

Capture and Simulation of Fluorescence in Fabrics

Estimating and simulating fluorescent appearance in fabrics using LEDs and an RGB camera.

Master's thesis in Systems, Control and Mechatronics

Peter McEvoy

MASTER'S THESIS 2020

Capture and Simulation of Fluorescence in Fabrics

Peter McEvoy



CHALMERS
UNIVERSITY OF TECHNOLOGY

Department of Electrical Engineering
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2020

Capture and Simulation of Fluorescence in Fabrics
Peter MCEVOY

© PETER McEVOY, 2020.

| | | |
|--------------|---------------------|--|
| Supervisors: | Elena Garcés | Universidad Rey Juan Carlos & Seddi Labs Inc. |
| | Marco Fratarcangeli | Department of Computer Science and Engineering, Chalmers University of Technology |
| Examiner: | Morten Fjeld | Department of Computer Science and Engineering, Chalmers University of Technology |

Master's Thesis 2020
Department of Electrical Engineering
Chalmers University of Technology
SE-412 96 Gothenburg
Telephone +46 31 772 1000

Cover: Three fluorescent fabrics under a consumer D65-like LED illumination. Beneath each image is the corresponding spectral re-radiance matrix that has been estimated using the method presented in this thesis. The upper-right half of each image shows a photograph of the actual fabric under D65-like LED illumination. The lower-left half of each image shows a rendered image of a Lambertian surface simulated using these re-radiance matrices to determine the fluorescent response to a light source with a similar illumination spectrum.

Typeset in L^AT_EX
Gothenburg, Sweden 2020

Capture and Simulation of Fluorescence in Fabrics
Peter McEvoy
Department of Electrical Engineering
Chalmers University of Technology

Abstract

In textiles, *fluorophores* (fluorescent molecules) are mixed with dyes to help achieve exceptionally vivid colours. Fluorophores help achieve such bright colours by absorbing light with a short wavelength (such as hard to see ultra-violet) and re-emit light with a longer wavelength (such as a blue, orange or red). The challenge when simulating fluorescence is that it places demands on the rendering-system (the system that simulates light interactions with the materials). An additional challenge is that methods to capture this fluorescent behaviour require specialised laboratory equipment. In this thesis, experiments are presented that aim to study if the multiple levels of structure in fabrics introduce any angular dependence of the fluorescence. The results suggest that in fluorescent yarns and fabrics there is no significant angular dependence. This result indicates that fluorescence can be captured from a flat sample from any distance, and that yarn-level resolution is not required in this regard. Using those conclusions a practical capture setup for fluorescence in fabrics is proposed. The method uses nine narrow-band LEDs, a consumer camera and an image analysis procedure that estimates the wavelengths of the detected light through numerical optimisation to give a *re-radiance* matrix. This matrix describes the fluorescent material's spectral response to wavelengths of incoming light. The results using this capture setup are compared against raw data measured with a spectrometer. The results show agreement. Finally, matrices acquired from the capture setup are used to generate computer generated images of the materials under different illuminations. These images are compared to photographs of the same scenes.

Keywords: computer graphics, spectral rendering, fluorescence, appearance capture, material modelling.

Acknowledgements

I had the opportunity to do my Master's Thesis Project at SEDDI Labs Inc. in Madrid, Spain. I would like to thank all of the brilliant people I met during that time for making me feel welcome and making my stay as good as it was. Thank you to David Pascual, Carlos Rodriguez and Sergio Suja for their help and support.

Thank you to Miguel A. Otaduy for first introducing me to this opportunity and Jorge Lopez Moreno for making it possible, supporting the project and for providing valuable feedback when the path forward was difficult to find.

I would like to thank Alejandro Rodríguez Aguilera and Carlos Castillo for helping with the fabric simulations and the voxalisation and Patricia Rodriguez Perez for supplying me with many interesting samples to investigate. Also, a big thank you to Javier Fabre for getting me up and running with the custom path tracer, for his assistance in the many debugging sessions, for the stimulating discussions and all of his help and support.

Throughout the work for this thesis, I also had the opportunity to visit an Optical Lab, run by Carlos David Heras at the University of Zaragoza to perform many of the experiments presented in this thesis. I would like to thank Carlos and Pilar Castillio for helping me with the measurements, building the setups and making me feel welcome during my stay at the Lab.

Last but certainly not least, thank you to my two supervisors. Marco Fratarcangeli at Chalmers, for his encouragement, advice and support. And thank you very very much to Elena Garcés, my supervisor at SEDDI Labs Inc., for always following up on questions and requests, keeping such good tabs on the progress of the project, for making it possible to perform the measurements at Zaragoza, for welcoming me into her team and her overall support.

I had a great experience.

Peter McEvoy, Gothenburg, October 2019

Contents

| | | |
|----------|---|-----------|
| 1 | Introduction | 1 |
| 2 | Theory | 4 |
| 2.1 | Light and Colour | 4 |
| 2.2 | Surface Reflectance and Material Appearance | 5 |
| 2.3 | Reflectance models | 7 |
| 2.4 | Fluorescence | 8 |
| 2.5 | Image Formation | 10 |
| 3 | Behaviour of Fluorescence in Fabrics | 12 |
| 3.1 | The Case for Suspecting Anisotropic and Directionally Dependent Fluorescence | 13 |
| 3.2 | Angular Dependence of Fluorescence in Yarns | 15 |
| 3.2.1 | Results and Conclusions | 16 |
| 3.3 | Angular Dependence of Fluorescence in Fabrics | 18 |
| 3.3.1 | Results and Conclusions | 19 |
| 3.3.2 | Limitations and further work | 20 |
| 3.4 | Summary | 21 |
| 4 | Practical Capture of Fluorescence | 22 |
| 4.1 | Related Work | 22 |
| 4.2 | Capture Setup | 23 |
| 4.3 | Estimating Spectral Reflectance | 26 |
| 4.3.1 | Fitting Camera Radiance Response | 27 |
| 4.3.2 | Estimating the Re-radiance Map | 29 |
| 4.4 | Evaluation | 31 |
| 4.5 | Conclusion | 37 |
| 5 | Simulating Fluorescence | 39 |
| 5.1 | Background | 39 |
| 5.2 | Spectral Rendering | 41 |
| 5.3 | Bi-spectral Reflection | 43 |
| 5.4 | Renders | 45 |
| 5.5 | Conclusion | 48 |
| 6 | Summary | 50 |



Statue of Victor Hasselblad (inventor of the Hasselblad camera) guarding vividly coloured coats and backpacks belonging to pupils on a nearby school outing in Gothenburg, Sweden.

1

Introduction

In textiles, *fluorophores* (fluorescent molecules) are mixed with dyes to help achieve exceptionally vivid colours that may be seen in the neon clothes from the 80s and more recently the vivid greens, yellows, pinks, and oranges in athletics ware. These fluorophores help achieve bright colours by absorbing light with a short wavelength (such as hard to see ultra-violet) and re-emit light with a longer wavelength (such as blue, orange or red), see Figure 1.1. Fluorophores also see use in more reserved applications, such as optical brighteners that give white fabrics an extra boost of blue, making them appear brighter and even more white in many types of illumination.




Figure 1.1: Light interaction with a fluorescent material. Incident light from a certain direction with ingoing wavelength λ_i interacts with the material. The light is then reflected in the outgoing direction with the outgoing wavelength λ_o . In a fluorescent material this outgoing wavelength will be longer than the ingoing.

When trying to faithfully recreate the appearance of such fabrics in computer graphics, this fluorescent behaviour is important to consider, especially when the scene illumination changes. In computer graphics, rendering-systems simulate light interactions to generate realistic images. The challenge with simulating fluorescence is that it places demands on the rendering-system, and the methods to capture data needed for these simulations require special laboratory equipment. The rendering-system needs to support *spectral rendering*, where the light interactions are simulated in terms of wavelengths instead of a set of three colours; red, green and blue, known as the *RGB-primaries*. Fluorescence also requires *re-radiance* data. That is, contrarily to most rendering systems and appearance capture devices, it cannot be assumed that incident light and reflected light from such materials have the same wavelength, making effects such as fluorescence difficult to accurately capture.

