EMSC. The result of these transformations is visible in table 3.3 (10), where spectra show excellent overlap in the spectral regions least affected by analyte. The PCA plot shows good clustering of identical samples. t-SNE shows the most even clustering throughout the concentration range of the analyte compared to any single pre-processing method applied individually. This combination

of pre-processing techniques will therefore be used in fitting a model to the cotton calibration dataset

Model fitting

Predicting stain concentration from an input set of pre-processed reflectance spectra requires a function which can transfer a vector of input variables into a single predicted stain concentration variable.

A linear transformation could theoretically achieve this in a single step by using linear regression to fit a hyperplane, with each preprocessed spectral channel as a discrete input variable. The resulting function would take the form:

$$\begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ \dots \end{bmatrix} = \begin{bmatrix} 1 & r_{1,1} & r_{1,2} & r_{1,3} & \dots \\ 1 & r_{2,1} & r_{2,2} & r_{2,3} & \dots \\ 1 & r_{3,1} & r_{3,2} & r_{3,3} & \dots \\ \dots & \dots & \dots & \dots & \dots \end{bmatrix} \begin{bmatrix} \beta_0 \\ \beta_1 \\ \beta_2 \\ \dots \end{bmatrix} + \begin{bmatrix} \epsilon_1 \\ \epsilon_2 \\ \epsilon_3 \\ \dots \end{bmatrix} \quad (3.10)$$

where $r_{m,n}$ is the preprocessed DRS spectrum data of the m^{th} sample and reflectance at n^{th} wavelength. y_m is the set stain concentration of the m^{th} sample. β is a vector of coefficients which needs to be determined by minimising ϵ , a vector of residuals.

The size of n is determined by the instrument used to collect the reflectance spectrum. In this study the FT-SWIR spectrometer used has an effective spectral range of 900 nm to 2500 nm, and the bandwidth is limited to 8 cm^{-1} . This means 160 discrete channels are required to effectively sample the input space, combined with over-sampling, this yields over 900 reflectance variables recorded in each spectrum per sample and therefore the same amount of β variables to determine. With a dataset of only 34 samples, this approach is difficult as the number of coefficients which need to be fitted is far greater than the number of samples in the dataset. Ultimately, this leads to a over-fitted

model which performs poorly on samples external to the training set. Due to the constraints described above, it is therefore necessary to reduce the number of variables input into the model, and a PCR approach has been utilised in this section of work. The implementation will first be discussed before gauging its accuracy.

It was previously shown (section 1.5.1) that the majority of the variance in a dataset can be described by only a few PCs using a dimensionality reduction technique such as PCA (figure 3.7). In PCR, samples' corresponding PC scores are used for model inputs in place of pre-processed DRS spectra. The transform function behind such an approach is nearly identical to equation 3.10, where the reflectance variables are substituted by PC scores.

The simplified inputs do not need to sacrifice on a large amount of information to greatly reduce the number of input variables. The use of this operation is also necessary as inputs into a machine learning model should be orthogonal, or free from multicollinearity, which has been inherently introduced in the oversampling used during measurement of the dataset.

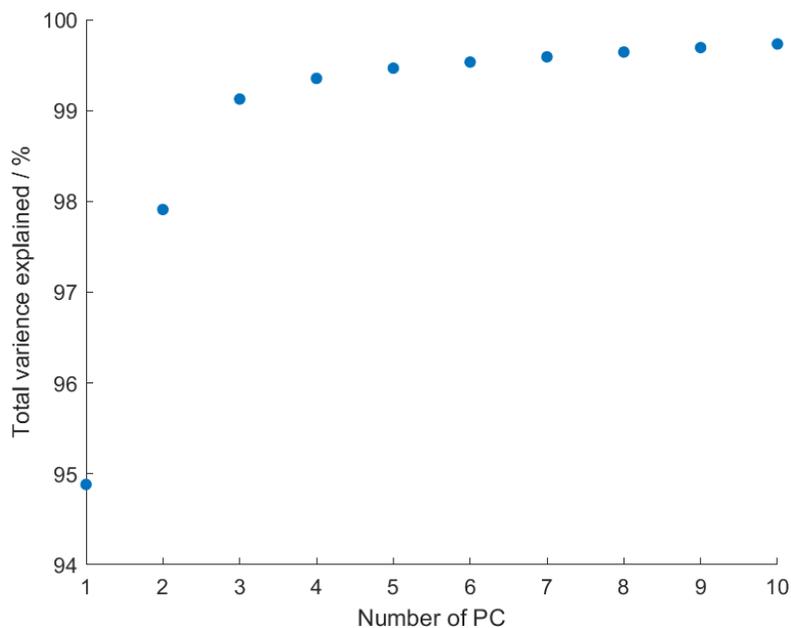


Figure 3.7 Percentage variance cumulatively explained in the pre-processed cotton training dataset using an increasing number of PCs.

It was found that in the pre-processed cotton dataset, almost 95 % of the variance can be described by a single variable, the score of the first PC (PC1). If the first 3 PCs are also used, then the total amount of variance described by the three variables increases to over 99 %, and diminishing returns then show that the PC scores below the 3rd rank add very little information to the data presented to the model input.

The function of PCA can be further investigated by inspecting the coefficients used to calculate the scores of the most significant PCs. Figure 3.8 shows the weighting of each input variable used to calculate the scores. PC1 weightings resemble the absorbance spectrum of the analyte, whilst the weightings forming PC2 appear similar to the absorbance spectrum of fabric. PC3 weightings show features present from both fabric and analyte which are difficult to interpret, but still contain information which will be used in model fitting.

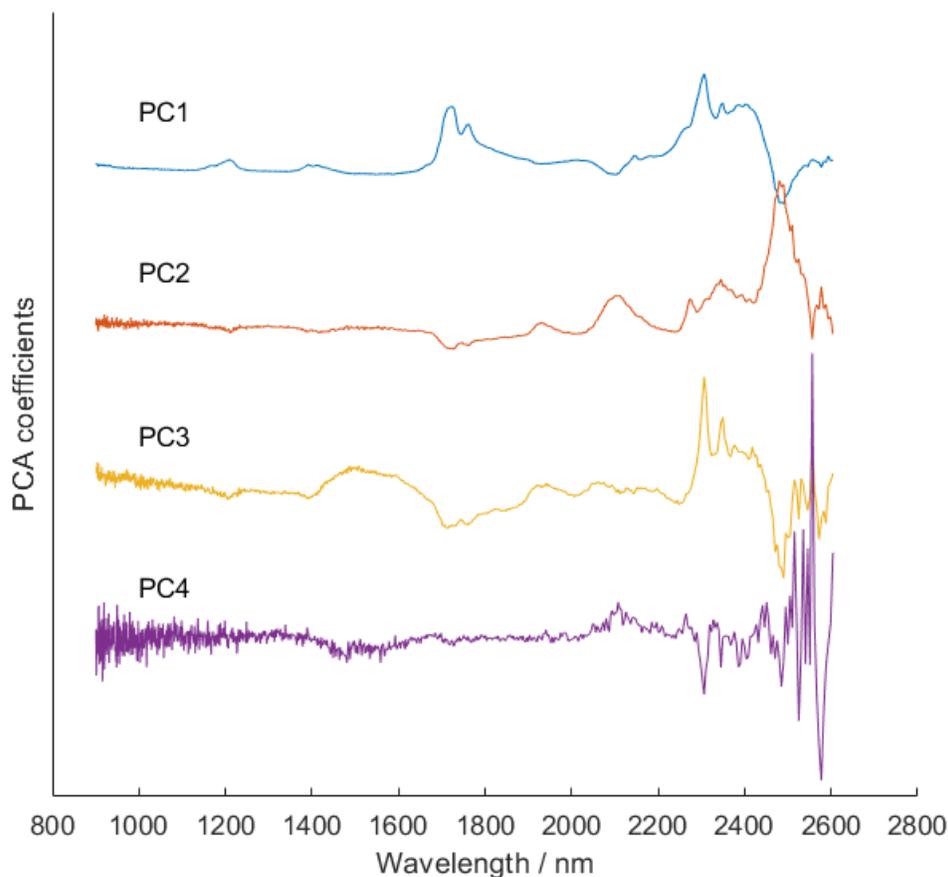


Figure 3.8 PCA coefficients calculated from the reflectance spectra of the cotton calibration set. Each trace shows the coefficients for a corresponding PC, which are used in linear combinations with each spectrum to calculate the PC scores. The coefficients are mean-centred and have been offset for clarity.

The coefficients of PC4 show large contributions from the extremities of the spectral range scanned, where SNR is the lowest. The inclusion of PC4 in the model will therefore likely result in over fitting as these spectral regions also contain little information about the substrate or analyte. The remaining 30 PC coefficients appear similar to PC4, and are too noisy to meaningfully interpret.

The vast majority of the information present in the dataset is contained in a single principal component. A chemometrics model can therefore be obtained by fitting a function to the score of each sample with respect to its set concentration of analyte. The results (figure 3.9) appear to show a linear relationship, meaning a first order polynomial can

be used to form the chemometrics model. This result has been facilitated by the pre-processing step performed on the dataset which increases linearity of the response with increasing analyte concentration. The calibration samples at 10 wt % appear lower than the trend in the rest of the dataset, this will be discussed further in section 3.8. A single calibration sample at 70 wt% appears offset from the trend, potentially due to an error in its preparation.

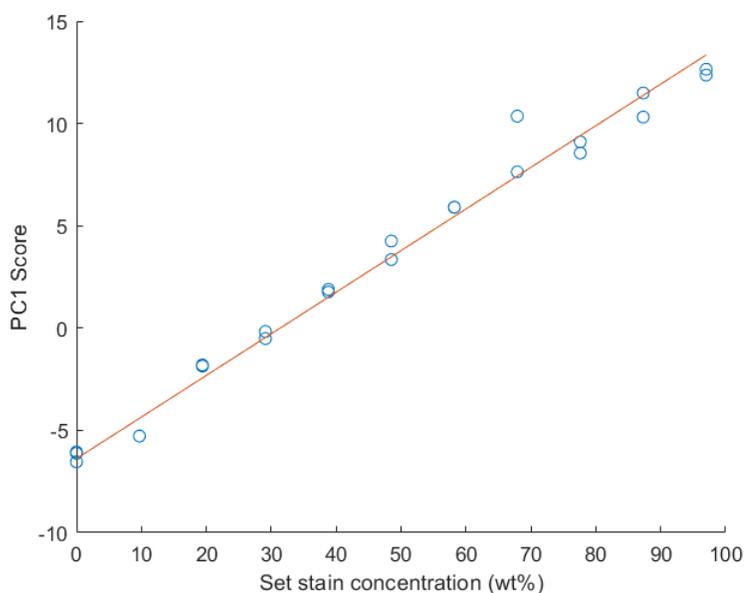


Figure 3.9 Basic linear PRC model fitted to the cotton calibration dataset.

The amount of information input to the model can be increased by including the scores from the remaining principal components as separate channels input to the model. Including the score from PC2 adds an additional dimension to the distribution, which can be fitted to a plane via a double first order polynomial function. The resulting transfer function is therefore a plane with the scores of PC1 and PC2 in the x and y axes, and stain concentration in the z axis:

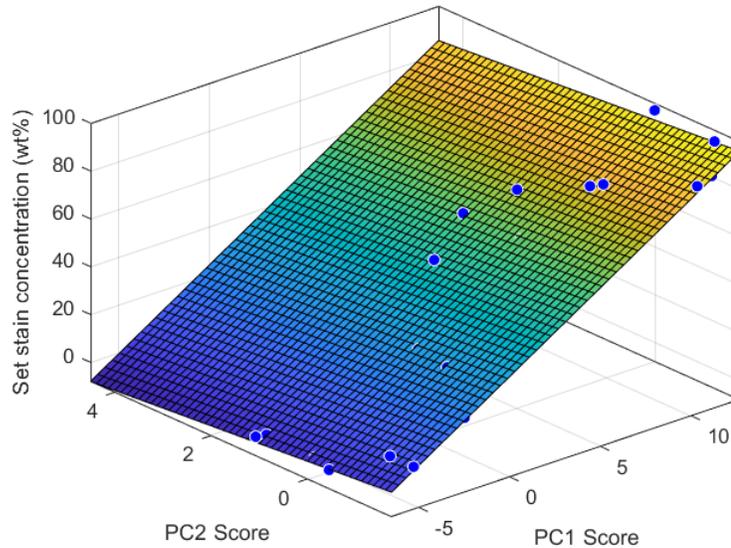


Figure 3.10 PCR model fitted to the first two principal components calculated from the cotton calibration sample set using PCA.

The model can be made more complex by including the scores of even more PCs into the input, hence using more of the information captured in the reflectance spectra in model fitting. The resulting transfer function fitted to a variable amount of PC scores then takes the mathematical form of a hyperplane, similar to that described in equation 3.10 with PC scores replacing the preprocessed reflectance values. The fitting procedure remains unchanged: the least squares technique is employed to optimise the coefficient vector by minimising the vector of residuals.

A series of five PCR models were fitted to the cotton dataset, with each model using an increasing number of PC score variables as inputs. The fitted functions forming the first two models have already been discussed. The remaining PCR models are difficult to visualise as they span into more than three dimensions, they are instead presented in terms of their performance. The accuracy of a model can be initially assessed by the RMSE calculated with 3-fold cross-validation, and these have been tabulated in table 1.14. It can be seen that a simple single component model can achieve prediction

with an RMSE of 3.8 wt%, and the best performing model fitted was that based on three PC inputs, with a calculated RMSE of 2.3 wt%. The inclusion of additional PC scores may appear to improve the fit by minimising residuals, yet validation shows that models become over-fitted after more than three PC scores are included in the model. The accuracy of prediction subsequently decreases as seen in the RMSE of the four and five component models.

Table 3.4 Error of prediction (from 3 fold cross-validation) as a function of the number of PCA components used in fitting of a PRC model to the cotton calibration sample set. * A three component model fitted using the same dataset but with the absence of pre-processing.

Number of PCs used for model	Model RMSE (wt%)
1	3.8
2	3.2
3	2.3
4	2.6
5	2.7
3*	4.2

The value of pre-processing raw reflectance data can also be observed. A separate model has been trained using the same architecture as the most accurate model trained on pre-processed data. It can be seen in table 3.4 (3*) that the RMSE of prediction of this model is 4.2 wt %, significantly higher than the 2.3 wt% found for an identical model fitted to pre-processed data. The use of pre-processing has therefore reduced the RMSE of prediction by 1.8 wt%. This is a relative increase in model performance of 42 %, which can be considered as a rough metric of pre-processing efficacy.

The three component model which shows the best overall performance, or the lowest RMSE, will be used for the rest of experiments in this section. The application of this model to the set of samples used in its training can be seen in figure 3.11, where the model has been applied to each sample in the calibration set. A perfectly fitted model

would show a $y = x$ relationship.

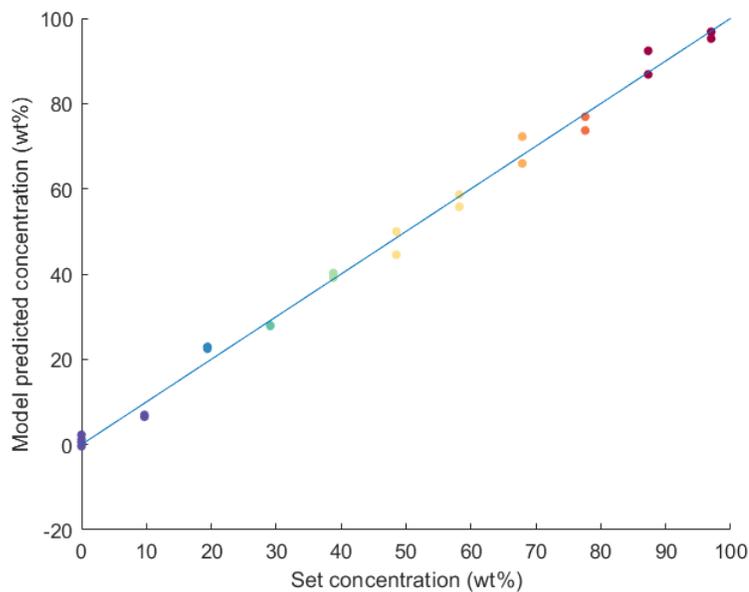


Figure 3.11 Predicted stain concentration from a calibration set of samples using a three component PCR model.

The prediction accuracy shows a linear relationship throughout the concentration range of the analyte. Further analysis of the fit can be done by looking at the distribution of residuals. The residuals show a random distribution around the ideal $y = x$ line, and has a standard deviation of 2.1 wt %, which is consistent with the 2.3 wt % RMSE calculated during cross-validation.