These challenges have been the focus of the work for this thesis.

The structure of the thesis is as follows:

- Chapter 2: Theory. A short overview of the main concepts used throughout the thesis: Fluorescence, camera image formation, rendering and material appearance models.
- Chapter 3: Behaviour of Fluorescence in Fabrics. Experiments for studying if the multiple scales of structure in fabrics introduce any angular dependence for fluorescence. Here, we can see that in fluorescent yarns and fabrics, there is no significant angular dependence. As a consequence, fluorescence can be captured from a flat sample from any distance. Yarn-level resolution is not required, provided that the fabric does not have differently coloured yarns woven with each other.
- Chapter 4: Practical Capture of Fluorescence. I present a method for capturing the fluorescent behaviour of a fabric. The method uses nine narrow-band LEDs, a consumer camera and an image analysis procedure using numerical optimisation, that estimates the wavelengths of detected light from fluorescent materials using the capture RGB values from the camera.
- Chapter 5: Simulating Fluorescence. I outline how a typical rendering-system can be extended to support the simulation of fluorescence. Finally, rendered results are compared to photographs of fluorescent fabric samples under different illuminations.

Background and Previous Work

Previously, several approaches have been presented that allow for simulation of fluorescence. Jung et al. [15] introduce a parametric model that uses a fluorescent compound's *absorption spectrum* and *emission spectrum* to inform a reflectance model. Such spectra are freely available online [3]. However, when the goal is to recreate the appearance of a physical sample it is difficult to determine exactly what compound was used in the dying process. For this use case, a more direct capture of the appearance from the actual sample is beneficial.

Hullin et al. [13] propose a setup that thoroughly samples both the *bi-directional* (measurements for each incident and outgoing direction) and the *bi-spectral* (sample for ingoing and outgoing wavelengths) intensity of light from a spherical sample with a homogeneous fluorescent paint. The setup proposed uses laboratory equipment; a strong wide-spectrum light source and two tunable optical narrow-band filters, one for the light source and one for the camera. Since this proposed appearance capture procedure samples over different directions *and* different ingoing and outgoing wavelengths the acquisition time is very long (5-8 hours). In their work Hullin et al. [13] found that the fluorescent component of reflected light showed little angular dependence, with an exception in the "Speckled" dataset. This speckled dataset was a spherical sample painted with an orange-fluorescent paint on top of which a different, yellow fluorescent, paint was sprinkled. Blurring the camera simulates seeing the object from a distance, making the sprinkled paint dots emulate microstructure (small details on the surface). Figure 3.1 shows an illustration of this dataset, note how towards the edges of the blurred sphere the colour shifts, indicating a change in fluorescence with angle. This experiment and its significance for this thesis will be revisited in Chapter 3.

Since the authors found that fluorescence mainly has very little angular dependence except potentially for when the object has microstructure, the question arises: Is it necessary to sample directional behaviour when capturing fluorescence in fabrics? If so, what would be an efficient and practical way to perform such measurements? These are the question that will be addressed in this thesis.



2

Theory

In this chapter, the general terms and phenomena of light with which this thesis concerns itself with is introduced. Additionally the phenomena of fluorescence is presented in more detail. Finally, the process with which a common camera produces images is described. In Chapter 4, this description of the camera imaging process is used to perform capture of fluorescence.

2.1 Light and Colour

Light is electromagnetic radiation with mainly two properties that affect how we visually perceive it; the intensity of the radiation and the wavelength. The light that we normally encounter from everyday light sources or that is reflected from objects is made up of a distribution of different wavelengths. In the framework of radiometry, such emitted or reflected light are described with the term *radiance*.

Human eyes are sensitive to this electromagnetic radiation in the wavelength interval of 360 nm to 830 nm, usually referred to as the *visible spectrum* of light. Within the eyes are two types of photoreceptors: *rods* and *cones*. Rods are very sensitive to light intensity but do not respond significantly to differences in wavelength. Cones, on the other hand, are sensitive to certain wavelength ranges. Cones can be divided into three sub-types; short-wavelength sensitive cones, middle-wavelength sensitive cones and long-wavelength sensitive cones. These three types of cones allow us to sense the intensity of light within three ranges of the electromagnetic spectrum. The fact that we are sensitive to three wavelength ranges suggests that three *primary colours* should be sufficient to describe all of the colours that we humans can perceive. This is referred to as the *tristimulus theory*. This *trichromatic* nature of our vision allows us to sense many different colours but also has the consequence that certain very different spectra can be perceived as being the same colour.

In 1931 the International Commission on Illumination (CIE) published findings resulting from experiments that attempted to find a relationship between how the average person perceives colour and how it relates to the electromagnetic spectrum [7]. They performed this study by asking subjects to sit in a room and to look at the light at a given wavelength. The subjects were then tasked to match the colour of this reference by adjusting the intensity of three differently coloured lights mixing red, green and blue light. From this work, the CIE RGB colour space and in turn the widely used CIE XYZ 1931 Colour Matching Functions were derived, which can be seen in Figure 2.1.




Figure 2.1: The CIE XYZ 1931 Colour Matching Functions. These functions were derived from the CIE RGB colour space, which in turn was obtained by experiments that aimed to show how an average person needs to mix light with three colour primaries in order to match light at a given wavelength. These results can also be interpreted as the colour sensitivity curves for the average human.

The XYZ Colour Matching Functions give a quantitative definition of what we mean by red, green and blue and allow us to translate the physical properties of light into perceived colour. This specific system is fundamental to many of the ways currently used to describe colour in digital systems and is the underlying assumption in most digital colour spaces that are used today, such as CIE RGB, and sRGB (see Section 2.5 for more on colour spaces).

For more details on the current understanding of colour and its perception, the first chapter in *Digital Color Handbook* [24] gives definitions, a rigorous overview on the perception of colours and also explains in more detail how the CIE XYZ and other colour matching functions were obtained.

2.2 Surface Reflectance and Material Appearance

When looking at a red ball, what factors are involved for this ball to have its appearance? Firstly, the material of the ball itself has a central role in determining its colour and appearance, which, in case of the red ball, probably reflects most light with wavelengths that correspond to red and absorbs the rest. This behaviour can

be properly expressed by a surface *reflectance spectrum*. Such a spectrum describes the relative amount of light that is reflected off the surface at a given wavelength. A critical underlying assumption here is that the material does not change the reflected light's wavelength, meaning the reflectance is only dependent on a single wavelength variable. The reflectance spectrum is thus a so called *uni-spectral* function.

Next, if we change the light illuminating the ball to a coloured light, the appearance of the ball will also change (e.g. a white paper under red light will appear red), meaning that the illumination also affects the appearance. An *illumination spectrum* can describe the distribution of wavelengths emitted from a light source in a similar fashion to a reflectance spectrum.

Finally, imagine that instead of a human eye looking at the ball, there was some other observer (such as a camera or a cat) looking at the same ball. The sensors in this observer might not sense light in the same way as humans, they may be sensitive to different parts of the spectrum. Thus, the sensor of the observer also affects the perceived appearance. For this, *spectral sensitivity functions* are used to describe what distribution of wavelengths contribute to a perceived *primary colour* (referred to as a *channel* in the context of a digital camera), in many standard colour cameras and for humans, these primary colours are red, green and blue. See the CIE XYZ colour matching functions (Figure 2.1) for an example, since these functions can also be interpreted as spectral sensitivity functions.

We thus have three main components that contribute to an object's colour appearance:

- Surface reflectance spectrum $r(\lambda)$ – How much light a surface reflects at a given wavelength λ .
- Illumination spectrum $I(\lambda)$ – Intensity of light emitted from a light source at a given wavelength λ .
- Receptor channel k 's spectral sensitivity $c_k(\lambda)$ – Distribution of wavelengths that contribute to the observer's channel k . There is one function for each channel k .