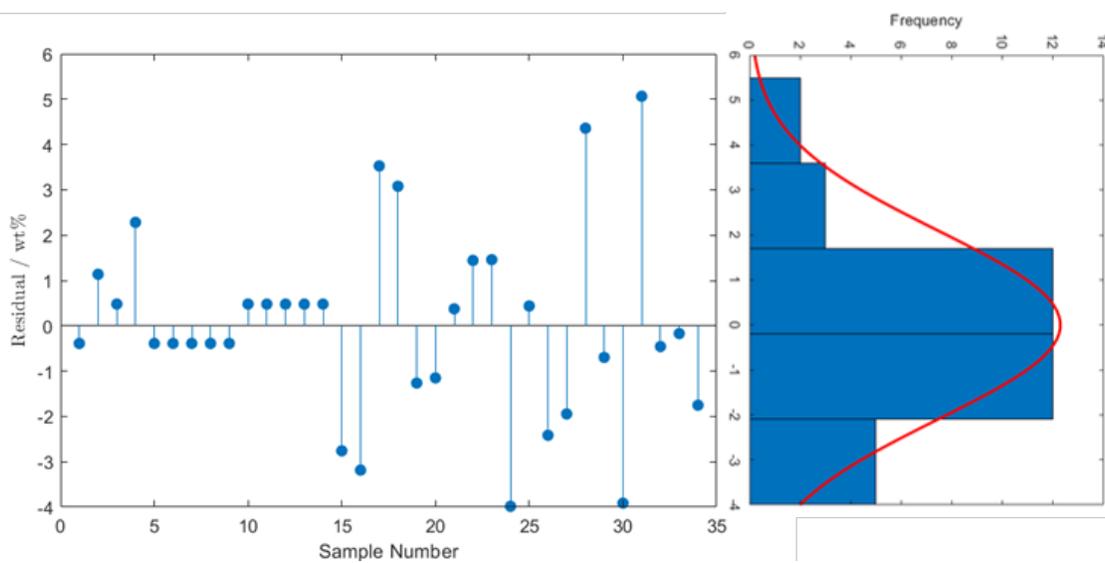


Figure 3.12 (Left) Residuals plot from 3 component PCR model fit with the x axis referring to the index of the reflectance spectrum in the dataset. (Right) Histogram of residuals (blue) with Gaussian distribution fit (red).

Error analysis & assessment of model accuracy

It is important to assess the accuracy and reliability of predictions output by the fitted PCR model by analysing the errors associated with prediction. The accuracy of the fitted model is limited by the magnitude of the errors in the dataset used to create it. The dataset used in this trial was estimated to have a standard error of 3% associated to the set surface lipid concentration. For comparison to the calculated RMSE, an extra calculation step is necessary.

To calculate the error contribution from practical construction of the training set, the error in each sample of the set can be combined to measure its contribution to the RMSE calculated during fitting. This has been done as an RMS sum of absolute errors from each sample in the dataset.

In the calibration set, summation of the 3 % preparation errors across all samples used yields an RMSE contribution of 1.8 wt %. This is lower than the 2.3 wt % RMSE of prediction of the model found through cross-validation. This is a good indication that the PCR model is not over-fitted and suggests that the accuracy may be improved using a more complex model, or there is a significant source of unaccounted for error in the dataset.

To investigate this further, the quality of the fit can be assessed by weighting the residuals to standardise the preparation error contribution through the analyte concentration range, and the outcome of the procedure is displayed in figure 3.13. The resulting plot would ideally show an even distribution throughout the range, yet in this case it is observed that low concentration samples show an unusually large deviation, suggesting an unaccounted for source of error.

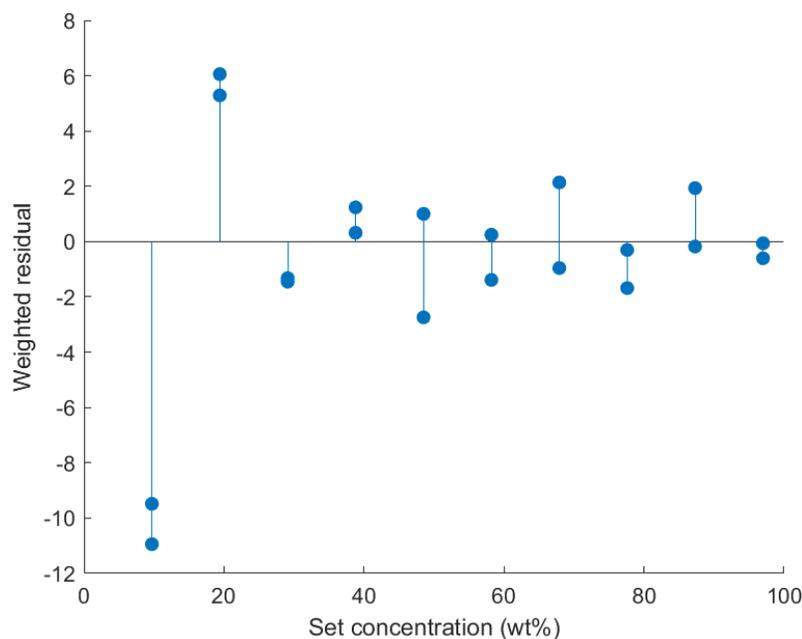


Figure 3.13 Weighted residuals from three component PCR model applied to the cotton calibration dataset. The residuals have been scaled to account for the error in sample preparation by dividing the value of each point by the error in its associated set concentration.

These unusually large errors on low concentration samples is likely a manifestation of the error associated with the dilute solution staining method used for sample preparation. The concentration gradient resulting from increased evaporation at the fabric swatch edges leads to a lower model stain concentration at the center of the swatch compared to the sides. Samples at 20 wt % appear positively biased, which may be a result of model non-linearity due to the unaccounted for errors in the 10 wt % samples.

Aside from accuracy, another important limitation of the model is the performance at extremes of the input range. The smallest concentration of model stain which can be reliably detected is called the limit of detection (LOD). The smallest concentration of model stain which can be reliably quantified is called the limit of quantification (LOQ). LOD and LOQ can be experimentally determined by measuring the range of results when blank samples are repeatedly scanned, with the assumption that the prediction model is linear throughout the analyte range. LOD and LOQ are also device dependant, as each instrument will have unique fluctuations in the output of varying

magnitude. The LOD can be defined as 3 x standard deviation of the blank, and limit of quantification as 10 x standard deviation of the blank.¹⁴⁷ A further set of 45 clean cotton swatches was measured and the PCR model used to infer the lipid concentration, the results can be seen in figure 3.14.

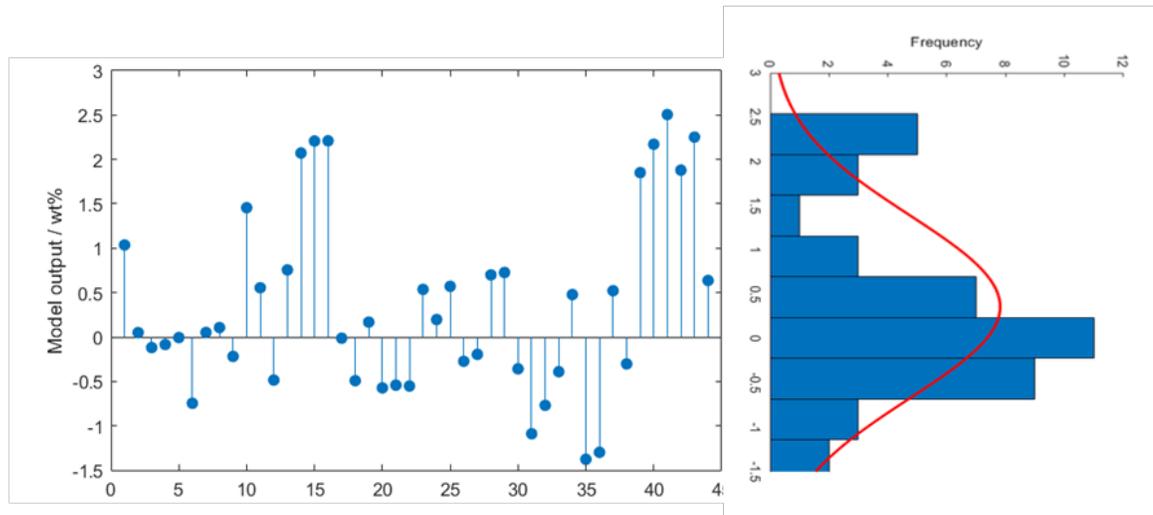


Figure 3.14 (Left) Output from 3 component PCR model on 45 blank samples more with the x axis referring to the index of the reflectance spectrum in the dataset. (Right) Histogram of distribution (blue) with Gaussian distribution fit (red) of mean 0.36 wt % and standard deviation 1 wt %.

The distribution of measurements on blank samples yields a mean of 0.36 wt% and standard deviation of 1.03 wt %. The LOD is therefore calculated as 3.5 wt %, and LOQ calculated as 10.7 wt %. The distribution also appears to show a small cluster of samples with an abnormally high predicted concentration, however, the reason behind in still unclear.

3.4.2 Multi-fabric model

Preprocessing

Prior to modelling the reflectance spectra measured on the multi-fabric calibration sample set, it was again necessary to process the raw reflectance data in order to aid model fitting. The combination of techniques which was found to be effective was conversion to KM space, followed by standardisation of the resulting spectrum. The result of this transformation on each type of fabric in the calibration sample set, as well as a selection of representative samples with different levels of staining, can be seen in figure 3.15.

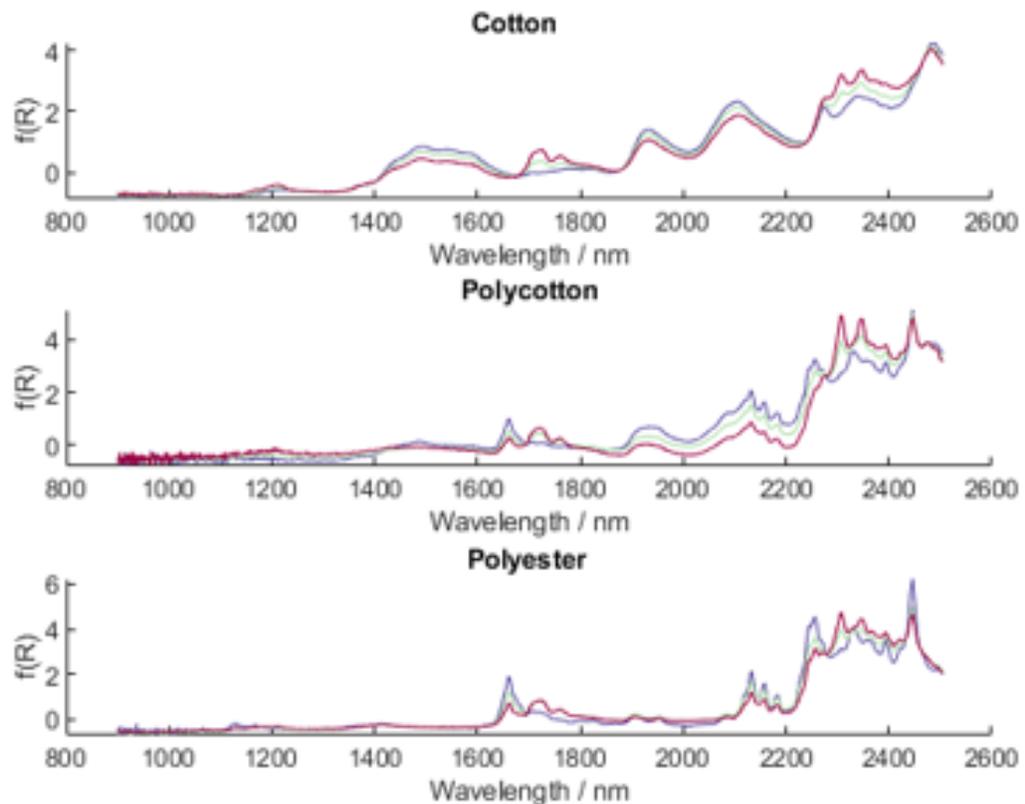


Figure 3.15 Representative spectra of cotton, polycotton, and polyester in the calibration sample set after pre-processing by KM transformation and SNV, at set surface lipid concentrations of: 0 wt% (blue), 30 wt% (green), and 60 wt% (red).

The effect of the preprocessing procedure was assessed using unsupervised machine

learning techniques, PCA and t-SNE, the application of which is displayed in figure 3.16. The PCA scores plot shows a good improvement in structure, with samples broadly forming three gradients. It is interesting that without preprocessing spectra, PC1 mainly describes the lipid concentration while PC2 varies most with lipid concentration. This pattern is reversed after KM and standardisation, along with a reduced spread for concentration gradients on the same type of fabric. t-SNE plots show separation between set of different fabric type while concentration gradients remain closely packed. Overall, these results show that the choice of preprocessing method in this case is valid.

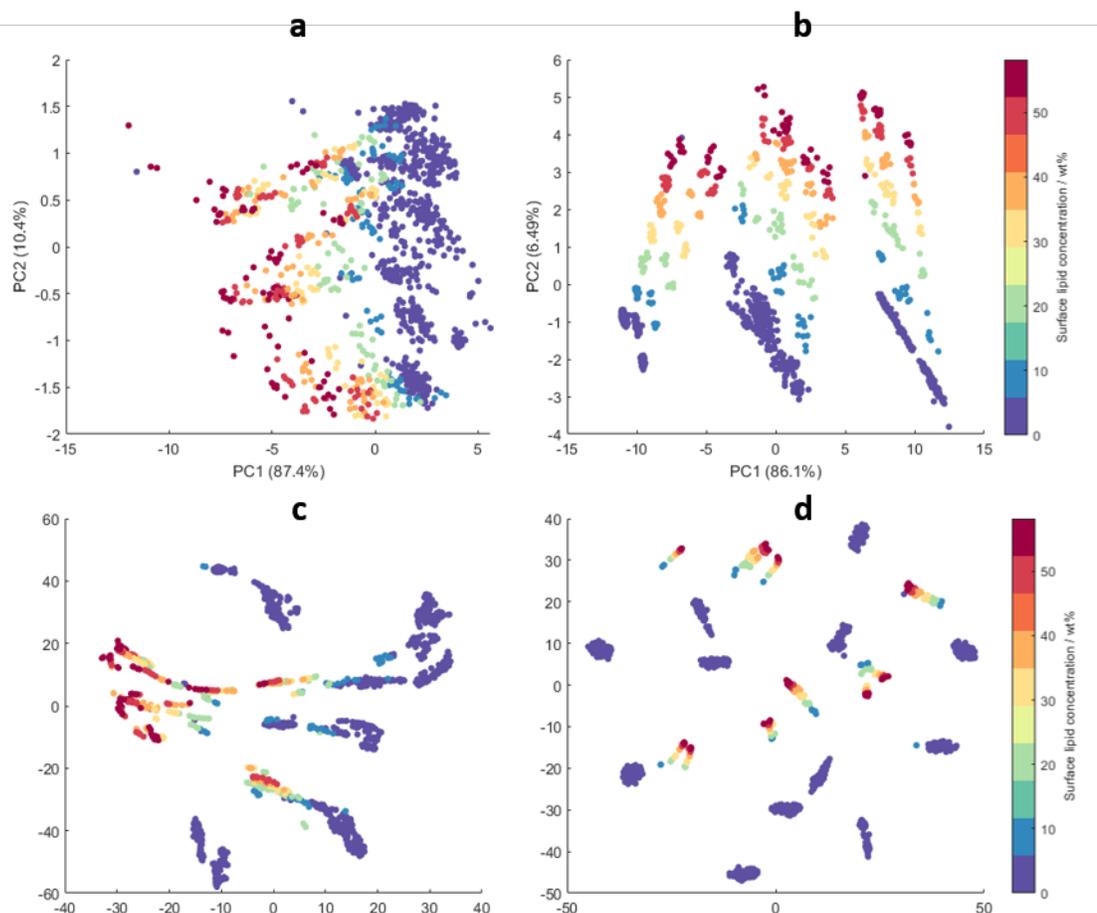


Figure 3.16 (Top) PCA scores plot of raw spectra in advanced training set prior to preprocessing (a) and after preprocessing (b), the value given in the parentheses refers to the percentage variance described by the corresponding PC. (Bottom) t-sne before (c) and after (d) pre-processing.