Mathematically the relationship between these components can be modelled by,

$$I_k = \int r(\lambda)I(\lambda)c_k(\lambda)d\lambda, \quad (2.1)$$

where I_k is the observer's perceived intensity of the colour primary (or channel) k .

This model neglects many things that could be added to more accurately model the appearance of materials. Especially the surface reflectance $r(\lambda)$, which is considerably simplified here as being a function of one wavelength only. The model could be extended to include the position of the point of reflection on the surface as well as the light's ingoing and outgoing directions to that point on the surface. As we will see in Section 2.4, for effects such as fluorescence, we must also separate the dependence of the wavelength into a dependence on both the ingoing wavelength

and outgoing wavelength. The ingoing wavelength is the incident light’s wavelength, that is light coming directly from a light source or that has been reflected of some other object and then bounced towards this point on the surface. The outgoing wavelength, is the wavelength of the resulting light from the surface interaction, that is simply reflected light or light that has been emitted via fluorescence (from now on the term “reflected light” can also encompass emitted light from fluorescence). This new dependence on both ingoing and outgoing wavelength makes the reflectance of a fluorescent material *bi-spectral*.

2.3 Reflectance models

A more sophisticated way of describing how a surface reflects light is achieved by using a *Bidirectional Reflectance Distribution Function* $f_{\text{BRDF}}(\lambda, \omega_i, \omega_o)$. This function describes the intensity of the reflected light from a surface given a wavelength λ , the ingoing (light) direction ω_i and the outgoing (viewing) direction ω_o . The light’s ingoing and outgoing directions are specified relative to the surface’s normal vector n . See Figure 2.2.




Figure 2.2: A BRDF model describes the relative intensity of light being reflected from a surface given the ingoing ω_i and outgoing ω_o light direction relative to the surface normal n , for wavelength $\lambda = \lambda_i = \lambda_o$. For a bi-spectral BRDF, we can no longer assume that $\lambda_i = \lambda_o$.

The simplest BRDF is the Lambertian Diffuse surface model that has no dependence on ingoing or outgoing direction and gives a constant value b for these variables. It does have a function of wavelength to describe the surface’s colour d ,

$$f_{\text{Lambertian}}(\lambda, \omega_i, \omega_o) = b \cdot d(\lambda).$$

This BRDF has no angular dependence and can be used to model the appearance of objects which have no shiny reflections.

Many other BRDF models exist to mimic different types of materials with parameters for surface features. Some BRDF functions are also purely based on measured data [21].

To model the radiance received by an observer from a point on this object, the BRDF must be integrated similarly to Equation (2.1), however, here viewing and illuminating directions are taken into account. In computer graphics, this formulation is referred to as the rendering equation [17], which is used to simulate light interactions in a scene,

$$L_o(\omega_o, \lambda) = L_e(\omega_o, \lambda) + \int_{\Omega} f_{\text{BRDF}}(\omega_i, \omega_o, \lambda)(\omega_i \cdot n)L_i(\omega_i, \lambda) d\omega_i. \quad (2.2)$$

Here, $L_o(\lambda, \omega_o)$ is the spectral radiance from the object directed toward the viewing direction ω_o . $L_e(\omega_o, \lambda)$ is an emissive term if the object emits light. $L_i(\lambda)$ is the spectral

radiance of the light incident on the object, this incoming light can come directly from a light source or from another object, evaluated recursively by using the same equation (2.2). The factor $(\omega_o \cdot \mathbf{n})$ is the dot product between the surface normal \mathbf{n} and the angle between the incident light and the surface normal. This factor will cause the reflected radiance to be less when light direction and surface normal are not aligned, giving a shading effect, even in the case of the simple Lambertian BRDF.

A final important principle often utilised in reflectance models and in the study of light is that reflected light can be split into two components; *diffusely* reflected light and *specularly* reflected light. A rough and unsmooth surface will cause incoming light to be reflected in many different directions. Additionally, light that enters into a surface and interacts with the material below the surface before exiting out again will also have the light direction altered drastically in different directions. These sorts of interactions will give diffuse, light that is not very dependent on the viewing or illumination direction. On the other hand, when light hits a surface that is very smooth, the light will most likely be reflected along the ideal angle of reflection, without any significant deviation from this angle. This gives a specular reflection, reflection that is very dependent on viewing and illumination direction. Many materials give *both* diffuse and specular reflections and it is often useful to be able to theoretically separate these behaviours as two different components when studying material appearance.

2.4 Fluorescence

Fluorescence occurs when fluorophores absorb light with short wavelength and re-emit light with a longer wavelength. The extent to which these molecules absorb incoming light at a specific wavelength is determined by the fluorophore's *absorption spectra* $a(\lambda_i)$. The absorbed energy causes the fluorophores to attain an excited state. Fluorophores release the absorbed energy by relaxing to different energy states and thereby emit light with a wavelength corresponding to the energy difference between the states. The molecular structure of the fluorophore determines the levels of these energy states and the number of energy state transitions available. However, in photochemistry it is known that fluorescence only occurs in a significant amount as the molecule relaxes from the next lowest energy state to the base state, since higher energy states are placed too close to each other for significant photon emission. This principle is referred to as *Kasha's rule* and has the consequence that regardless of the wavelength of incoming light, the appreciable re-emitted energy will be of a certain wavelength depending on the fluorophore. Thus, fluorophores will emit the same wavelengths of light no matter what wavelength was used to illuminate them with, given that the illuminating wavelength is shorter than the fluorophores emitting wavelength (since the exciting energy needs to be greater than the emitted energy). However, the *intensity* of the emitted light, is determined by the wavelength of the incident light via the absorption spectra $a(\lambda_i)$ of the material. The distribution of the re-emitted light is described by an *emission spectra* $e(\lambda)$.




Figure 2.3: Spectra showing the response from the three sample fabrics that will be studied throughout this work. The samples have been illuminated with 380 nm light. From the spectra we see that some of the light has been reflected and not been absorbed by the fluorophores, thus keeping the wavelength of the light source. At the higher wavelengths an additional peak of light is detected, this is light that has been absorbed and then re-emitted by the fluorophores. Near 620 nm for “Pink Elastic”, 500 nm for “Yellow Fluorescent” and 440 nm for “White Gaberdine”. The centre wavelength and shape of the re-emitted light is determined by the molecular structure of the fluorophores.

The radiance from fluorescence L_f can be mathematically described as,

$$L_f(\lambda_o) = e(\lambda_o) \int_{\Lambda} l(\lambda_i) a(\lambda_i) d\lambda_i, \quad (2.3)$$

where $l(\lambda_i)$ is the illumination spectrum of the incident light and Λ is the visible spectrum which we integrate over. To demonstrate, Figure 2.3 shows how three fluorescent fabrics respond to ultra-violet light.




Figure 2.4: The image capture process with a digital camera. Light is reflected off an object and enters the lens. The lens has different levels of transmission depending on the light’s wavelength, this is described by the lens’ *transmission spectrum*. Once through the lens, the light hits the camera’s sensor. Over the sensor is a pattern of colour filters for red, green and blue. The sensor will record the amount of light under each of the filters. The interplay between the filters and electronics determines the camera’s *spectral sensitivity curve*. Finally, the detected values are converted into a standard colour space that can then in turn be converted to more specific colour spaces, suitable for an image file, such as a JPEG.

2.5 Image Formation

In Chapter 4, we will look at how a normal digital camera can be used to infer the wavelength of the detected light. To know how to infer this, it is useful to review what happens with the wavelengths of light when taking a photograph of a scene. Figure 2.4 shows an illustration of the steps involved in sensing light for a photograph and turning the light reaching the camera to a digital image file.

Camera and Lens

Light first enters the camera’s lens (middle of Figure 2.4). The lens’ primary role is altering the path that light takes, but lenses can also affect the spectrum of incoming light by allowing more or less light through according to wavelength. That is to say it act as as a filter. This is often reported in the lens’ datasheet as a *transmission spectrum*.