Model fitting

In modelling lipid concentration from the preprocessed multi-fabric dataset, the same limitations which made it necessary to use dimensionality reduction prior fitting in the cotton set, apply here. PCA was again employed here and the cumulative variance which can be explained by an increasing number of PC can be seen in figure 3.17. In the preprocessed multi-fabric dataset, 86 % of the variance can be explained by the first PC. This is a large majority, although far lower than the 95 % calculated for the cotton dataset, but this is to be expected given the added complexity of the sample set. The plot also shows that the increase in amount of variance explained as extra PC are included stagnates after 5 PCs.

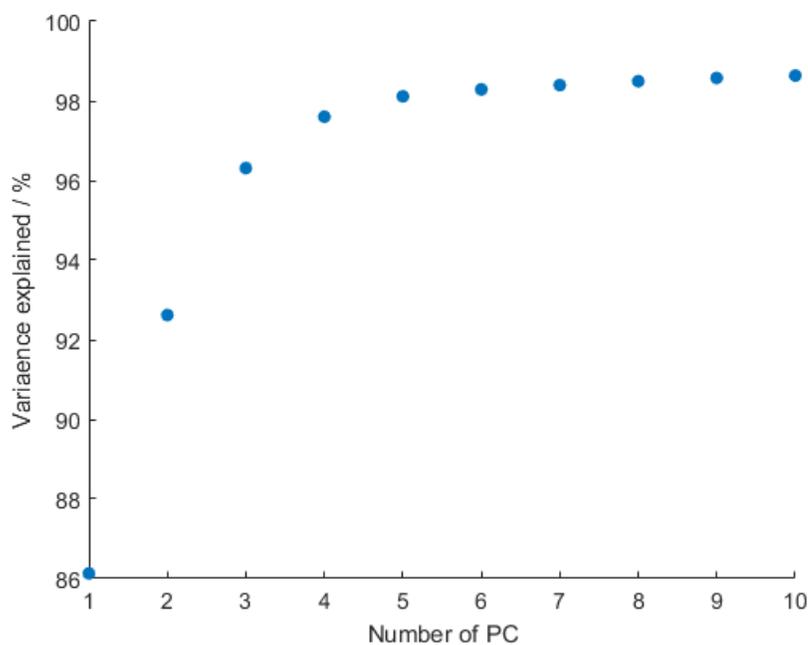


Figure 3.17 Percentage cumulative variance explained in the preprocessed multi-fabric dataset as a function of increasing number of PCs.

Examination of the coefficients used to calculate the PC scores have been plotted in figure 3.18 to investigate the workings of PCA in extracting information from spectra in the multi-fabric set. The coefficients of PC1 show features most prevalent in polyester and cotton, which is why PC1 best differentiates different fabrics. PC2 appears to have

a major component of features observed in the optical spectra of lipids (section 3.3.4), as well as minor features attributed to the different fabrics types. The remaining PC coefficients resemble complex shapes with features similar to those found in the optical spectra of model compounds, and the contribution of the low-SNR signal extremities increases with decreasing PC rank.

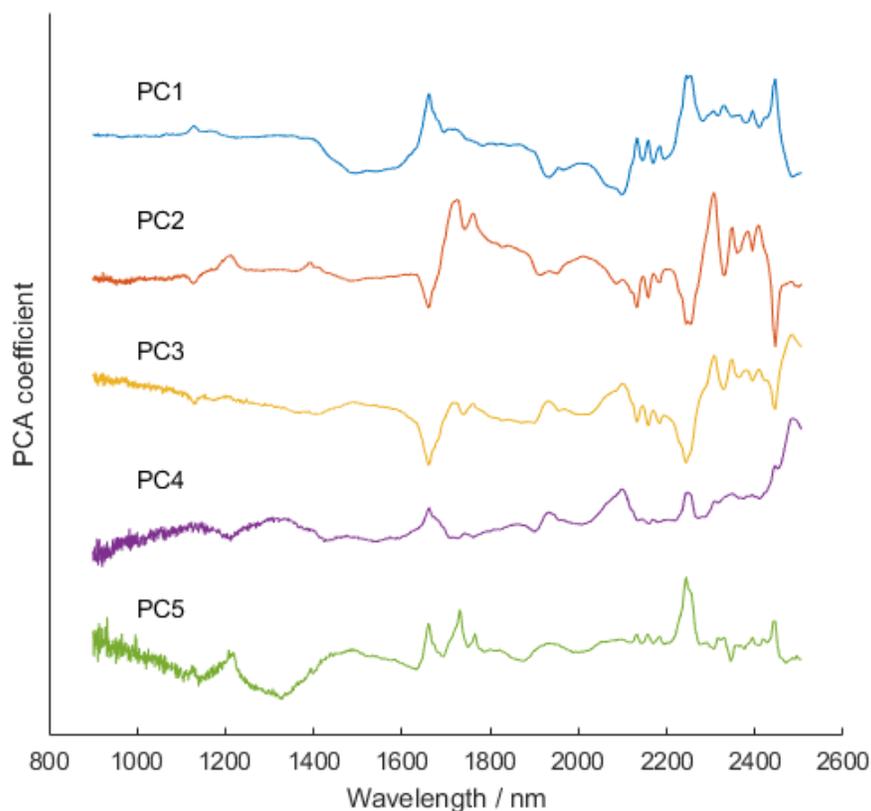


Figure 3.18 PCA coefficients used to calculate the first 5 PCs on the pre-processed multi-fabric calibration data set. Data has been standardised and vertically offset for clarity.

Attempts to fit PCR type models to the preprocessed multi-fabric dataset were unsuccessful, and the models were found to perform poorly in validation. Gaussian process regression (GPR) is a different machine learning technique which was found more suitable to this application. Five GPR models were fitted to the pre-processed multi-fabric dataset with an increasing number of PCs included into the model. A linear basis function and squared exponential kernel were chosen for the GPR model, and the RMSE of

prediction calculated from 5-fold cross validation for each model has been summarised in table 3.5.

Table 3.5 Error of prediction in GPR models trained using a varying number of PC fitted to the preprocessed multi-fabric calibration dataset. RMSE was calculated using 5-fold cross validation.

Number of PCs used for model	Model RMSE (wt%)
1	13.3
2	3.9
3	1.8
4	1.8
5	1.9

The fitted models show a drastic increase in performance as the second and third PC is added to the input. The highest performing models were found to be those fitted to three and four PC inputs, with a RMSE of prediction at 1.8 wt %. The model using three PC scores will be used in further testing as it is less likely to be prone to over-fitting than the four component model. This is an improvement on the 2.3 wt% RMSE calculated for the cotton model, which likely stems from the increased size of the dataset as well as the improvements to the sample preparation method previously described in section 3.3.3.

The output of the model applied to its own training set can be visualised in figure 3.19, along with the residuals to the ideal $y = x$ line. The fit shows a good distribution of residuals evenly spread around zero, and highest deviation recorded is on the order of ± 8 wt %.

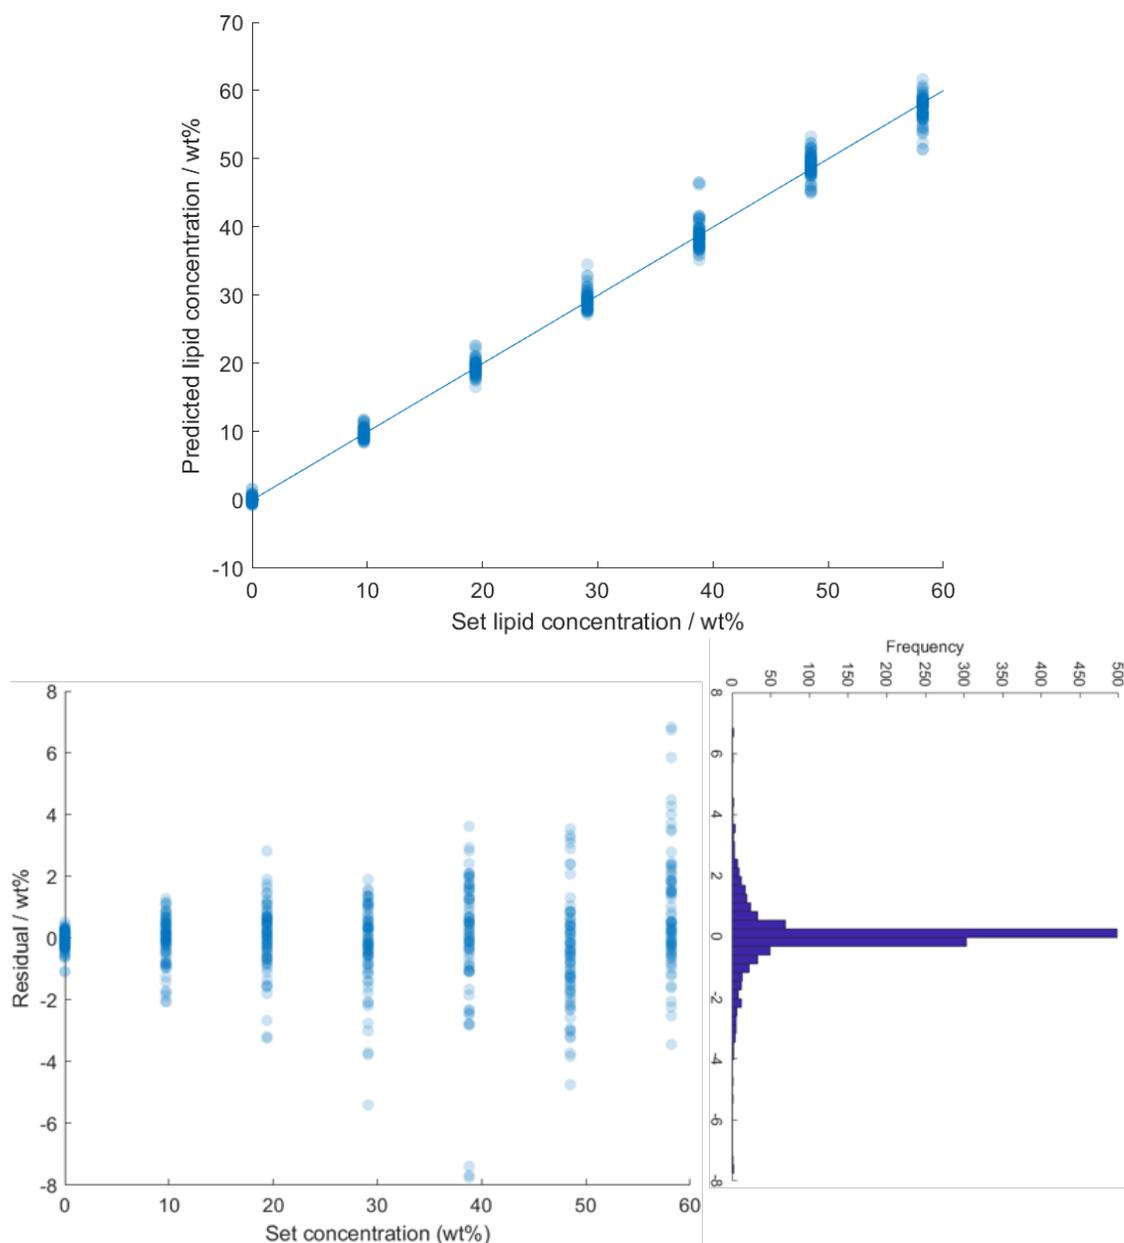


Figure 3.19 (Top) Predict lipid concentration from three component GPR model on multi-fabric calibration sample set. (Bottom-left) Residual plot from ideal $y = x$ fit. (Bottom-right) Histogram from residuals distribution. Points have been plotted with 20% opacity.

The residuals can be further analysed by weighting them in order to standardise the contribution of the 3% preparation errors. The outcome can be seen in figure 3.20 where the distribution of the points shows no significant pattern, yet at lower stain concentrations the error rate appears to be higher.

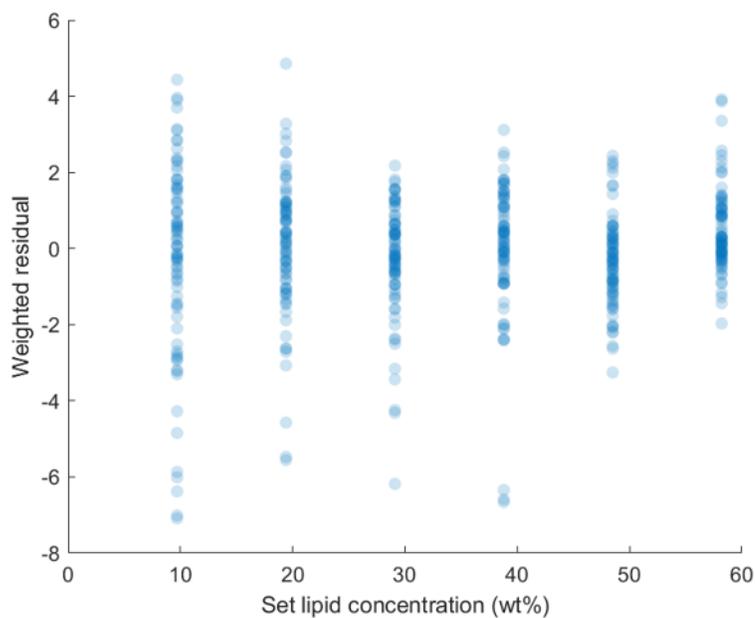


Figure 3.20 Weighted residuals from fitted gpr model prediction on training dataset

Looking at the results of the model on unstained fabric (figure 3.21) shows no significant sign of bias in the model for any particular fabric, as the mean for each batch of fabrics remains very close to zero. The standard deviation of the set was calculated to be 0.16 wt %, leading to a corresponding LOD of 0.5 wt % and LOQ of 1.6 wt%, although these values were calculated from the dataset used in training and therefore require further validation.

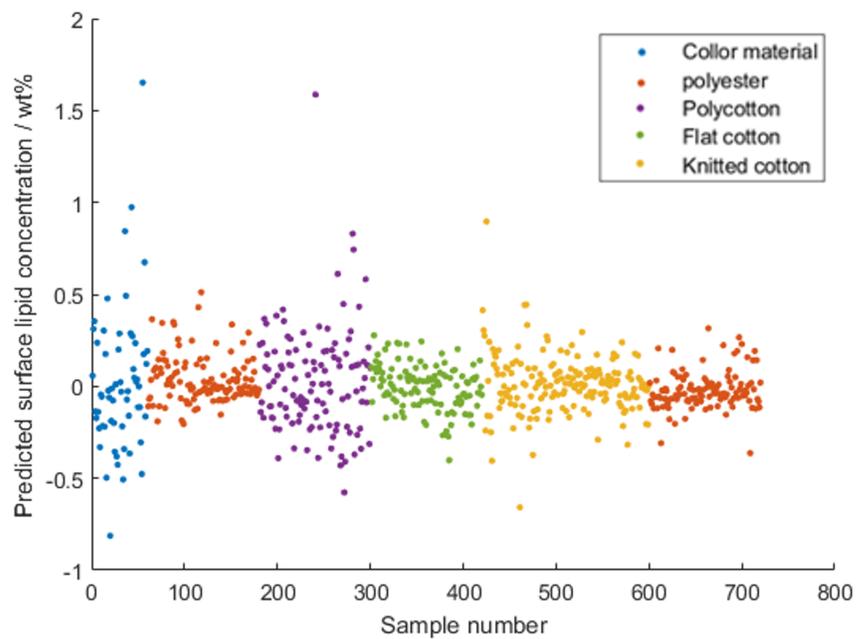


Figure 3.21 Output from three component GPR model on clean samples in the multi-fabric calibration sample set.

3.5 Model testing

The validity and workings of the developed lipid prediction model (three component GPR type) were further examined. Firstly, a set of samples with a known lipid concentration that was prepared separately to the multi-fabric calibration set will be analysed, followed by investigation into the depth probed using this new reflectance-based technique.

3.5.1 External set

A set of polyester samples was prepared and measured in the same manner as that used for the multi-fabric calibration set, the set consists of 34 individual samples (14 stained / 20 clean) measured at 3 points, leading to 102 spectra. The dataset used to train the cotton model, which was not presented to the GPR model during training, was combined with this new polyester set to create an external set of samples for testing the model. The resulting predictions by the developed three component GPR model are displayed in figure 3.22. The RMSE of prediction on samples within the concentration range of the GPR model was calculated to be 2.1 and 2.0 wt % in knitted cotton and polyester respectively. This is marginally larger than the 1.8 wt % RMSE of prediction calculated during fitting using cross-validation, which indicates that there may be a small degree of over-fitting.