After passing through the lens, the light hits the camera sensor, consisting of many thousands of small photodiodes (sensors). This interaction will induce a charge in the diode that can then be measured, giving a value proportional to the intensity of light.

For colour capture, different filters are placed over each photodiode. In many consumer cameras, a Bayer filter mosaic is used, a set of red, green and blue filters that are placed over the photodiodes in a certain pattern. By reading the sensed light from the photodiodes corresponding to the filter pattern, three channels of colour

are obtained. It is at this point we lose the spectral nature of our light and are now dealing with colour in terms of channels. The model for channel intensity I_k , Equation (2.1) as introduced in Section 2.1, also describes the same process with a camera.


How the camera senses the colour from a spectrum is described by the camera's *spectral sensitivity function*. This function alters between camera manufacturers and models since it depends on the filters used and the way the electronics alter the signal. Often this spectral sensitivity is reported in the camera's datasheet, but there are many cases where they are not or that the accuracy of the data is questionable.

Digital Colour

With the photodiode measurements, the camera will sort the values corresponding to the filter mosaic pattern to get values for each channel (red, green and blue), a process referred to as *de-Bayering*.

When photographing, we wish for colours that match what we perceive. However, the camera's spectral sensitivity functions are seldom consistent between models or even similar to the sensitivity functions of the average human. Therefore the channels are rescaled to try to compensate for the differences in spectral sensitivity and then transform the colours into a standard *colour space*, such as CIE XYZ. This transformation is determined by the manufacturer via a calibration process as well as the camera's white balance setting.

Finally, the camera converts the colours to a colour-space suitable for an image file and stores it on the camera's memory card.



3

Behaviour of Fluorescence in Fabrics

Before trying to design a good way to capture and simulate the fluorescence in fabrics, we first need to know how fluorescence behaves, especially in the context of fabrics.

Fluorescence occurs when certain types of molecules (*fluorophores*) absorb light and re-emit it as light with a longer wavelength. Fluorescent compounds often come in the form of homogeneous liquids that can see direct use in medical imaging. In these liquids the fluorescence will be scattered and emitted equally in all direction, there is no special directional dependence. In textiles, such fluorescent solutions are mixed with dyes to help achieve vivid colours. Textiles are highly structured with multiple levels of scale, where each level can introduce interesting new behaviour.

Previous radiance models provide several alternatives to model fluorescence in surface-like materials or as scattering in homogeneous liquid or gaseous media. However, most of these models have treated fluorescence as an *isotropic* or diffuse effect, i.e. fluorescence is assumed not to show dependence on ingoing and outgoing light direction (it has little or no angular dependence). However, as will be argued for in more detail in Section 3.1, the angular and anisotropic behaviour of complex volumetric materials, such as fabrics, deserve to be studied in isolation.

To help decide an appropriate method in capturing and simulating fluorescence in fabrics we want to investigate how and if the different scales of structure in fabrics introduce significant changes in the fluorescence dependent on the light's ingoing and outgoing direction and if so, at what level.

In this chapter, the following questions will be addressed:

- What, if any, are the effects on fluorescence when altering the angle of incident light with respect to a yarn (a yarn consists of multiple fibres twisted around each other). (Section 3.2)
- Does the structure where multiple yarns woven into a single fabric give any angular dependence for fluorescence when the fabric is viewed from afar? (Section 3.3)

3.1 The Case for Suspecting Anisotropic and Directionally Dependent Fluorescence

A material's surface reflectance being *directionally dependent* means that the amount of light reflected is highly dependent on the light's ingoing and/or outgoing direction, contrast this to the Lambertian diffuse reflectance introduced in Section 2.3, where there is no angular dependence. A surface material being *anisotropic* means more specifically that its reflectance changes when rotated around its surface normal. Brushed metal and cloth are materials that exhibit this behaviour and are thus anisotropic. A flat piece of brushed metal or shiny silk fabric on a flat table will reflect light in different amounts depending on the rotation of that material on the table.

Confusingly, when talking about volumetric media, such as liquids and gases, the terms isotropy and anisotropy have different meanings and can be used in two different contexts. They can describe the medium itself: an anisotropic medium means that the particles are aligned in a coherent structure. This can introduce directional dependence on how light is reflected and scattered through the medium. The way the particles that make up the medium scatter light is described by a *phase-function*. Here, isotropy and anisotropy can also describe if the particles scatter light equally in all directions or is weighted to scatter more in a certain direction. See [23, Ch 11.2] for more information. In addition to liquids and gases, volumetric media can be used to model materials that demonstrate significant light scattering inside the material such as in skin and fabrics.

What is the motivation for suspecting that fluorescence in fabrics might exhibit directional and anisotropic behaviour?

In previous work on modelling the radiance from fluorescence, Wilkie et al. [27] analysed the behaviour of fluorescence from a cardboard surface and introduce an analytical BRDF for simulating the fluorescence and its behaviour on diffuse surfaces. To inform their model, the authors performed experiments with a monochromatic laser aimed, with different angles of incidence, at an orange fluorescent flat cardboard sample. They note that, as the laser beam was brought closer to grazing angles, the amount of fluorescence decreases. To recreate this behaviour in their analytical model, the surface was modelled as two layers, a transparent “varnish-layer” (a Fresnel micro-facet model) and a fluorescent non-angular dependent, Lambertian




Figure 3.1: Left and middle: The "Speckled"-dataset photos from by Hullin et al.¹[13]. Two layers of fluorescent paint are used on this sample, a yellow base and orange sprinkled on top, emulating microstructure. When viewing from afar (simulated by the blurry image) we see that colour shifts to orange as we get closer to the edge of the sphere. An indication that fluorescence can show angular dependence if there is microstructure. Right: Dual yarn colour fabric, the orientation of the surface shows a different combination of warp and weft yarns in the weaving pattern due to self-occlusion, this is an example of the structure of the fabric introducing anisotropic diffuse behaviour.

base layer, thus attributing the angular dependence of the fluorescence to Fresnel effects.

In the work "Acquisition and analysis of bispectral bidirectional reflectance and reradiation distribution functions" [13], Hullin et al. measure fluorescent paints on spherical samples and their re-radiance for different viewing and lighting angles. The authors introduce the concept of a *Bi-spectral Bi-directional Reflectance and Reradiance Distribution Function* (BRRDF, the bi-spectral equivalent of the BRDF) and use the measured data for analysis and later rendering. From their analysis of different paints, they note that the fluorescent component of the re-radiated light demonstrates weak angular dependence with regard to illumination and viewing direction, somewhat contradicting the results from [27]. However, when looking at their "Speckled" dataset, a dataset mimicking material with microstructure, the fluorescent component demonstrated more significant angular dependence, see Figure 3.1. These insights are interesting when considering fabrics.

Abdellah et. al. consider fluorescence in scattering participating media [1]. Synthetic or publicly available *absorption* and *emission* spectra are used to inform the fluorescence behaviour. The scattering of the fluorescence is considered to be isotropic, making this model well suited for simulating liquid or gaseous media with fluorescence. In computer graphics, fibre level simulation of fabric radiance is possible using volumetric scattering models [18, 2], giving impressive results. However, the medium of fabrics and fibres, with its many different scales of structure, is far

¹Matthias B. Hullin, Johannes Hanika, Boris Ajdin, Hans-Peter Seidel, Jan Kautz, and Hendrik P. A. Lensch. 2010. Acquisition and analysis of bispectral bidirectional reflectance and reradiation distribution functions. In ACM SIGGRAPH 2010 papers (SIGGRAPH '10), Hugues Hoppe (Ed.). ACM, New York, NY, USA, Article 97, figure 7. DOI: <https://doi.org/10.1145/1833349.1778834>

from the homogeneous gases and liquids studied in [1].

The three levels of structure in fabrics are *fibres*, *yarns* and the woven or knitted *fabric* itself. At each of these levels, the structure can potentially introduce new directional behaviour. The different levels can also hide or obfuscate directional behaviour from smaller levels, such as when fibres are twisted around each other to form a yarn.

Fibre Fibres can be thought of as very small and long cylinders. How these fibres reflect light is a subject of active research that shares many principles in common with hair-fibre scattering [20, 5, 2]. The mechanism by which light is scattered in a fibre depends on the fibre's shape, roughness, dye and other properties. An additional interesting behaviour we see from the scattering of fibres is that light that enters the fibre, i.e. does not directly reflect off the surface as a specular reflection but instead scatters internally through the fibre, can also show directional dependence with respect to illumination and viewing angles [2, 18]. We thus see *anisotropic* scattering in the diffuse (non-specular) components of the scattered light. Would fluorescence in such fibres also demonstrate such anisotropic behaviour?

Yarn The next level is a yarn, this consists of multiple fibres that have been twisted around each other. Yarn introduces structured complexity and, just as with fibres, the specular reflections will be highly dependent on lighting and viewing direction, however, as discussed in fibres, the diffuse non-specular reflections could potentially see directional dependence depending on the structure of the yarn. If it is a type of yarn where the fibres are highly aligned, such as in silk, some anisotropic behaviour of the diffuse light might be expected. But if the yarn fibres are not aligned, i.e. has a lot of twisting or is fuzzy, much of this directionality is lost by the random irregularities.

Fabric The final scale encompasses the pattern with which the yarns are woven or knitted. This pattern can heavily influence the reflective behaviour of the fabric and introduces even more structure. Also, should a weaving-pattern incorporate two different types of yarn, even more complex behaviour can be introduced, see the rightmost image in Figure 3.1 for an example of how a weaving pattern with two differently colour yarns can introduce directionally dependent diffuse behaviour.

3.2 Angular Dependence of Fluorescence in Yarns

In this section, we perform an experiment in order understand if the fluorescence of a yarn, composed by multiple fibres twisted around each other, varies with 1) the direction of the light source and; 2) the position of the yarn. Given the cylindrical geometry of a yarn, we assume that the observed radiance is symmetric in view ω_o and light directions ω_i .




Figure 3.2: Photo and schematic of the measurement setup for yarn measurements.

Setup The practical setup (see Figure 3.2) is built on top of a standard optical plate where a fixture holds and keeps the yarn stretched. This fixture can be rotated to position the yarn in a vertical or horizontal orientation. A spectrometer is aligned to measure the reflected light from the yarn. A Xenon-arc lamp mounted on a rotating arm emits light collimated through a lens and filtered with a band-pass filter allowing light near 400 nm to pass through. When the yarn is in the horizontal position the rods holding the ends of the yarn blocks the view for the spectrometer. To avoid this blocking, the spectrometer, when the yarn is in the horizontal position, is rotated clockwise around the yarn so that it has a 45° angle from its original position. Thus, when the yarn is in the horizontal position and the light is at 45° , then the angle between the spectrometer and light direction is 90° .

Measurement First, the yarn is fixed in the vertical or horizontal position. Then, light is rotated so that it pivots around the yarn at angles 30° , 45° , 60° , 75° , 90° and 105° . For each of these angles, the spectrum detected by the spectrometer is recorded. Since we are illuminating with narrow-band light near 400 nm, the reflected spectrum will have two peaks, one *reflected peak* near 400 nm (this includes specular and regular diffuse reflections) and one *fluorescent peak* near the wavelength for the fluorophores (in this case near 500 nm).

3.2.1 Results and Conclusions

The peak values for the measurements per light angle and yarn orientation are plotted in Figure 3.3. The results show that the intensity of the fluorescence barely changes due to angle when compared to light arising from specular and diffuse reflections. This is true both for when the yarn is in a vertical and in a horizontal orientation.

Thus, at least for single yarns, the experiments show that angular dependence of fluorescence is not significant relative to the intensity of non-fluorescent reflected light. Even though this test does not rule out angular dependence of fluorescence at the fibre-level, it does suggest that, at the yarn-level, even if there were such an anisotropic behaviour, it would be lost because of the yarn's fibres being twisted

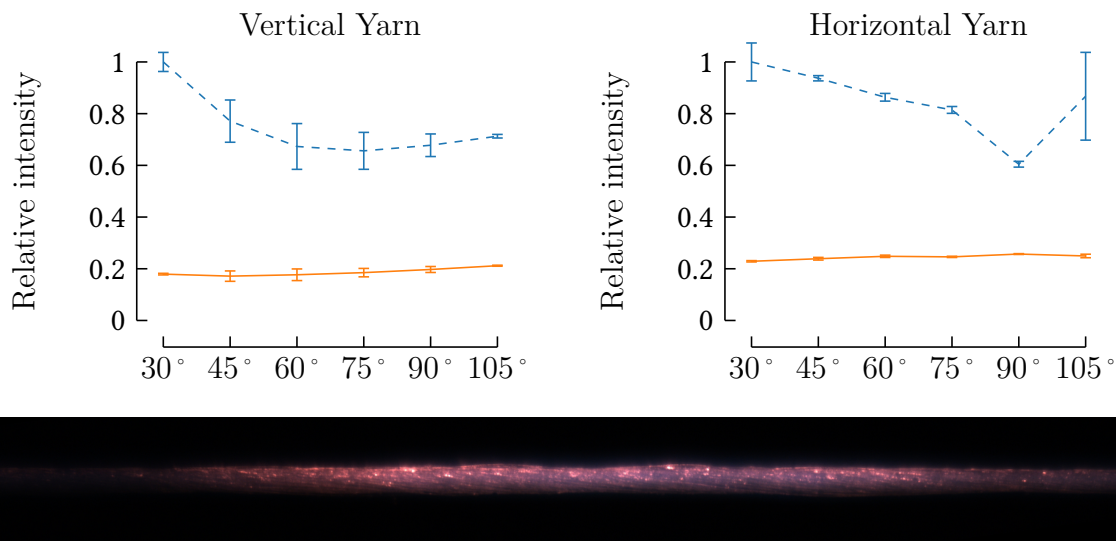


Figure 3.3: Plots of the peak relative intensities for reflected light from fluorescence (—) and other reflected light (- - -) from the yarn sample as a function of illumination angle. The plotted result is an average of 5 and 2 repeated measurements, respectively for vertical and horizontal orientations. The bottom horizontal image shows the yarn used for the measurements.

around each other.

With that said, there is a slight linear relationship of higher fluorescence as the light angle increases (more clearly seen in the vertical case). An explanation for this is that it is due to more of the yarn surface is being illuminated as the angle increases and is not due to any directionality in the fluorescence.

Limitations and further work The measurements made here are somewhat crude and can only tell us about significant angular dependence of fluorescence. As seen in the bottom image of Figure 3.3, the yarn used during measurements has some twist but still has its fibres somewhat aligned and is not fuzzy. However, it would have been interesting to also examine yarns/threads known for their high directionality (such as silk). However, fluorescent silk was unavailable during the experiments. To give more definitive answers to the question of anisotropic fluorescence at the fibre-level, we would ideally make measurements directly on fibres or on yarns with highly aligned fibres using a monochromatic laser. Such measurements would be similar to the those performed by other studies [27].




Figure 3.4: Schematic of the measurement setup for identifying potential angular dependence of fluorescence in fabrics. The setup consists of a styrofoam sphere onto which the fabric sample was wrapped. In the state shown there is a 45° angle between the hyperspectral camera and the direction of the incident light. The camera is capable of measuring light from 400 nm to 1000 nm. The Xenon-arc light source, optical bandpass and focusing lens are mounted on a movable arm.

3.3 Angular Dependence of Fluorescence in Fabrics

Next, we study the phenomena in full fabrics, using a hyperspectral camera and spherical samples. The concern here is if the micro-structure introduced by the many woven yarns introduces some special angular behaviour for the fluorescence. This has been demonstrated to be possible in the “Speckled” dataset from [13] (Figure 3.1). We expect this to be the case if the fabric were to be woven using multiple types of yarns (differently dyed), but is such behaviour apparent in a cloth woven with identical yarns?