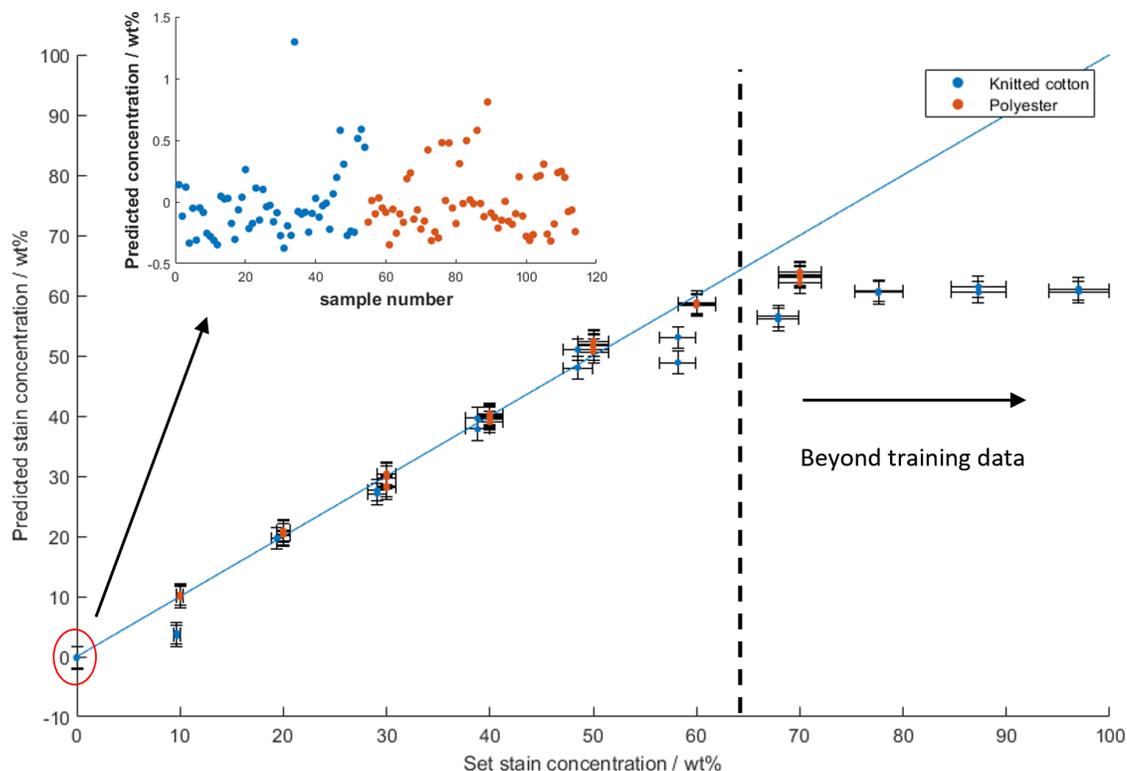


Figure 3.22 Predicted lipid concentration of three component GPR model applied to the cotton calibration and an external polyester set. Lipid concentration beyond 68 wt % were not shown to the model during training. The error bars represent the RMSE of the model (vertical) and the error in sample preparation (horizontal).

On unstained samples, the distribution of predictions shows a standard deviation of 0.19 wt %, which leads to an estimated LOD of 0.6 wt% and LOQ of 1.9 wt %, again marginally higher than that calculated from cross-validation during fitting.

The majority of the predicted outputs lie within their associated error of the $y = x$ line, however, several points in the cotton set are outliers which cannot be accounted for. These may be a result of defects in the calibration samples, but with such a small dataset it is difficult to trace.

The larger maximum lipid concentrations set on cotton fabric allow us to test the behaviour of the fitted model to data which is outside the range of its training set. Over-fitted models which are weakly related to physical principals typically perform erratically in such scenarios, however, in this case the output seems to plateau at

around 60 wt %

3.5.2 Depth penetration analysis

In order to understand how the distribution of lipid inside a fabric matrix contributes to the level of staining quantified by DRS and the subsequent data processing, an experiment was devised to measure how the predicted stain concentration varies with the depth at which the lipid is buried in the fabric. A stack of thin tissue paper was used to imitate a cotton fabric substrate, with each sheet in the stack having a thickness of approximately 0.1 mm. The stack consists of 20 sheets in total with a black backing material. Tissue paper was selected as a suitable model for fabric as it is primarily composed of cellulose, and therefore the reflectance spectrum closely resembles that of cotton fabric, making it suitable for analysis by the developed lipid quantification model. A single sheet was stained with hot coconut oil in the centre (figure 3.23), and after allowing the stain to solidify, the stained sheet was placed at the top of the stack and sequentially moved down while the SWIR reflectance of the stack's center was measured at each position of the stained sheet.

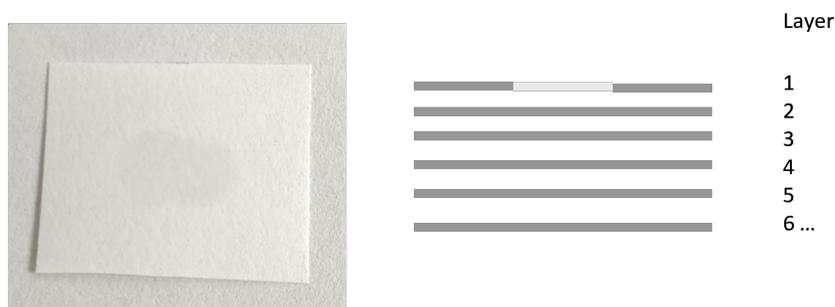


Figure 3.23 Diagrammatic experimental setup used to measure the depth probed by SWIR DRS. (Right) Single sheet stained with coconut oil. (Left) Diagram of stack used to imitate cotton fabric.

The reflectance spectra show an overall increase in reflectance as the stained sheet is placed further into the stack, while absorbance bands associated with the lipid stain are visible in the first 5 samples, with the intensity fading quickly going deeper into

the stack. The predicted stain concentration shows an exponential decrease as the stained sheet is placed further into the material, and the contribution to the measured reflectance spectrum becomes negligible after the stain is around 0.5 mm deep.

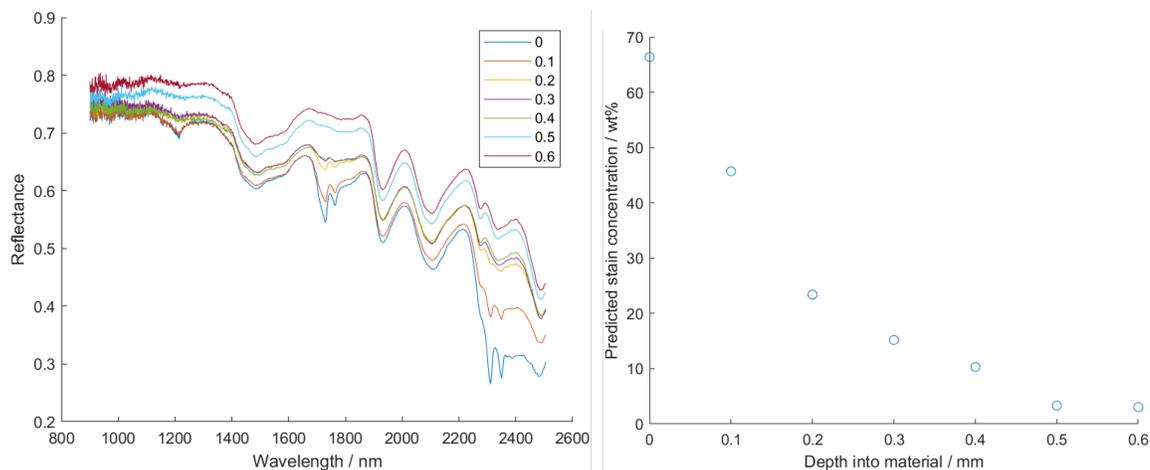


Figure 3.24 (Right) Overlaid reflectance spectra of a thin paper stack used to imitate cotton fabric. A single sheet has been stained with oil and a reflectance spectrum of the whole stack recorded with the sheet at various depths in the stack (displayed in mm in the legend) . (Left) Results from GPR model applied to reflectance spectra of the paper stack, with the stained sheet a varying depth.

The results show that the vertical position of the stain in the fabric matrix has a major impact on how it is quantified, and lipids near the surface dominate the measured response. In this setup, the total depth which can be probed by DRS was found to be around 500 μm , although the cellulose fibres that comprise tissue paper are smaller than typical clothing fibres. This results in a larger degree of scattering encountered in the model material than what would be expected in real fabrics, therefore the penetration depth of DRS in fabrics is likely larger than what was measured in this experiment.

The results also suggest that output from the model should be treated with care as measurements show a non-linear weighting with the depth of lipid stain into the fabric. This means that a single reading can potentially be attributed to either a light lipid soiling near the surface of the fabric, or heavy soiling at the core of the fabric with a clean surface. Although, to understand how the predicted lipid concentration varies with vertical distribution of the lipid requires more in-depth research as the vertically

separated lipid components would almost certainly affect each other.

All spectra recorded from calibration sample, as well as fitted models are available in 5.2.

3.6 Comparison to colourimetric method

Another commonly used technique to measure the lipid concentration on fabric is to stain the oil with a trace amount of dye, and track concentration across the fabric using colourimetry. This approach was discussed in section 3.2.2. In order to compare the accuracy, it is first necessary to explore the assumptions that are key to its function. These will be discussed in-turn:

The center of the stain is representative of the entire stain. When a model stain is applied to a fabric then is it subject to chromatographic effects, and the result is usually an uneven distribution where the stain concentration falls radially from the stain center. The effect can be seen as the change in colour of a single stain in figure 3.3. This effect is the reason that direct application of the model stain was not used for preparation of the calibration sample set.

There is no interaction between the dye and the fabric, or the affinity of the dye to the stain is far greater than that to the fabric or the wash liquor. This must be obeyed as it is the dye being measured and assumed that its concentration is proportional to the concentration of model stain. If the dye has a greater affinity for the fabric, then the oil may be unnoticeably washed out of the fabric. This has been tested experimentally using a chromatography setup, with fabric as the substrate, sunflower oil as the solvent, and oil violet as the dye. The results are shown in figure 3.25, and it can be seen that the retention factor is near unity, with streaking likely due to high concentration of dye which was required to clearly visualise the stain, meaning that the initial assumption is valid.

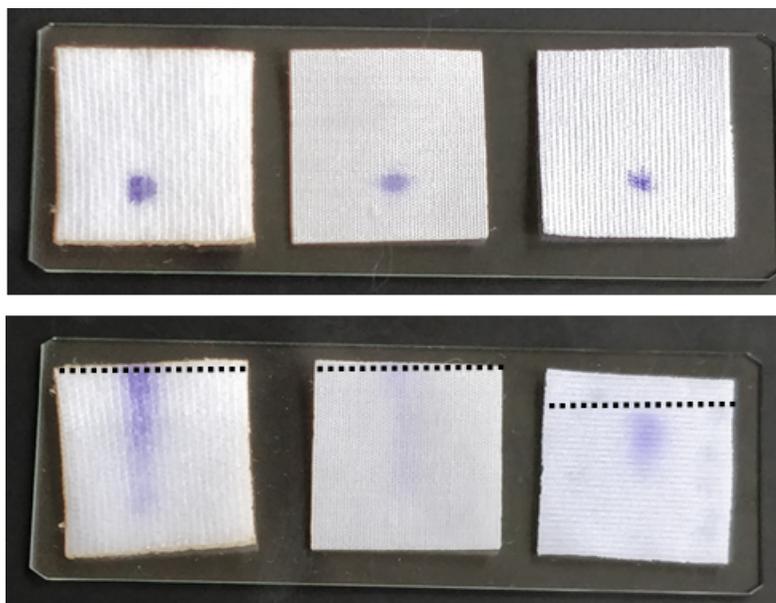


Figure 3.25 Chromatography experiment to test oil violet dye affinity for various fabrics. (Top) Before eluting, (Bottom) After eluting. Substrates left to right: knitted cotton, poly cotton(65/35), and poly ester. Dashed line represents the solvent line.

The dye does not wash out of the stain, spread across the fabric, or degrade during the wash cycle. This is difficult to test without a technique to measure the oil concentration independently of the dye. The problem of dyes transferring between garments during a wash cycle has been an issue in the detergent industry for centuries, and laundry formulations typically include bleaching agents which could potentially discolour the dye without removing the stain. This is therefore a likely source of error in using a colourimetry approach.

The measured ΔE^* between the dyed stain and clean fabric is linearly dependant on the concentration of stain present on the fabric. In order to assess the performance of laundry detergent, the SRI parameter should be linearly dependant on the concentration of stain present. However the CIE $L^*a^*b^*$ colour space was designed to be perceptually uniform as opposed to scaling linearly with colourant concentration. This can be validated by measuring the SRI against a set of fabrics

which have known amounts of stain on their surfaces, the cotton calibration set in this case. $L^*a^*b^*$ coordinates and the subsequent ΔE^* values were measured by diffuse reflectance spectroscopy (figure 3.26).

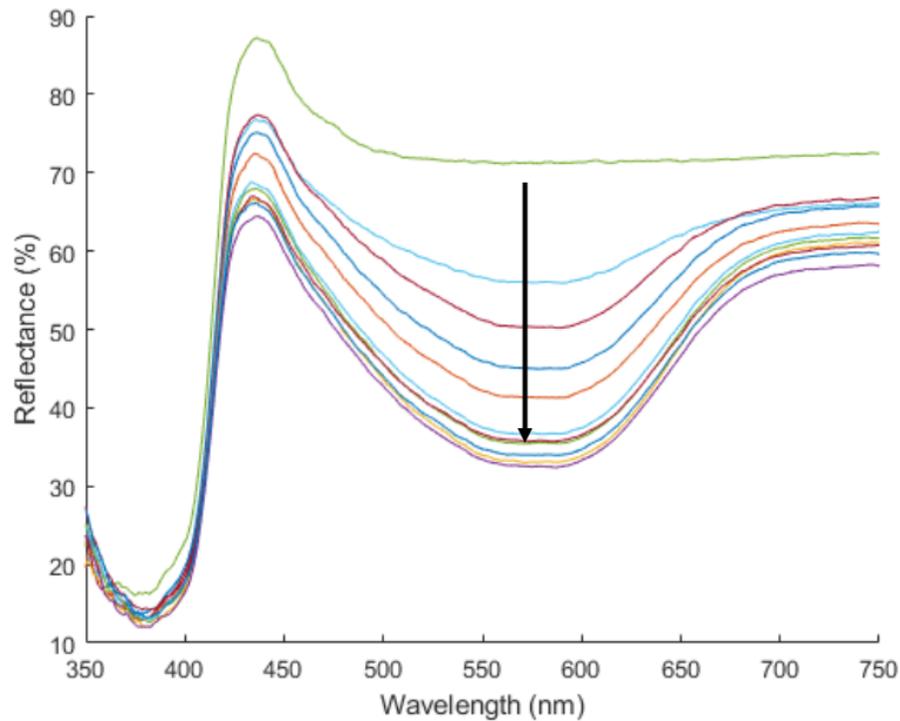


Figure 3.26 Diffuse reflectance spectra of knitted cotton fabrics stained with an increasing amount of oil violet dye (indicated with arrow) from cotton calibration set. Measurements are relative to a Spectralon reference.

The reflectance spectra show two concentration dependent features: a trough at 580 nm which corresponds to absorbance of the oil violet dye and a peak at 440 nm which is attributed to fluorescence from the optical brighteners used to treat the fabric during production. The measured spectra were processed using the CIE system of colourimetry introduced in section 1.3.2 to calculate their colour appearance as described in the $L^*a^*b^*$ space. The difference between the coordinate of clean fabric and the sample in each case was used to calculate the ΔE^* values are shown in figure 3.27. The relationship between ΔE^* and set lipid concentration appears highly non linear, making the SRI metric based upon colourimetry invalid for use in meaningful assessment of lipid removal during a wash cycle.

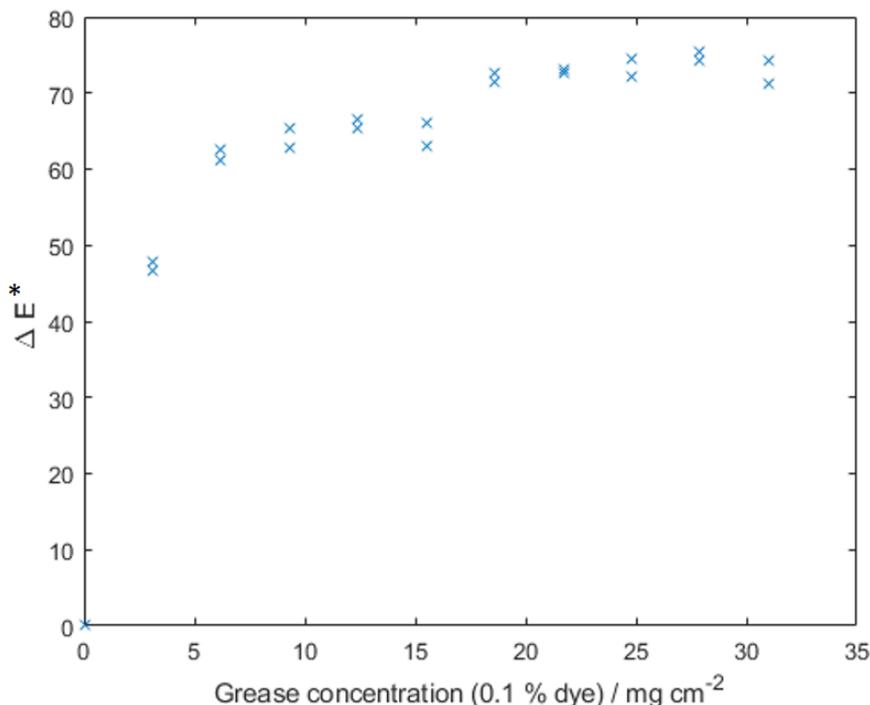


Figure 3.27 ΔE^* values calculated between clean cotton fabric and cotton swatches of set surface lipid concentration, prepared by the solvent evaporation method described in section 3.3.3.

It should also be noted that in this example, the presence of FWA in the fabrics combined with the near-UV component of probe light source used to measure ΔE^* leads to unreliable readings; and this is also the case in other colourimetry setups. Since the surface is emitting light as well as reflecting it with a measured intensity that is proportional to, among other factors, both the near-UV component and the visible component of the probe light source. This can be described by equation 3.11:

$$R(\lambda)_{obs} = \frac{I + I_f}{I_{ref}} = \frac{I}{I_{ref}} + \frac{I_f}{I_{ref}} = R(\lambda) + \frac{I_f}{I_{ref}} \quad (3.11)$$

where I is the intensity of reflected light, I_{ref} is the intensity of the probe light source, I_f is the intensity of fluorescence, $R(\lambda)_{obs}$ is the observed reflectance spectrum, and $R(\lambda)$ is the fluorescence-decoupled reflectance spectrum.

This complex relationship between the reflected light and the probe light source means that the measured ΔE^* values are device dependant, and identical samples measured in different colorimeters would not give comparable results.

Combining these assumptions with the restriction that the colourimetric method can only be used with model lipid compounds stained with dye, as opposed to real-world stains, the colourimetry approach shows little suitability for assessment of laundry detergent performance on lipid stains.

3.7 Imaging

The lipid prediction model developed in this section operates on reflectance spectra which have been collected from a single point around 2 mm in diameter. In order to study the cleaning process on real items of clothing, a more complete analysis needs to be carried out by scanning an area of fabric, and looking at the distribution of a lipid across the surface. This can be achieved by recording a grid of reflectance spectra at each point on the fabric surface, followed by applying the developed model to each reflectance spectrum in turn, giving a concentration map ($c(x, y)$) that can be used to visualise the distribution of otherwise invisible lipids on the fabric. Integration of the concentration map over the scanned area of the fabric can be used to calculate the total mass of lipids detected on the fabric (m_{stain}):

$$\int_{y_{min}}^{y_{max}} \int_{x_{min}}^{x_{max}} c(x, y) dx dy = \frac{m_{stain} 100\%}{\rho_a} \quad (3.12)$$

where x and y are the Cartesian coordinates of the scanned area and ρ_a is the area density of the fabric. In artificial stains, the total mass of lipid detected on the fabric can be experimentally measured using gravimetric analysis.

The SWIR reflectance instrumentation used in development of the new technique been modified in order to allow imaging of lipid stains on fabric via the addition of an automated two-axis moving stage underneath the reflectance probe (figure 3.28). The stage acts to reposition the sample being analysed so that a grid of reflectance spectra can be collected across the surface. This technique allows for hyper-spectral imaging with high spectral resolution and sensitivity but low spatial resolution, which limits the size of the smallest features that can be resolved to around 2 mm.

For an initial test, 120 μ l of sunflower oil was deposited onto the center of a 5cm x 5cm knitted cotton swatch, with the mass of the swatch measured before and after staining.

The center of the swatch was then scanned with a grid of 9x9 points separated by 5 mm. While the swatch appears to show almost no colouration on the surface, the stain can clearly be identified on the concentration map. The distribution appears plateaued near the center and is quickly reduced to zero at the edges of the stain, showing limited diffusion of the model stain across the fabric. The integral of the concentration map was approximated using trapezoidal numerical integration. The integral of the scanned 16 cm² area was calculated to be 322 wt% cm². Using the area density of knitted cotton (33.0 mg cm⁻²), the total mass of lipid on the fabric swatch is found to be 106 mg, which is a 3.7 % deviation from the applied amount of 110.4 mg.

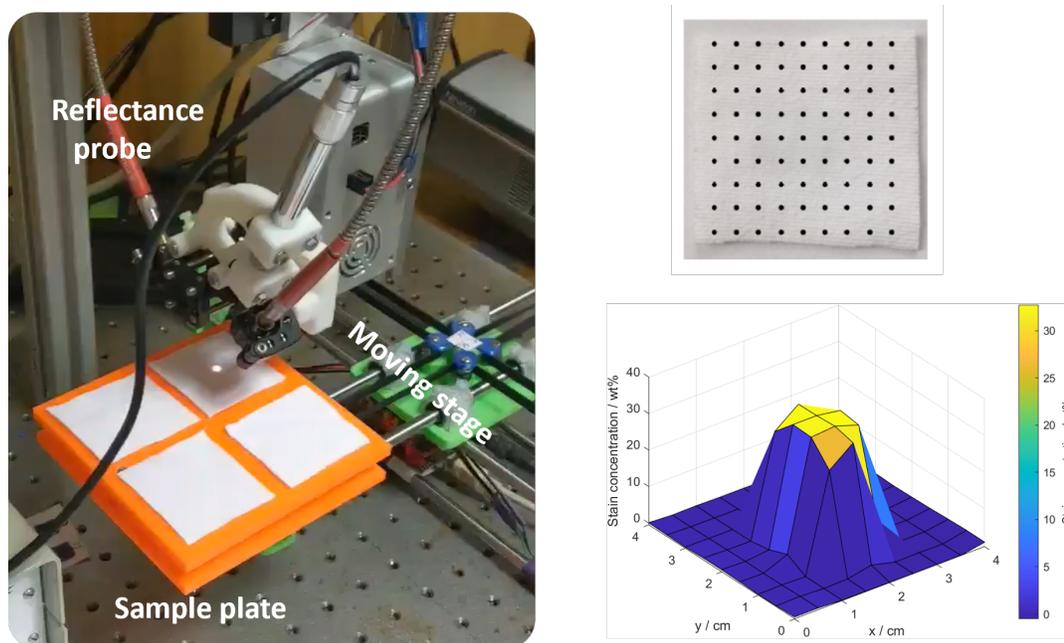


Figure 3.28 (Left) Modified reflectance setup for imaging of stained fabrics. (Right) Knitted cotton swatch stained with 120 ul of sunflower oil with overlaid grid showing scanned points and the measured concentration map below.

Grid scans are useful when the distribution of soil on a garment is completely unknown, however, they are also quite time consuming and not always necessary. The symmetry of artificial stains used in wash tests was found to provide a convenient way of simplifying imaging requirements. Stains tend to be symmetrical around their center and spread at different rates in directions of the weave and against the weave, leading to an elliptical distribution of the soil. This distribution can be probed without a full grid of points,

instead a cross of points can be scanned with the stain aligned to the center of the cross. The diagonal elements can then be estimated using linear interpolation in polar coordinate space. This radial form of interpolation is shown in figure 3.29 where the stain concentration at a query point q (c_q), which is equidistant from the center of the cross as scanned points a and b , can be described as:

$$c_q = \left(1 - \frac{2\theta}{\pi}\right)c_a + \left(\frac{2\theta}{\pi}\right)c_b \quad (3.13)$$

where c_a and c_b is the stain concentration at points a and b respectively, and θ is the angle at the center of the scanned cross between a and q . This technique was applied to the initial test swatch using 14 out of the 81 scanned points which form a cross aligned to the center of the stain; a further 28 points were then calculated using interpolation and the resulting concentration map is shown below.

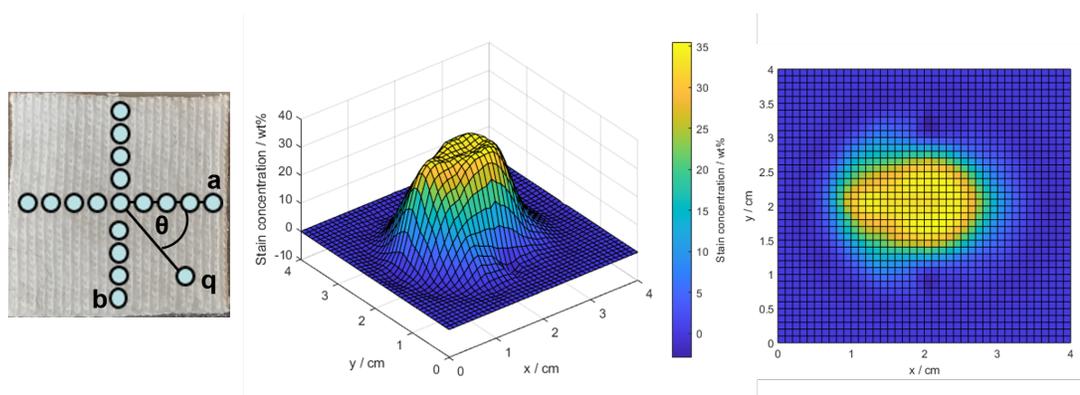


Figure 3.29 (Left) Diagram of scanned points used for radial interpolation. (Middle / right) Interpolated concentration map of cotton swatch.

The shape of the stain appears very similar to that obtained with a full grid scan, and integration of the surface yields a total lipid mass of 103 mg. This is a further 2.5% error compared to the integral calculated using a full grid of point, however, this measurement requires over 80 % less data as an input. The number of points to scan also scales linearly with the size of the scan ($n \times n$ points) in cross scanning, as opposed to grid scanning where the number of points to scan increases with the square of n .

Cross scanning in combination with interpolation can therefore be used to greatly increase the sample throughput at the expense of accuracy. The accuracy of quantitative imaging was further assessed with a another 4 samples were made up in the same manner as the initial sample, and each swatch was scanned with an 11x11 point cross. Interpolation was then used to estimate a further 132 points and the concentration maps integrated to calculate the total mass measured on the fabric. The results are summarised in table 3.6, the concentration maps all appear similar to the initial test sample and are available in section 5.1. The error was calculated by summation of the relative errors from each point in the image.

Table 3.6 Results of quantitative imaging of stained knitted cotton swatches using point scanning SWIR reflectance spectroscopy.

Sample	Mass of lipids deposited / mg	Mass of lipids measured / mg	error %
1	110.4	103	6.7
2	108.0	106.1	1.6
3	103.8	104.1	0.3
4	109.9	108.0	1.7
5	105.4	98.3	6.7

The results show good agreement between gravimetry and quantitative imaging. The average error in the total mass of lipid on fabric as measured by scanning reflectance spectroscopy is 3.4 % and appears to be mainly weighted at under-predicting the total mass deposited.

3.8 Low-concentration calibration samples

The measurement of lipid concentration on fabric is ultimately limited to the accuracy of the calibration set. As discussed in section 3.3.3, the method for preparation of calibration samples works well on samples with a stain concentration of above 10 wt %. Samples prepared using the solvent-evaporation method with a concentration below 10 wt% tend to have an un-even distribution of lipid across the fabric swatch. The effect is similar to the ‘coffee-ring’ effect observed in particle-laden liquids. When a dilute solution of oil is deposited on a fabric swatch, the edges of the fabric have a larger surface area compared to the center of the fabric, which results in increased evaporation at the edges and in-turn creates a concentration gradient that pulls more oil to the edges of the swatch.

To gauge the effects of this on the concentration of the calibration samples, a set of 20 polyester calibration samples was made using the same approach as that used in making the multi-fabric set. The concentration range was set from 1 - 10 wt %. The samples were measured using DRS at the center of each swatch twice and the developed GPR model applied. The results show that at low stain concentrations the lipid concentration on calibration samples is under-predicted, which is consistent with the idea that the model stain is more accumulated at the edges, or the swatch is less concentrated in middle.

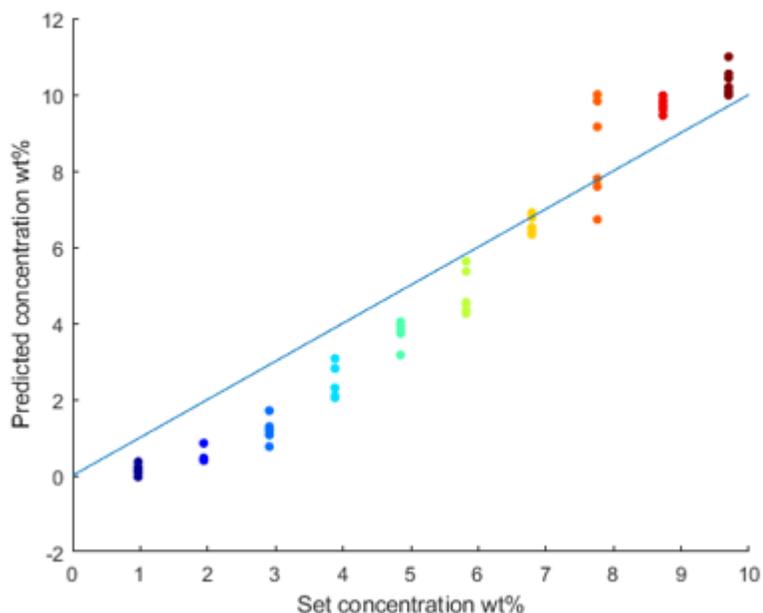


Figure 3.30 GPR model prediction on low-concentration set of polyester calibration samples, including $y = x$ line for reference.

Numerous strategies have been unsuccessfully attempted to mitigate the effect of this issue such as blocking the edges of the fabrics during drying, and using mixtures of solvents to create an opposing flow to the center.

A different approach to staining was therefore considered. If the model stain could be deposited onto the fabric in small droplets, individually applied to the fabric, then the solvent gradient effects should be irrelevant. Experimentally, this can be achieved using micro-dispensing, for which exist a number of commercial products. In this study, a piezo-based inkjet dispenser was constructed and implemented with an xy stage for a custom printing setup. The jetting head can dispense individual droplets of sunflower oil with mass as low as $7 \mu\text{g}$, while the moving stages underneath can dynamically reposition the sample to an accuracy of 0.1 mm. Further details about construction of the unit are available in section 3.10 and the setup is demonstrated in figure 3.31.

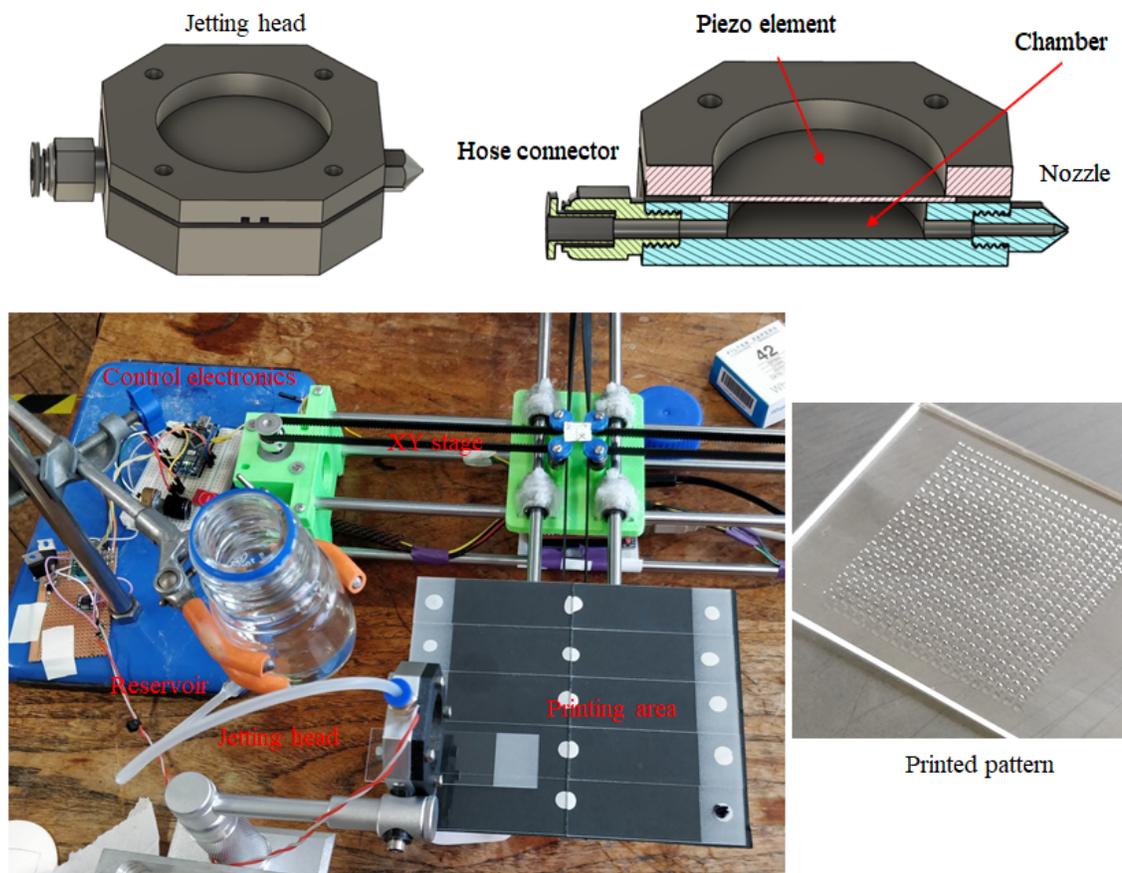


Figure 3.31 (Top) CAD drawing of custom made dispensing head. (Bottom) Custom printing setup used for dispensing droplets of oil, an example of a dispensed grid pattern is included.