Setup The setup (see Figure 3.4) consists of a 60 mm diameter styrofoam sphere fixed at the corner of the optical plate where the fabric sample is wrapped around. A hyperspectral camera² was used to measure spectral samples with 3 nm of resolution, in the range from 400 nm to 1000 nm, for each pixel. The camera lets us simultaneously measure the spectral radiance from the material at multiple illumination and viewing directions (ingoing and outgoing directions) by looking at the different points of the sphere. The same Xenon-Arc lamp with lens and filter from the previous experiment was used here. The lens was moved closer to the light source in order to get a wider focused light. This lens adds some complexity to the analysis since now the light cannot be considered as arising from a point source, but rather it results in the source having a central high-intensity area and a peripheral fall off region. This requires the data to be compensated. In hindsight this lens should have been removed to better approximate a point light and avoid these complications in the setup.

²Hyperspectral camera used: SpecimIQ, <http://specim.fi/iq>.

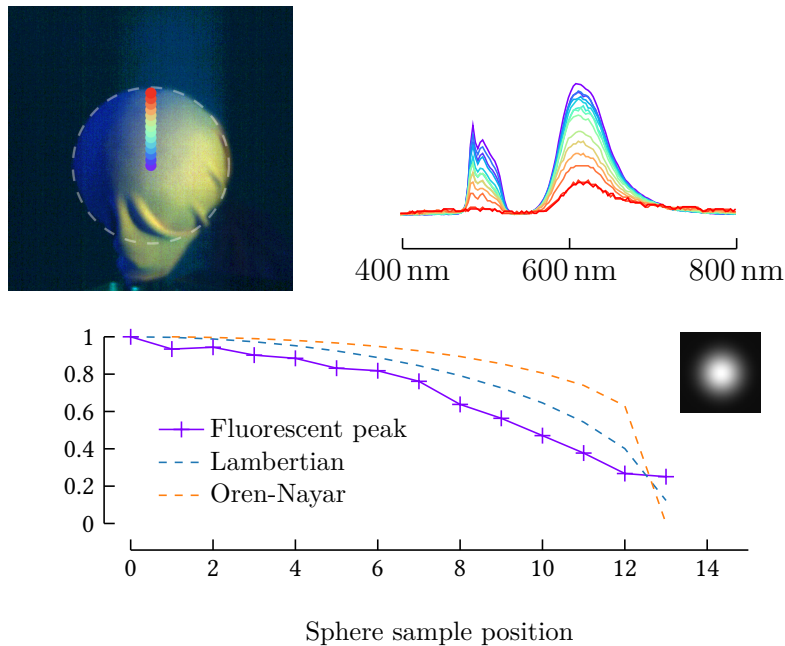


Figure 3.5: Top left shows an RGB representation of a hyperspectral photo of a fluorescent fabric where sample points have been marked using colours. The spectra of each of these sample points are shown in the top-right plot. Each spectrum has been scaled to compensate for the light’s angular falloff. The bottom left plot shows the relative intensity of the fluorescent peak and how it changes with sample position (position 0 is nearest the centre). For comparison, inferred normals at those positions on the sphere are used to calculate the Lambertian and the micro-facet Oren-Nayar diffuse responses. The image with the circular shape shows the estimated falloff region of the lens and light source used to compensate the measurements.

Measurements A hyperspectral photograph was taken of the spherical fabric sample illuminated by filtered light near 500 nm from a 45° angle (from the point of view of the camera). Using the setup dimensions, the normal vector for each point of the photographed sphere can be inferred which is later used to compare the falloff with standard shading models.

3.3.1 Results and Conclusions

The top row of Figure 3.5 shows spectral measurements taken from the centre of the sphere towards the perimeter. We observe that normal reflection occurs near 500 nm and fluorescence near 600 nm. It is shown that both peaks behave similarly, reducing in intensity as the samples move further up the sphere. We then plot the relative intensities of the fluorescent peaks at different positions in the sphere (bottom of Figure 3.5). We see that the intensities gradually reduce following the shape of a Lambertian fall off, meaning that that the reduction can be attributed to the geometry of the sphere. For completeness, we also compare with a micro-facet model (Oren-Nayar).

For capture, that means that we can measure the fluorescent behaviour from afar, at the fabric level, without worrying that the structure of the fabric alters the fluorescence in a significant manner. This is true, at least for fabrics that are coloured by one dye, i.e. is not made of differently coloured yarn (such as the one in Figure 3.1).

3.3.2 Limitations and further work

Fabrics are complex materials and are not modelled well by simple surface appearance models such as a Lambertian diffuse or a micro-facet model. However, the fluorescent component in isolation might exhibit such a simple Lambertian-like behaviour. Based on the results from the measurements presented in this chapter, this appears to be an agreeing model. However, the setup has limitations that makes it hard to rule out angular dependence or anisotropy in fluorescence. The following points will discuss some of these limitations and possible approaches to address them.

Fresnel As light is incident on a surface with increasing grazing angle a higher proportion of the light will be reflected specularly, i.e. will not enter the material. In the measurements, comparing with the Lambertian and micro-facet surface material models, Fresnel effects are not considered. This was done because a fabric with its many yarns will probably not have such a strong Fresnel effect as a smooth surface would. Taking the results from the yarn measurements into account we have even less reason to expect strong Fresnel effects. However, these effects should ideally be considered.

Pattern orientation One of the concerns at the fabric level was that the repeated structure of the woven yarns would introduce some new directionally dependent and anisotropic behaviour of fluorescence at the macro scale. From the measurements, this does not seem to be the case. However, the possibility anisotropy in fluorescence was not tested well. The same sort of measurements should have been done for several different rotations of the fabric, causing the pattern to be aligned differently.

Forward-scattered light In the setup used for the fabric measurements, the light source is beside the camera. This means that our measurements result from light that has hit the fabric's yarns and is then reflected back to the camera. That is, we see the light that has mostly been *back-scattered*, which often means more diffuse-like behaviour. The setup should ideally be extended to also have the light behind the spherical sample as a backlight, thus giving more *forward-scattered* light for measurements.

BRDF measurement device With a proper spectral BRDF measurement device, the concerns of pattern orientation and forward/back-scattering could be taken


care of all at once. The device presented in “An adaptive parameterization for efficient material acquisition and rendering” [6] is a good example of such a measurement device and setup that would give highly accurate data on any potential angular dependence. Such a setup could be extended with a band-pass filter for the illuminating light, separating the fluorescence in the measured spectrum, allowing analysis of only the fluorescence.

3.4 Summary

From the measurements we saw:

- At the level of the yarn we see no special angular dependence for the intensity of fluorescence, suggesting it is an isotropic effect. This might not be the case at the fibre-level, but could still be modelled as one, since such behaviour is likely lost when going up to the yarn-level.
- For fabrics, the repeated micro-structure from the woven yarns does not appear to introduce any special angular dependence for fluorescence and its behaviour mimics that of a diffuse surface shader. The angular dependence we do see in the images can be attributed to the shape of the geometry.

Based on these measurements, it should be safe and accurate to simply capture the fluorescent behaviour from a fabric sample that lays flat.



4

Practical Capture of Fluorescence

We now shift our focus towards finding a practical way to capture the necessary information required to simulate the fluorescent effect. We are looking for an efficient way to perform *appearance capture* of a fluorescent material.

Previous work by Hullin et al. [13] has been able to recreate fluorescence by capturing a full bispectral BRDF. However, given that fluorescence seems to have isotropic angular behaviour (as shown in Chapter 3), we can achieve similar performance without regard for view and illumination direction. Instead, sparse spectral samples are captured using estimated with narrow-band LEDs and a conventional RGB camera.

4.1 Related Work

A material BRDF or scattering BSDF can be captured in order to feed a data-driven model for rendering [21] or to help find values for a parametric shading model [18]. Such BRDFs can be measured using a *gonioreflectometer*, a specialised instrument for sampling the reflectance from a material while altering the viewing and illuminating direction. See Dupuy et al. [6] for an overview of recent appearance capture setup using a gonioreflectometer with new sampling techniques to minimise capture time. Such a setup is costly and the capture times are often prohibitively long.