The printing setup was tuned to dispense even grids of model stain, an example of such a grid deposited onto a glass slide is shown in figure 3.31. Fabric swatches were then placed to align with the printed pattern, and a grid of droplets dispensed evenly across its surface. The swatch was then weighed to determine the mass of lipid added, and staining procedure repeated until the desired stain concentration was achieved.

Using this method, a set of 13 calibration samples was made using swatches of cotton, polycotton, and polyester at a stain concentration of 0-10 wt %. The swatches were measured using DRS at three locations: in the center, in the corner, and midway between the two. This was done in an attempt to identify uneven distributions in the fabric. The spectra were analysed using the three component GPR model, and

compared to mass difference recorded during printing. The results on polyester and polycotton can be seen in figure 3.32. In general, a good correlation is observed, however, there are significant deviations between sample, as well as between reading on different locations of a single sample.

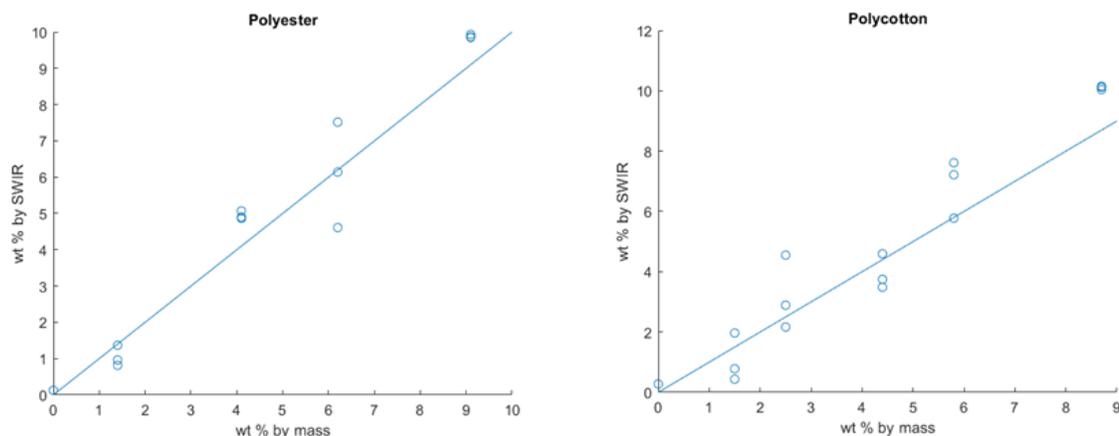


Figure 3.32 Lipid concentration of polyester and polycotton inkjet printed calibration samples, as predicted by the developed GPR model and compared to concentration measured by mass after deposition.

In the inkjet printed cotton calibration samples, the lipid concentration predicted using the new technique returned unusually high values compared to the amount dispensed by mass (figure 3.33-left). It was hypothesised that the model stain remained close to the surface after printing, contributing more to the predicted concentration. The fabric samples were heated for a few minutes using a hot-air tool and left overnight to homogenise, the measurements were then repeated (figure 3.33-right). This time the predictions are closer to the ideal $y = x$ fit, and support the theory of the lipid being confined to the surface. Even after this homogenisation period, however, the GPR model still over-predicted the lipid concentration compared to the set mass, which needs to be further investigated. This effect is only seen in the knitted cotton sample likely as a result of the fabric thickness, which was by far greatest in knitted cotton fabric.

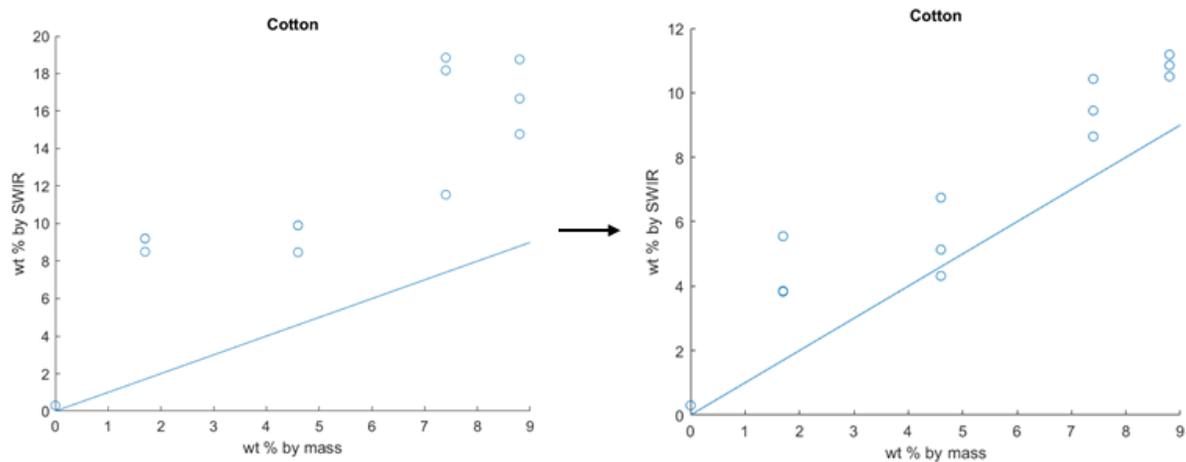


Figure 3.33 Lipid concentration of cotton inkjet printed calibration samples as predicted by the developed GPR model, and compared to concentration measured by mass after deposition. The samples on the right have been heated and left to homogenise.

3.9 Summary and conclusion

There has been identified a need for a technique capable of measuring the concentration of lipid present on a fabric. A unit system for such a metric was first defined to yield a wt % scale based on the lipid proportion by mass. A procedure for making calibration samples with a known and consistent lipid concentration was then devised, and subsequently used to make two calibration sample sets with a range of set lipid concentrations. One set consisted purely of stained cotton swatches, and another also containing samples of polyester and polycotton.

After consideration to preprocessing of reflectance spectra for machine learning, these sample sets were used to firstly train a PCR model in predicting the lipid concentration on the cotton fabric only set. It was found that a PCR model could be reliably fitted to the cotton set with an RMSE of 2.3 wt%.

Modelling was then expanded to the multi-fabric calibration sample set, where PCR models were found unsuitable. A GPR model could be fitted with an RMSE of prediction calculated to be 1.8 wt% from cross-validation. The model was further tested on spectra that were not used for training. The results showed a marginally larger but consistent RMSE of around 2 wt%. The LOD and LOQ of this model were determined to be 0.6 and 1.9 wt % respectively. The performance of the GPR model was also found stable when samples of lipid concentration beyond that shown to the model in training were tested, no erratic behaviours were seen in the model output which shows that the model has a good grounding in physical principals.

The impact of the vertical distribution of lipid inside the fabric was investigated through a model system made of thin tissue paper to simulate fabric. This system showed that the depth of the lipid soil inside the fabric matrix scales non-linearly with contribution to predicted lipid concentration. This is not directly comparable to textiles with larger weave size compared to tissue fibers, Although it gives an approximate value for the depth probed by the SWIR DRS method, around 0.5 mm.

Previously, measuring the level of lipid on fabric has been achieved by adding a reporter molecule to the model stain prior to simulated staining. The validity of this method was studied using the calibration set of samples which had a trace amount of dye included in the model stain. The results concluded that the ΔE^* metric, which is sometimes used to measure the relative performance of laundry formulations, scales non-linearly with the concentration of lipid stained onto fabric. This property makes it unsuitable in wash tests as a single ΔE^* value can originate from vastly different levels of lipid removal. Furthermore, techniques which rely on reporter chemicals are fundamentally limited in their applications, as real-world stains cannot always be labelled.

The constraints of single-point scanning were also overcome via automated scanning over the sample surface, and integration of the resulting concentration map was utilised to calculate the mass of stain in a scanned fabric swatch. This is the first example known to the author of an absolute measurement of lipid concentration on fabric.

The main limitation of the SWIR method is the issues in making low concentration samples of uniform lipid concentration across the surface. These were identified to originate from a coffee ring-like effect which occurs using the dilution method of sample preparation. In an attempt to resolve this problem, a custom inkjet printing setup was employed to dispense a uniform grid of small droplets evenly across the surface of the calibration swatches. Measurement of the subsequent samples showed promising results with thinner fabrics showing a good agreement between the SWIR method and gravimetry. However, the deviations were significant and on cotton samples, the initial measurements resulted in a higher than expected lipid concentrations. The discrepancy is thought to originate from the poor penetration of the model stain into the fabric, but inkjet printing still remains the most promising solution to making calibration samples.

Overall, the developed technique based on SWIR DRS is shown to be well suited to the task of studying lipid residues on fabric. The technique has clear advantages compared

to the methods using reporter chemicals, mainly that unlabelled stains can be measured.

However, before the technique can be meaningfully applied to assessment of laundry detergents, the fitted model must be verified for accuracy at low concentrations.

3.10 Experimental

3.10.1 Sample preparation

The fabrics used in this study were provided by Proctor & Gamble and their identity verified by spectroscopy before use. The sunflower oil used as model stain was purchased from a local supermarket and again verified by comparison of SWIR spectra to literature.

3.10.2 Reflectance spectroscopy

Reflectance Spectra were recorded using a custom fiber-optic reflectance spectroscopy system which was primarily designed for the measurement of ancient manuscripts using low optical power density. The system has previously been characterised with the results available elsewhere³ and a brief summary will be given here. The core of the system consists of an ARCOptix FTIR OEM module fitted with a FTIR-OEM200 light source and a FTIR-OEM000 interferometer. The output of the interferometer is coupled via a bifurcated fiber to a remote reflectance probe which is in-turn free space coupled to a detector (ARCOptix FTIR-OEM100). The probe consists of two kinematically-mounted fibre optic collimator focused to a single spot 50 mm below the probe. A 90° off-axis parabolic mirror then collects and directs the diffuse reflections into the detector.

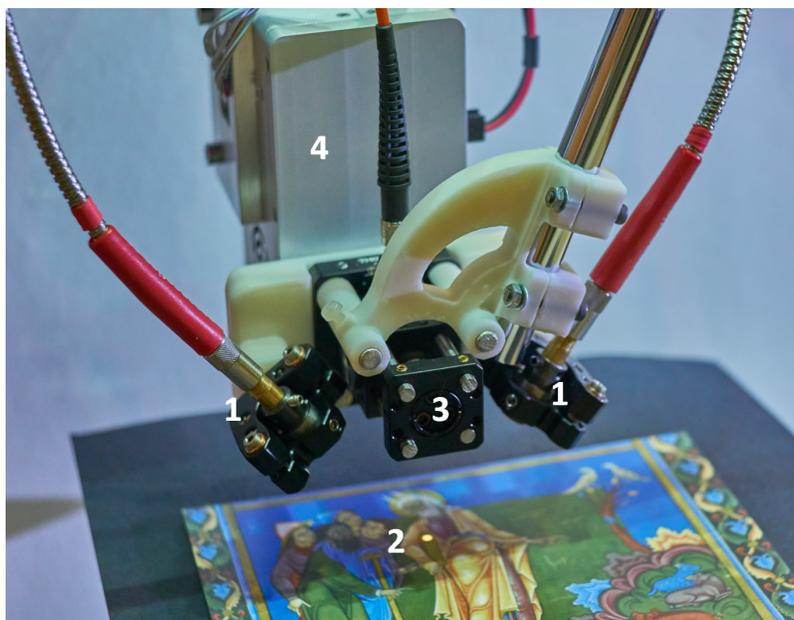


Figure 3.34 Setup of custom reflectance probe. (1) Probe optics (2) Probe spot (3) collection optics (4) detector.

Reflectance data was acquired using the OEM ‘ARCOptix Rocket’ software provided with the instrument. The detector was set to medium gain, with a spectral resolution of 8 cm^{-1} and three spectra were averaged for each measurement. The spectra were then exported to MATLAB for processing.

3.10.3 Imaging

Automated point-scanning was facilitated by the inclusion of a custom xy stage. The design was adapted from ref. 148 and features a dual stepper motor belt drive, controlled by an Arduino microcontroller. A computer side interface was developed using MATLAB app designer,¹¹⁸ which simultaneously connects to the ARCOptix spectrometer via the OEM API as well as the xy stage. The vertical focus was controlled manually as the algorithm repositions the sample and records a DRS spectrum.

3.10.4 Inkjet printing

The custom jetting head used for the production of uniform calibration samples had its design based on a literature device.¹⁴⁹ The device was constructed from a CNC machined aluminium frame, to which a common audio pick-up piezo crystal (purchased from RS components) was attached via compression from a 3D-printed bracket. The crystal was driven using a H-bridge MOSFET configuration, controlled by an Arduino microcontroller. The commonly available 3D printer nozzles were used in this implementation. The size of droplets could be primarily controlled by: pulse duration, pulse amplitude, and nozzle size. These parameters were tuned at the beginning of each printing batch until consistent drop ejection was observed. The smallest droplet mass that could be achieved with a 0.2 mm nozzle was 7 μg . For printing grid patterns, the same xy stage as used for imaging studies was employed here.

Chapter 4

Conclusions and future work

Throughout this work, it has been shown that modelling and machine learning, in combination with reflectance spectroscopy form a powerful tool for extracting information from the physical world.

In the first section of this work, we set out to understand the colour change observed when the illumination of certain lanthanide salts (mainly neodymium and holmium) is exchanged between sunlight and fluorescent lighting. The previous attribution of this phenomenon to the ‘Alexandrite effect’ was found inconsistent with experiments and the origin was concluded to be illuminant metamerism. Apart from nomenclature, these findings show that more consideration should be given to quantification of phenomena when reporting inconsistencies in visual appearance. Moreover, the study of such extreme phenomena lead to the development of tunable illuminant metamerism.

Tuning the spectral profile of light reflected from neodymium chloride, exploits the metamerism to provide an entirely novel method of tunable colouration. Granted the applications of such an approach are limited, the very nature of its operations serves to remind us that colour must be considered as an object-illuminant pair, as

opposed to an intrinsic property of the object. The remaining challenges of balancing spectrally tunable light sources require further work in order to reduce the off-white hue observed in some illumination profiles. The solutions will likely require a stable and well-calibrated spectral feed-back in addition to accurate and well-resolved spectral control of the illuminant. Neodymium chloride was the only sample considered in this study for tunable illuminant metamerism, and it was found adequate to fully control the hue of the appearance, however, the maximum possible saturation was limited. Other samples may very well be more suited for this task and allow a greater range of colour representation. The samples of holmium acetate hydrate and erbium acetate hydrate measured in this study both satisfy the criteria for use in tunable illuminant metamerism, although this would require a suitable spectrally tunable light source to be designed.

Given the differences between individual human observers, the feasibility of indistinguishable metamers of ‘white’ light that can control the colour appearance of a specific object remains to be seen.

The CIE colourimetry system was extensively used in the study of extreme colour changes with a good level of accuracy. The colour appearance models were utilised to simulate the impact that oxidative degradation of fabric has on its colour appearance. The resulting simulations were used to effectively analyse the workings of FWAs and HDs from a spectroscopy perspective. This model was derived from a single set of samples and still requires calibration with experimentation. The platform herein described provides a versatile method of analysis of fabric appearance from spectroscopy data. If combined with a large enough dataset on real world: fabric reflectance spectra, illuminant spectra, and the spectral properties of laundry formulations could be collected, then it could be analysed using this platform to see the distribution of possible appearance scenarios. Such an analysis would be of great use in development of laundry detergent formulations, where the most effective formulations can be selected as the ones which are effective under a larger majority of scenarios. Similarly, formulations

could also be tailored to specific illumination condition, like the outdoors for example, and the correct additive selected through simulation.

The model itself may also be improved by using a less approximated form of the KM model, similar to the work of Stavenga and van der Kooi¹⁵⁰ who approximate flower leaves using a stacked layer model in order to study the vertical distribution of the chromophores. Such a model would need to be realised with each layer having its own wavelength-dependant attenuation properties. Future fabric appearance models may even feature layer-wise simulation of properties such as FWA excitation integral and luminescence emission, allowing for better prediction of FWA's behaviour. Such models would be computationally more intensive, but with the rapid developments seen in silicon integrated circuits, computation of larger and larger models is becoming accessible to a wider range of devices.

Following the colourimetry work, the attention of the study then turned to developing a method for quantifying the concentration of lipid in laundry stains. This is a mature topic in laundry literature, yet common practices still involve flawed assumptions. There is also no common unit system for comparing the level of lipid staining between studies, and this was the focus of initial work in this field.

A unit system called wt % was devised based on the ratio of masses between the lipid and the fabric it occupies. This quantity was then represented in physical samples which have a set and near-uniform lipid concentration across their surface. These samples were termed calibration samples, and were used in combination with DRS and GPR to form a prediction model. The accuracy of any model fitted to the calibration sample set is ultimately limited by the accuracy to which the set stain concentration is known to. The dilution method devised to make calibration samples was found to work well above stain concentration of 10 wt %, however, below this limit a coffee ring-like effect leads to an uneven distribution of model stain across the calibration sample. A solution to this model was suggested in the form of inkjet printing, which was shown to be a promising approach but further work is required to understand the lipid deposition process.

The developed model can estimate the lipid concentration on fabric by multi-variate analysis of the SWIR DRS spectrum. A GPR model fitted to a calibration set of cotton, polycotton and polyester was able to achieve an RMSE of prediction of around 2 wt%. Measurement of blank samples showed that the model has a LOD of 0.6 wt% and LOQ of 1.9 wt%. These values still need to be validated to ensure that the model is linear at lower concentrations.

The addition of moving stages to the DRS setup allowed for the technique to be applied in imaging, where a low-resolution concentration map could be measured across the surface of a fabric. This application is a powerful addition to the method, and allow for the variation of lipid across a fabric to be mapped, as well as the absolute mass calculated to a good degree of precision. The work can be further improved if applied with hyper-spectral cameras. Such cameras are now commercially available¹⁵¹ and capable of vastly greater spacial resolution than the point scanning method used in this study. Hyper spectral imaging systems, however, typically have lower sensitivity and spectral resolution compared to the FT spectrometer used in this study. The application of lipid prediction from hyper-spectral imaging systems therefore provides its own set of challenges for future scientists.

At the time of writing, further work in validation of SWIR DRS lipid prediction models has been carried out at Proctor & Gamble. Inkjet printing has been successfully used to validate such models across the full analyte concentration range, and the technique submitted for a patent application.

This work has therefore shown that DRS in the SWIR region is well suited to the study of lipid components of laundry stains. The technique works well on the model system created, and while its application to real-world stains remains strongly feasible, there are external factors which need to be explored. There are many textiles available other than the three explored in this work, and for them to be compatible with this method

they must not absorb strongly in the regions where lipid bands are most prevalent. The moisture content of the fabrics is also likely to vary with time, and the impact of this on the developed model remains to be investigated.

The technique is broadly sensitive to lipids due to the vast amount of C-H moieties they contain, this limits the scope of this approach in identifying the lipids present. However, it may even be beneficial in simply quantifying the amount present, which is important for testing laundry detergent performance.

In combination, the sections of work described in chapters 2 and 3 contribute to furthering the development of laundry technology. Given the rapid developments seen in the fields of optics and electronics, as well as the constant wide-spread demand for laundry, it is speculated that technology is going to become an increasing part of the laundry process.¹⁵² This may be in the form of washing machines with integrated garment analysis and tailored dispensing of formulation components, or colour correction products made for specific garments types or environments. The work presented in this thesis provides firm grounding for future scientists developing such technologies.

Chapter 5

Appendices

5.1 Supplementary information

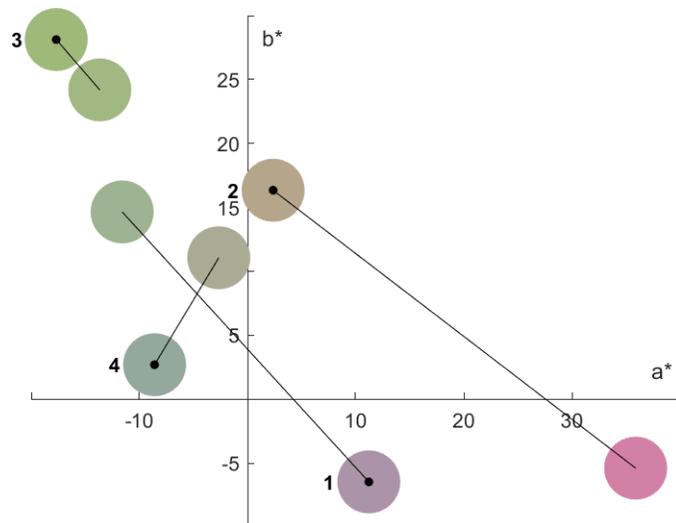


Figure 5.1 L*a*b* coordinates of diffusely reflected light from samples under sunlight (central black spot) and cool white fluorescent lights. (1) neodymium chloride (2) holmium sulphate (3) praseodymium chloride (4) Alexandrite.

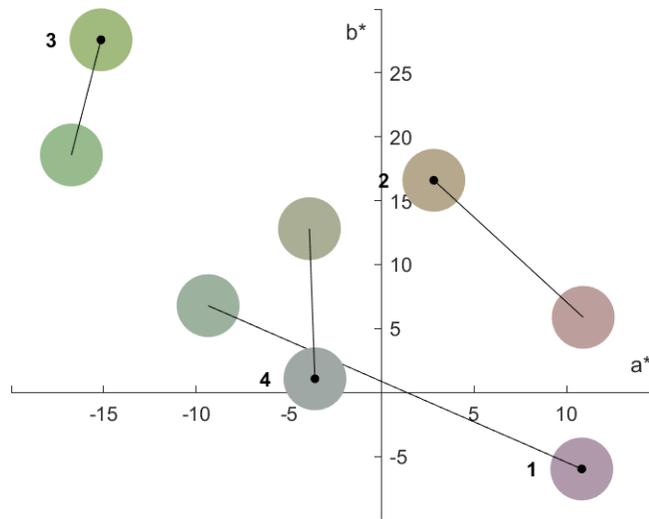
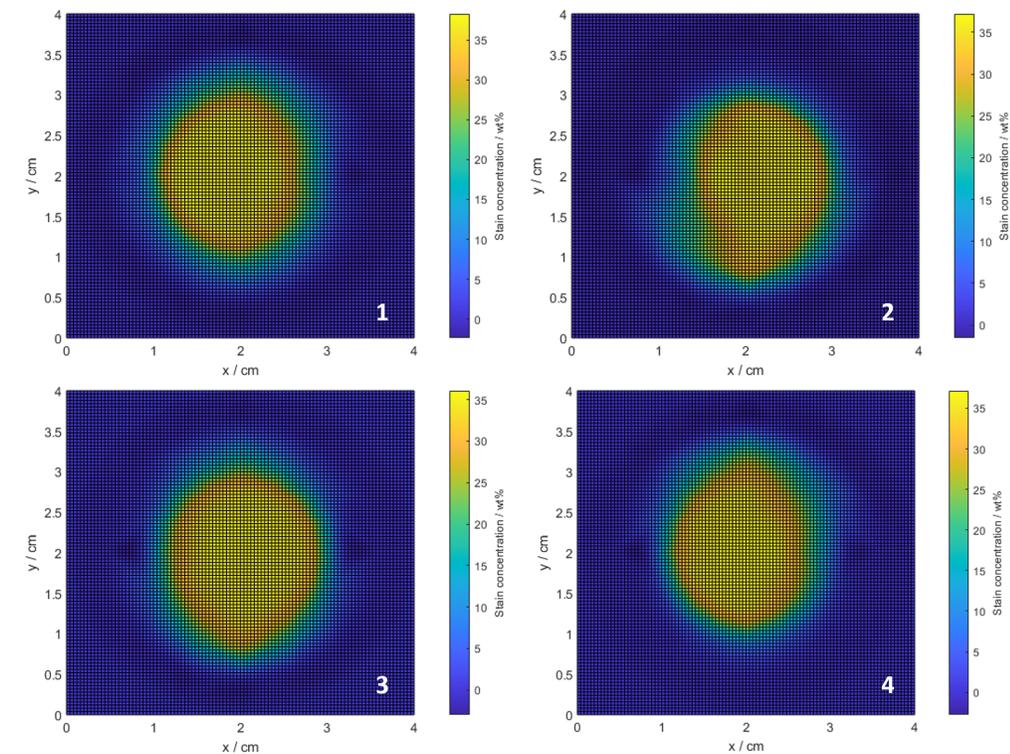


Figure 5.2 $L^*a^*b^*$ coordinates of diffusely reflected light from samples under 4310 K black body radiator (central black spot) and warm white fluorescent lights. (1) neodymium chloride (2) holmium sulphate (3) praseodymium chloride (4) Alexandrite.

Figure 5.3 Stain concentration maps of four stained knitted cotton swatches imaged by cross scanning in combination with interpolation. The sample numbers shown in the lower right corner of each plot correlates to the sample index in table 3.6.



5.2 Electronic supplementary information

Further information is provided in the ‘Supplementary Information’ folder.

References

1. R. W. G. Hunt and M. R. Pointer, *Measuring colour*, Wiley, Chichester, 1987.
2. A. Scheeline, How to Design a Spectrometer, *Appl. Spectrosc.*, 2017, **71**, 2237–2252.
3. H. Hayler, Visible-NIR-SWIR Fibre Optic Reflectance Spectroscopy in the Analysis of Colourants on Illuminated Manuscripts, PhD thesis, University of Durham, 2018.
4. M. Eismann, *Hyperspectral Remote Sensing*, SPIE, 2012.
5. J. Mooney and P. Kambhampati, Get the Basics Right: Jacobian Conversion of Wavelength and Energy Scales for Quantitative Analysis of Emission Spectra, *J. Phys. Chem. Lett.*, 2013, **4**, 3316–3318.
6. J. Geist, W. K. Gladden and E. F. Zalewski, Physics of photon-flux measurements with silicon photodiodes, *J. Opt. Soc. Am.*, 1982, **72**, 1068.
7. P. Wurfel, S. Finkbeiner and E. Daub, Generalized Planck’s radiation law for luminescence via indirect transitions, *Appl. Phys. A*, 1995, **60**, 67–70.
8. R. Basri and D. Jacobs, Lambertian reflectance and linear subspaces, *IEEE*, 2003, **25**, 218–233.
9. Spectralon Diffuse Reflectance Material, <https://www.labsphere.com/labsphere-products-solutions/materials-coatings-2/coatings-materials/spectralon>, (accessed April 2022).
10. K. Yamamoto, A. Masui and H. Ishida, Phosphor-Based InGaN/GaN White Light-Emitting Diodes With Monolithically Integrated Photodetectors, *IEEE*, 2021, **33**, 6285–6293.
11. A. Kelly, S. Halsey, R. Bottom, S. Korde, T. Gough and A. Paradkar, A novel transreflectance near infrared spectroscopy technique for monitoring hot melt extrusion, *Int. J. Pharma.*, 2015, **496**, 117–123.

12. S. A. Halsey, The use of transmission and transreflectance nearinfrared spectrscopy for the analysis of beer, *J. Inst. Brew.*, 1985, **91**, 306–312
13. P. Thongnopkun and S. Ekgasit, FTIR Spectra of faceted diamonds and diamond simulants, *Diamond and Related Materials*, 2005, **14**, 1592–1599.
14. D. G. Myers, *Psychology*, Worth Publishers, New York, 2008.
15. T. Young, The Bakerian Lecture: On the Theory of Light and Colours, The Bakerian Lecture, 1802, **92**, 12–48.
16. A. Walz-Flannigan, S. Stekel, H. Weber, D. Lanners, R. Jonsgaard, T. Peterson and S. G. Langer, Aging and Quality Control of Color LCDs for Radiologic Imaging, *J Digit Imaging*, 2011, **24**, 828–832.
17. P. Boher, T. Leroux, V. Collomb Patton and T. Bignon, Color accuracy of imaging using color filters, *J. Inf. Disp.*, 2011, **13**, 7–16.
18. J. Schanda, *Colorimetry: Understanding the CIE System*, Wiley, New york, 2007.
19. H. J. A. Dartnall, J. K. Bowmaker and J. D. Mollon, Human Visual Pigments: Microspectrophotometric Results from the Eyes of Seven Persons, *Proc. Royal Soc. B 220*, 1983, **220**, 115–130.
20. K. Nassau, *The Physics and Chemistry of Color: The Fifteen Causes of Color*, Wiley, New york, 1983.
21. S. S. Patterson, M. Neitz and J. Neitz, Reconciling Color Vision Models With Midget Ganglion Cell Receptive Fields, *Front. Neurosci.*, 2019, **0**, 1.
22. G. D. Hastings and A. Rubin, Colour spaces - a review of historic and modern colour models, *Afr. vis. eye health*, 2012 **71**, 133–143.
23. Development of the Munsell Color Order System, <https://munsell.com/about-munsell-color/development-of-the-munsell-color-order-system/>, (Accessed April 2022).
24. M. Jossierand, E. Meeussen, A. Majid and D. Dediu, Environment and culture shape both the colour lexicon and the genetics of colour perception, *Sci Rep*, 2021, **11**, 2045-2322.
25. D. H. Foster and A. Reeves, Colour constancy failures expected in colourful environments, *Proc. R. Soc. B.*, 2022, **289**, 2021-2037.
26. A. D. Logvinenko, B. Funt, H. Mirzaei and R. Tokunaga, Rethinking Colour Constancy, *PLoS One*, 2015, **10**, 1932-6203.
27. D. Weiss, C. Witzel and K. Gegenfurtner, Determinants of Colour Constancy and the Blue Bias, *IPE*, **8**, 2017.

28. Don't Overwash: It's Time to Change the Way We Care, <https://www.fashion-revolution.org/dont-overwash-its-time-to-change-the-way-we-care/>, (Accessed April 2022).
29. Environmental Impacts of the Fashion Industry, <https://www.sustainyourstyle.org/old-environmental-impacts>, (Accessed April 2022).
30. Global Greenhouse Gas Emissions Data, <https://www.epa.gov/ghgemissions/global-greenhouse-gas-emissions-data>, (Accessed April 2022).
31. J. S. Golden, V. Subramanian, G. M. A. U. Irizarri, P. White and F. Meier, Energy and carbon impact from residential laundry in the United States, *J. Integr. Environ. Sci.*, 2010, **7**, 53–73.
32. The Hidden Impact of Your Daily Water Use, <https://www.bbc.com/future/article/20200326-the-hidden-impact-of-your-daily-water-use>. (Accessed April 2022).
33. R. M. Mohamed, A. A. Al-Gheethi, J. Noramira, C. M. Chan, M. K. A. Hashim and M. Sabariah, Effect of detergents from laundry greywater on soil properties: a preliminary study, *Appl Water Sci.*, **8**, 16.
34. R. M. S. Radin Mohamed, A. Ahmed, C. M. Chan and A. H. Kassim, The Use of Natural Filter Media Added with Peat Soil for Household Greywater Treatment, *JET*, 2014, **2**, 33–38.
35. F. De Falco, E. Di Pace, M. Cocca and M. Avella, The contribution of washing processes of synthetic clothes to microplastic pollution, *Sci. Rep.*, 2019, **9**, 6633.
36. Plastic Fibres Found in Tap Water around the World, Study Reveals, <http://www.theguardian.com/environment/2017/sep/06/plastic-fibres-found-tap-water-around-world-study-reveals>. (Accessed April 2022).
37. N. Laskar and U. Kumar, Plastics and microplastics: A threat to environment, *Environ. Technol. Innov.*, 2019, **14**, 100352.
38. Y.S. Chi and S. K. Obendorf, *J Surfact Deterg*, 1998, **1**, 523–527.
39. K. Ahn, S. Zaccaron, N. S. Zwirchmayr, H. Hettegger, A. Hofinger, M. Bacher, U. Henniges, T. Hosoya, A. Potthast and T. Rosenau, Yellowing and brightness reversion of celluloses: CO or COOH, who is the culprit?, *Cellulose*, 2019, **26**, 429–444.
40. W. A. Reeves, T. A. Summers and R. M. Reinhardt, *Text. Res. J.*, **50**, 711–717.
41. M. S. Sk, R. Mia, B. Ahmed, A. Rahman and M. M. Rahman Palash, *Heliyon*, **7**, 2021.
42. G. Manasoglu, M. Kanik and K. Yildirim, *J. Eng. Fibers Fabr.*, **14**, 2019.