The majority of capture techniques are limited to uni-spectral materials (where incident wavelengths are the same as the outgoing wavelengths). Hullin et al. [13] proposed an image-based setup similar to that reported by Matusik et al. [21] by using tunable optical narrow-band filters in front of the light source and the camera, allowing for the spectral control of incident light reaching the camera, giving data that can be used to simulate the appearance of fluorescent surfaces. The extent to

which this proposed setup can be used to inform fluorescence in volumetric media has not been studied. When analysing data from this setup, the authors found that light from fluorescence shows little angular dependence in most samples, except potentially where microstructure is introduced (a relevant result when considering fabrics).

In computer vision, some methods have focused on isolating fluorescence from normal reflectance in RGB images by using the property of fluorophores that makes chromaticity (color, not intensity) of their emitted light from fluorescence invariant to illumination changes [29, 9]. Recently, Koyamatsu et al. [19] proposed such a fluorescence separation technique using narrow-band LEDs, which are becoming more accessible, and statistical prior knowledge on fluorescence. These separation techniques can be used to perform relighting operations of RGB photographs with fluorescent objects. However, the information retrieved is not sufficient to inform a simulation in a ray-tracing rendering context and cannot be used when creating completely new synthetic scenes. For rendering, knowledge of the underlying wavelengths and responses are needed to inform a simulation, especially when considering indirect illumination (light that has bounced from one or more objects before hitting the object we are interested in). An image-based relighting technique, such as the one proposed by [19] is not enough.

4.2 Capture Setup

The capture setup is shown in Figure 4.1. It consists of a set of narrow-band LEDs and a RGB camera. The LEDs have been chosen to have a reasonable spacing over the visible spectrum. Figure 4.2 lists model names for the LEDs and an approximation of their illumination spectra using Gaussian distributions, where mean λ_m and variance σ_m are chosen to fit the specification in the manufacturer’s datasheets. The use of narrow-band LEDs gives us control over one side of the bi-spectral capture (spectral control of the illuminating light). The LEDs have been modified by cutting off the top plastic dome and inserting them into a casing with a collimating lens to improve their effectiveness and the uniformity of their light. There is a potential risk that these lenses have will alter the spectra of the LEDs. However, since the LEDs are narrow-band, any such alteration will show as a change in scale and not significantly in shape. This change in brightness will be compensated during calibration.

Besides spectral control of the incoming illumination, we further need spectral measurements of the outgoing, reflected and emitted light (specular + diffuse + fluorescent). Ideally, a hyperspectral camera system would be used since this allows capture of the reflected radiance’s spectrum at different wavelengths. However, these camera systems are expensive and difficult to integrate into a practical low-cost setup. Some methods try to achieve spectral imaging using more accessible parts and equipment [4, 12, 11]. Inspired by such approaches (in particular by Park et al. [22]), we use a conventional RGB camera which responds to three spectral channels (red, green, and blue). As we will see presently, such cameras with controlled narrow-band il-



Figure 4.1: Capture setup. The fabric is wrapped around a planar surface. Left, a holder where the individual narrow-band LEDs can be mounted; middle, a RGB camera and lens that points to the fabric; right, a spectrometer that measures the spectrum of the reflected light from the fabric, which will serve as a ground truth.

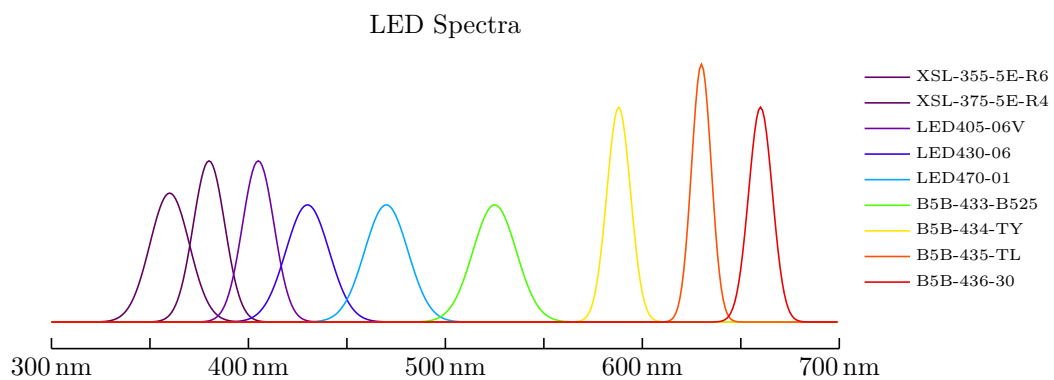


Figure 4.2: Emitted spectra up to scale from the LEDs according to their datasheets




Figure 4.3: Subset of the HDR images taken of three sample fabrics that have fluorescing behaviour. The three sample fabrics consists of the knitted “Pink Elastic Fluorescent”, “Yellow Fluorescent” that has a satin-like weave and “White Gaberdine” which has a gaberdine weaving pattern. For each sample nine HDR captures where performed, one for each illumination given by the available narrow-band LEDs, here preview images for only three out of the nine illuminations are shown.

lumination, are sufficient for estimating the wavelengths of the reflected light at a reasonable level of accuracy.

The RGB camera has a macro lens that magnifies the image so that individual yarns are clearly visible. As we concluded in Chapter 3, the fluorescent phenomenon of a fabric can be captured at the fabric level, provided that the geometry of the fabric is reasonably flat. Thus, having a level of magnification so that individual yarns are seen is not necessary. However, I found that in practice, a macro lens does help avoid issues such as narrow and non-uniform lighting that the LEDs gave. The camera is set to capture a set of images with varying exposure times that are then combined into a High Dynamic Range (HDR) image. A few of the captured images of the three fabrics used for validation can be seen in Figure 4.3. For each fabric, nine HDR images are captured, one for each LED. Additionally, on the right side of the camera in the setup, there is a spectrometer that is used to capture ground truth spectral data, used later only for evaluating the fitting performance.

Calibration The individual LEDs vary in brightness, therefore their relative intensities need to be measured to compensate for during estimation. To do this, we point the LEDs at a spectrometer from a fixed distance and record the amplitude




Figure 4.4: Spectral sensitivity functions for the camera used (Genie Nano c4900) and have been weighted by the Lens' spectral transmittance. Spectral sensitivity functions and spectral transmittance have been retrieved from the data-sheets for the camera and lens.

of each LED. Then, choosing one LED as a baseline, the brightness factors b_m of the other LEDs are specified relative to it. A spectrometer was used here but is not necessary, any method that can linearly measure the radiance would work.

4.3 Estimating Spectral Reflectance

The value captured by a camera when illuminated with a certain light source, m , can be modelled by a standard camera reflectance response model [22]:

$$I_{km} = \int f(\lambda_o) c_k(\lambda_o) l(\lambda_o) d\lambda_o, \quad (4.1)$$

where I_{km} is the sensed intensity for channel $k \in \{R, G, B\}$, c_k is the spectral sensitivity of the camera for channel k (see Figure 4.4), and f and l are the surface reflectance and illumination spectrum, respectively. This formulation, however, ignores the possibility that the surface reflectance f can respond by emitting light outside the illuminating spectrum of $l(\lambda_o)$, which is the case for fluorescent materials.