43. Ullmann's Encyclopedia of Industrial Chemistry, Wiley, New York, 2000.
44. A. Saeed, G. Shabir and I. Batool, *J. Fluoresc.*, 2014, **24**, 1119–1127.
45. J. B. Hall, J. J. Hefferren and N. H. Olsen, *J. Dent Res.*, 1970, **49**, 1431–1436.
46. D. Benjamin, Complex fluids for active delivery in laundry application Dye – Surfactant interactions, PhD thesis, University of Durham, 2018.
47. Statistics and Machine Learning Toolbox, <https://uk.mathworks.com/products/statistics.html>, (Accessed April 2022).
48. A Tutorial on Principal Component Analysis, <https://arxiv.org/abs/1404.1100v1>, (Accessed April 2022).
49. T-SNE - MATLAB & Simulink - MathWorks United Kingdom, <https://uk.mathworks.com/help/stats/t-sne.html>, (Accessed April 2022).
50. C. Cui and T. Fearn, *J. Near Infrared Spectrosc.*, 2017, **25**, 5–14.
51. C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*, MIT Press., Cambridge, 2006.
52. Exact GPR Method - MATLAB & Simulink - MathWorks United Kingdom, <https://uk.mathworks.com/help/stats/exact-gpr-method.html>, (Accessed April 2022).
53. A. Clarke and E. Waterton, *Landscape Research*, 2015, **40**, 971–992.
54. F. Xie, Y. Cao, C. Ranchon, A. Hart, R. Hansen, J. E. Post, C. W. Whitney, E. Dawson-Tarr, A. J. Drew and D. J. Dunstan, *Sci. Rep.*, **10**, 2020.
55. A. Tokunaga, L. M. Uriarte, K. Mutoh, E. Fron, J. Hofkens, M. Sliwa and J. Abe, *J. Am. Chem. Soc.*, 2019, **141**, 17744–17753.
56. Y. Inoue, P. Kuad, Y. Okumura, Y. Takashima, H. Yamaguchi and A. Harada, *J. Am. Chem. Soc.*, 2007, **129**, 6396–6397.
57. B. Y. Moon, S. Y. Kim and D. S. Yu, *PLoS one*, **15**, 2020.
58. S. Bardaa, N. Ben Halima, F. Aloui, R. Ben Mansour, H. Jabeur, M. Bouaziz and Z. Sahnoun, *Lipids Health Dis*, 2016, **15**, 73.
59. S. Kreft and M. Kreft, *Sci. Nat. OCLC no.*, 2007, **94**, 935–939.
60. G. O. Fruhwirth and A. Hermetter, *Eur. J. Lipid Sci. Technol.*, 2007, **109**, 1128–1140.
61. M. N. Taran and I. V. Naumenko, *Mineral. Mag.*, 2016, **80**, 705–717.

62. R. C. Powell, L. Xi, X. Gang, G. J. Quarles and J. C. Walling, *Phys. Rev. B.*, 1985, **32**, 2788–2797.
63. M. O. Hunault, Y. Harada, J. Miyawaki, J. Wang, A. Meijerink, F. M. F. de Groot and M. M. van Schooneveld, *J. Phys. Chem. A*, 2018, **122**, 4399–4413.
64. Y. Liu, J. Shigley, E. Fritsch and S. Hemphill, *Min. Mag.*, 1995, **19**, 186–191.
65. C. Poole, *Phys. Chem. Solids*, 1964, **25**, 1169–1182.
66. W. A. White, R. Rustum and J. Crichton, *Am. Min.*, 1967, **52**, 867–871.
67. H. Carstens, *Contrib. to Mineral. Petrol.*, 1973, **41**, 273–276.
68. I. Kasajima, *Sci. Rep.*, 6, 2016.
69. E. Gubelin and K. Schmetzer, *Gems. Gemol.*, 1982, **18**, 197–203.
70. Z. Sun, A. C. Palke and N. Renfro, *I*, 2016, **51**, 348–369.
71. D. Zhou, T. Lu, R. Sun, G. Shi, H. Chen and J. Ke, *Color Res Appl.*, 2019, **44**, 925–931.
72. Y. Liu, J. E. Shigley, E. Fritsch and S. Hemphill, *Min. Mag*, 1995, **59**, 111–114.
73. Y. Liu, Q. Chen, X. Bu and P. Feng, SPIE proceedings, 2002.
74. Lab Equipment and Lab Supplies — Fisher Scientific, <https://www.fishersci.com/us/en/home.html>, (Accessed April 2022).
75. Sigma-Aldrich: Analytical, Biology, Chemistry & Materials Science Products and Services., <https://www.sigmaaldrich.com/sigma-aldrich/home.html>. (Accessed April 2022).
76. L. R. Bernstein, *Am. Min.*, 1982, **67**, 356–359.
77. J. Lin, K. Diefenbach, M. A. Silver, N. S. Dalal and T. E. Albrecht-Schmitt, *Cryst. Growth Des.*, 2015, **15**, 4606–4615.
78. J. Lin, K. Diefenbach, J. N. Cross, J.-M. Babo and T. E. Albrecht-Schmitt, *Inorg. Chem.*, 2013, **52**, 13278–13281.
79. J. Diwu, D. J. Grant, S. Wang, L. Gagliardi and T. E. Albrecht-Schmitt, *Inorg. Chem.*, 2012, **51**, 6906–6915.
80. E. M. Villa, S. Wang, E. V. Alekseev, W. Depmeier and T. E. Albrecht-Schmitt, *Eur. J. Inorg. Chem.*, 2011, **25**, 3749 - 3754.
81. S. Wang, E. V. Alekseev, W. Depmeier and T. E. Albrecht-Schmitt, *Indoor Built Environ*, 2021, **51**, 7–9.

82. QP-Card, <https://argraph.com/QPcard/QPcard.htm>, (Accessed April 2022).
83. Lighting in Stores: What Is the Ideal Colour Temperature for Displaying Meat, Fish, Vegetables, Bread, <https://www.dmlights.com/help/lighting/>, (Accessed April 2022).
84. J. V. Cooper, B. R. Wiegand, A. B. Koc, L. Schumacher, I. Grun and C. L. Lorenzen, *Anim. Sci. J.*, 2016, 94, 4457–4462.
85. J. Yin, X. An, L. Chen, J. Li, J. Wu, Y. Luo, Q. Wang, H. Yu and K. H. Li, *IEEE Trans Electron Devices*, 2021, **68**, 132–137.
86. M. Meneghini, L.-R. Trevisanello, G. Meneghesso and E. Zanoni, *IEEE Trans. Device Mater. Reliab.*, 2008, **8**, 323–331.
87. K. W. Houser, M. Wei, A. David and M. R. Krames, *LEUKOS*, 2014, **10**, 165–180.
88. R. Feltman, Under LED Lights, Your Clothes Can't Get "Whiter-than-White", <https://qz.com/200743/under-led-lights-your-clothes-cant-get-whiter-than-white/>, (Accessed 2022).
89. R. F. Quartz, Why New Lightbulbs Will Make Your White T-Shirt Look Bad, <https://www.theatlantic.com/technology/archive/2014/04/why-new-lightbulbs-will-make-white-t-shirts-look-bad/361198/>, (Accessed April 2022).
90. Reveal HD+, <https://www.gelighting.com/led-hd-bulbs/reveal>, (Accessed 2022).
91. M. W. Jackson, *Ann Phys*, 2014, **526**, 65–69.
92. S. Konigs, S. Mayr and A. Buchner, *Ergonomics*, 2019, **62**, 1462–1473.
93. S. Mitra and S. Jana, *J. Fluoresc.*, 2015, **25**, 541–549.
94. K. Burns, K. B. Adams and J. Longwell, *J Opt Soc Am*, 1950, **40**, 339.
95. C. S. McCamy, *Color Res Appl*, 1992, **17**, 142–144.
96. D. G. Karraker, *Inorg. Chem.*, 1968, 7, 473–479.
97. G. W. Parker and P. M. Lantz, *J. Am. Chem. Soc.*, 1950, 72, 2834–2836.
98. G. R. Dillip, C. M. Reddy, M. Rajesh, S. Chaurasia, B. D. P. Raju and S. W. Joo, *Bull Mater Sci*, 2016, **39**, 711–717.
99. A. Ansari, H. A. Hussain and K. Iftikhar, *Spectrochimica Acta Part A: Molecular and Biomolecular Spectroscopy*, 2007, **68**, 1305–1312.
100. H. Mark and J. Workman, *Chemometrics in Spectroscopy*, Academic Press, 2018.

101. F. Reichert, F. Moglia, P. W. Metz, A. Arcangeli, D.-T. Marzahl, S. Veronesi, D. Parisi, M. Fechner, M. Tonelli and G. Huber, *Opt. Mater. Express*, 2015, **5**, 88.
102. P. Villwock, S. Siol and T. Walther, *Eur. Phys. J.*, 2011, **65**, 251–255.
103. S. W. Brown, J. P. Rice, J. E. Neira, B. C. Johnson and J. D. Jackson, *J. Res. Natl.*, 2006, **111**, 401.
104. V. Viliunas, H. Vaitkevicius, R. Stanikunas, A. Svezgda and Z. Bliznikas, *Light. Res. Technol.*, 2011, **43**, 321–330.
105. A. Kelber, C. Yovanovich and P. Olsson, *Philos. Trans. R. Soc.*, 2017, **372**, 2016065.
106. A. Brown, M. Corner, D. P. Crewther and S. G. Crewther, *Front. Hum. Neurosci.*, **12**, 2018.
107. H.-Y. Chou and T.-H. Yang, *J. Light Vis. Environ.*, 2008, **32**, 183–186.
108. G. Finlayson, M. Mackiewicz, A. Hurlbert, B. Pearce and S. Crichton, *J. Opt. Soc. Am. A*, 2014, **31**, 1577–1587.
109. J. Li, P. Hanselaer and K. A. G. Smet, *LEUKOS*, 2021, 1–23.
110. A. Petruilis, P. Vitta, J. Aglinskaite, R. Vaicekauskas and A. Zukauskas, *Proc. Latv. Acad. Sci. B: Nat.*, 2017, **71**, 366–371.
111. Fundamental Chromaticity Diagram with Physiological Axes. Pt. 1, ed. I. CIE Central Bureau, Vienna, 2006.
112. M. G. Nagle and D. Osorio, *Proc. Royal Soc. B*, 1993, **252**, 209–213.
113. G. He, J. Xu and H. Yan, *AIP Advances*, **1**, 2011.
114. A. M. Villalba, J. M. Monteoliva and A. E. Pattini, *Indoor Built Environ*, 2021, **30**, 970–984.,
115. L. G. Coppel, M. Andersson and P. Edstrom, *Appl. Opt.*, 2011, **50**, 2784.
116. A. Arecchi, T. Messadi and R. J. Koschel, *Field Guide to Illumination*, SPIE, 2007.
117. A. Beeby, L. Garner, D. Howell and C. E. Nicholson, *J. Inst. Conserv.*, 2018, **41**, 142–153.
118. MATLAB App Designer, <https://uk.mathworks.com/products/matlab/app-designer.html>, (Accessed April 2022).
119. Future Eden Supplier of Premium LED Components, Drivers and Heatsinks – Future Eden Ltd, <https://futureeden.co.uk/>, (Accessed April 2022).

120. O. W. Neiditch, K. L. Mills and G. Gladstone, *Bull Mater Sci*, 2016, **57**, 426–429.,
121. Y. Yu, M. Cai, L. Zhou and J. Shao, *Fibers Polym*, 2016, **17**, 305–309.
122. A. James, J. Casey, D. Hyliands and G. Mycock, *World J Microbiol Biotechnol*, 2004, **20**, 787–793.
123. M. Chauhan, R. S. Chauhan and V. K. Garlapati, *BioMed Research International*, 2013, 1–6.
124. D. Boskou, *Olive Oil: Chemistry and Technology*, AOCS Publishing, 2006.
125. R. Garces, J. J. Salas-Linan, M. Venegas-Caleron and E. Martinez-Force, *Whiley*, 2009, **21**, 79-82.
126. N. A. A. Ghani, A.-A. Channip, P. Chok Hwee Hwa, F. Jaafar, H. M. Yasin and A. Usman, *Food Sci. Nutr.*, 2018, **6**, 1298–1306.
127. M. Azir, S. Abbasiliasi, T. Tengku Ibrahim, Y. Manaf, A. Sazili and S. Mustafa, *foods*, 2018, **6**, 98.
128. E. Makrantonaki, R. Ganceviciene and C. C. Zouboulis, *Dermato-Endocrinology*, 2011, **3**, 41–49.
129. A. Pappas, S. Johnsen, J.-C. Liu and M. Eisinger, *Dermato-Endocrinology*, 2009, **1**, 157–161.
130. R. M. B. Mackenna, V. R. Wheatley and A. Wormall, *J. Invest. Dermatol.*, 1950, **15**, 33–47.
131. M. Picardo, M. Ottaviani, E. Camera and A. Mastrofrancesco, *Dermato- Endocrinology*, 2009, **1**, 68–71.
132. E. Jurado, Bravo, Bailon, J. Nunez Olea and Altmajer, *Ind. Eng. Chem. Res.*, 2003, **42**, 4303–4310.
133. E. Jurado, V. Bravo, G. Luzon, M. Fernandez-Serrano, M. GarcRoman, D. Altmajer-Vaz and J. M. Vicaria, *J Surfactants Deterg*, 2007, **10**, 61–70.
134. E. Jurado, M. Garcia-Roman, G. Luzon, D. Altmajer-Vaz and J. L. Jimenez-Perez, *Ind. Eng. Chem. Res.*, 2011, **50**, 11502–11510.
135. E. Jurado-Alameda, M. Garcia-Roman, D. Altmajer-Vaz and J. L. Jimenez-Perez, *J. Food Eng.*, 2012, **110**, 44–52.
136. I. Hughes, *Measurements and Their Uncertainties: A Practical Guide to Modern Error Analysis*, Oxford University Press, New York, 2010.
137. A. Garrido-Varo, M.-T. Sanchez, M.J. De la Haba, I. Torres and D. Perez-Marin, *Sensors*, 2017, **17**, 2642.

138. S. P. Mishra, A Text Book of Fibre Science and Technology, New Age International, New Delhi, 2000.
139. A. Basch, T. Wasserman and M. Lewin, *J Polym Sci A Polym Chem*, 1974, **12**, 1143–1150.
140. A. D. de Oliveira, V. H. da Silva, M. F. Pimentel, G. M. Vinhas, C. Pasquini and Y. M. B. de Almeida, *Mat. Res.*, **21**, 2018.
141. J. Gerretzen, E. Szymanska, J. J. Jansen, J. Bart, H.J. van Manen, E. R. van den Heuvel and L. M. C. Buydens, *Anal. Chem.*, 2015, **87**, 12096–12103.
142. J. H. Nobbs, *Wiley*, 2008, **15**, 66–75.
143. S. K. Loyalka and C. A. Riggs, *Appl. Spectrosc.*, 1995, **49**, 1107–1110.
144. K. H. Norris, *NIR news*, 2015, **16**, 10–13.
145. C. E. Miller and T. Naes, *Appl. Spectrosc.*, 1990, **44**, 895–898.
146. H. Martens, J. P. Nielsen and S. B. Engelsen, *Anal. Chem.*, 2003, **75**, 394–404.
147. D. A. Armbruster and T. Pry, *Clin Biochem Rev*, 2008, **29**, S49–S52.
148. T. Martinsen, 4xiDraw, <https://github.com/bullestock/4xidraw>, (Accessed April 2022).
149. D. M. Harris, T. Liu and J. W. M. Bush, *Exp. Fluids.*, **56**, 2015.
150. D. G. Stavenga and C. J. van der Kooi, *Planta*, 2016, **243**, 171–181.
151. SWIR-384 Hyperspectral Camera — 900 - 2500nm — ClydeHSI, <https://www.clydehsi.com/hyperspectral-cameras/swir-384-hyperspectral-camera>, (Accessed April 2022).
152. K. Abeliotis, C. Amberg, C. Candan, A. Ferri, M. Osset, J. Owens and R. Stammering, *Household and Personal Care Today*, 2015. **10**, 22.