We reformulate the equation to allow for bispectral radiance via fluorescence,

$$I_{km} = \int r_m(\lambda_o) c_k(\lambda_o) d\lambda_o. \quad (4.2)$$

Here $r_m(\lambda_o)$ is the reflected radiance from a fluorescent material under illumination m , defined as

$$r_m(\lambda_o) = \underbrace{f(\lambda_o) l_m(\lambda_o)}_{r_{m,uni}(\lambda_o)} + e(\lambda_o) \overbrace{\int a(\lambda_i) l_m(\lambda_i) d\lambda_i}^{s_m}_{r_{m,bi}(\lambda_o)}. \quad (4.3)$$

The first component, r_{uni} , is the normal reflected light from the material under illumination m , resulting from a uni-spectral interaction, i.e. did not undergo a

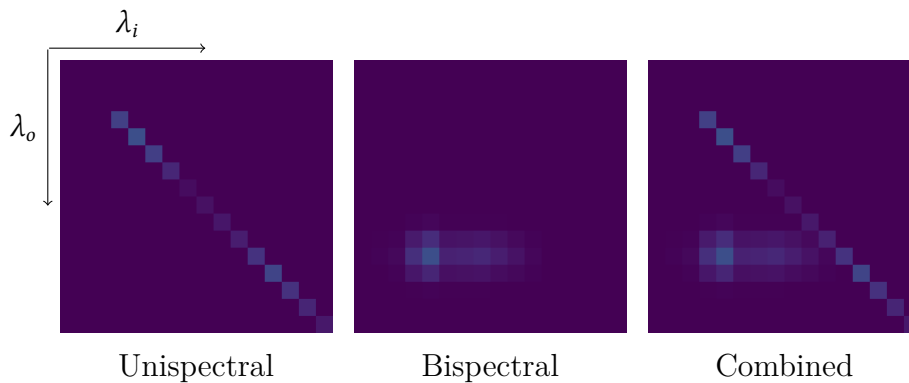


Figure 4.5: The two components and the combined estimated re-radiance matrix for "Pink Elastic".

shift in wavelength when being reflected. It is directly shaped by the illumination l_m . The second component, r_{bi} , is the re-radiated light resulting from fluorescence, the bi-spectral interaction, which *does* undergo a shift to a longer wavelength. Its spectrum is determined by the fluorophore's emission spectrum e and is not shaped by the illumination l_m other than scale via the integral s_m .

The goal is to use the camera response under multiple illuminations m to estimate the two radiance components r_{uni} and r_{bi} , for each illumination. From this and utilising the fact that our illuminations are narrow-band and Gaussian-like, we can estimate the uni-spectral reflectance $f_{uni}(\lambda)$ and the bi-spectral reflectance $f_{bi}(\lambda_i, \lambda_o)$ of the material that then gives us a *re-radiance map* $R(\lambda_i, \lambda_o)$, describing the full reflectance behaviour of the fluorescent material. See Figure 4.5 for an example radiance map with the two components.

Important to note is that the reflected radiance $r_m(\lambda_o)$ will be comprised of both specularly and diffusely reflected light. Our model (Equation 4.3) does not consider viewing or illuminating directions and thus cannot capture the highly directionally dependent specular reflections, the model is better suited to describe the diffusely reflected light with low directional dependency. Thus, to avoid affecting the estimated diffuse behaviour, the specular component should ideally be filtered out. This filtering can be done by illuminating with polarised light and having a polarising filter in front of the camera. However, we were concerned that the LEDs would not be powerful enough, and that using a pair of polarising filters (reducing the sensed light by quarter) would make capture too difficult. Instead, for this setup, the camera is placed facing straight onto the fabric and the illuminating light is placed 45° away from the camera's axis in order to avoid having the camera (the viewing direction) in or near the ideal angle of reflection, thus reducing the specular reflections. The solution is far from ideal and can be a source of error in the estimation.

4.3.1 Fitting Camera Radiance Response

There are two assumptions made for the fitting process:

- The fabric material has a smooth reflectance spectra. Thus, when illuminated by a light source with a narrow-band Gaussian-like emission spectrum, the uni-spectral reflected light r_{uni} can be well approximated by a Gaussian function.
- The bi-spectral re-radiated light from fluorescence r_{bi} , whose shape is determined by the the fluorophore's emission spectrum e can also be approximated with Gaussian a function (remember that r_{bi} is only scaled the illumination I).

The first assumption is justified by the fact that generally all natural surfaces have smooth spectra [14, 22]. The latter assumption however, does not always hold, such as with optical brighteners (as seen in the the *White Gaberdine* fabric Figure 2.3, where the fluorescent peak does not have a clean Gaussian shape). However, later we will see that it can provide plausible visual results when rendering, at least in a surface rendering setting.

Using our two assumptions, we can approximate the reflected radiance $r(\lambda_o)$ (Equation 4.3), by fitting a pair of Gaussian functions,

$$\tilde{r}(\lambda_o) = \underbrace{a_1 \exp\left(-\frac{(\lambda_o - \mu_1)^2}{2\sigma_1^2}\right)}_{\text{Gaussian 1}} + \underbrace{a_2 \exp\left(-\frac{(\lambda_o - \mu_2)^2}{2\sigma_2^2}\right)}_{\text{Gaussian 2}}. \quad (4.4)$$

The Gaussian parameters (a, σ, μ) will be adjusted by optimisation (numerical gradient descent) to minimise the difference between the approximated camera response \tilde{I}_k and the actual values captured by the camera I_k .

For each illumination m , the following minimisation problem needs to be solved,

$$\begin{aligned} \underset{\substack{(a_1, \mu_1, \sigma_1, \\ a_2, \mu_2, \sigma_2)}}{\text{argmin}} \quad & \sum_{k \in R, G, B} \left\| \frac{1}{b_m} I_k - \tilde{I}_k \right\|^2 \\ \text{subject to} \quad & a_1 > 0.1 \\ & |\mu_1 - \lambda_m| < \epsilon \\ & \mu_2 > (\lambda_m + \epsilon) \\ & \sigma_1 \in [\sigma_{1,\text{min}}, \sigma_{1,\text{max}}] \\ & \sigma_2 \in [\sigma_{2,\text{min}}, \sigma_{2,\text{max}}] \end{aligned} \quad (4.5)$$

Here I_k are the actual RGB values captured from the camera that have been scaled with the LED brightness factors b_m (described in the calibration step). \tilde{I}_k is the model described in Equation 4.2 that uses the gaussian approximation (Equation 4.4) for the reflected radiance, it models how our estimated reflectance translates to RGB values. The value ϵ is empirically set to 15 nm. The first two conditions force Gaussian 1 in Equation 4.4 to have at least some amplitude and to be around the normal reflectance peak of the light wavelength λ_m , thus estimating the uni-spectral reflected light r_{uni} . The third condition makes Gaussian 2 have a greater wavelength than Gaussian 1, estimating the fluorescence r_{bi} . Gaussian 2 has a greater wavelength since the fluorescent light will always re-emit at a longer wavelength than that of the incident light, a process referred to as *Stoke's shift*, meaning that the interaction can




Figure 4.6: Example of fitting process using synthetic RGB values. Two Gaussian functions are fitted to find uni-spectral radiance and bi-spectral radiance. Combined, the functions describe the total response under narrow-band illumination. The synthetic RGB values used for estimating the Gaussian functions are obtained by integrating the measured spectrum with the spectral sensitivity functions of the camera.

not increase the energy. The last two conditions constrain the variances to reasonable ranges. Values $\sigma_{1,\min}$, $\sigma_{1,\max}$ are determined by the variances we see in the LED illuminations (since uni-spectral reflected light will be shaped by the illumination). $\sigma_{2,\min}$, $\sigma_{2,\max}$ are empirically set to 10 and 40, respectively, to avoiding sharp spikes and overly flat estimates for the fluorescence. This minimisation problem is solved using numerical gradient descent.

Figure 4.6 shows a demonstration of this estimation method used with RGB values that have been obtained by integrating a measured spectrum of a fluorescent material illuminated by narrow-bandwidth light with the spectral sensitivity functions for a standard camera. The spectra of the illuminated fluorescent material was captured using a spectrometer. From this early test, we see that the estimation method (Equation 4.5) can detect and find the two peaks, giving us a mean-wavelength, a variance and an amplitude for the two components.

4.3.2 Estimating the Re-radiance Map

After fitting we now have multiple sets of parameters for Gaussian functions that estimate the uni-spectral r_{uni} and the bi-spectral r_{bi} reflected radiance from the material under each LED m . Our goal is now use these parameters to estimate the uni-spectral $f_{\text{uni}}(\lambda_o)$ and bi-spectral $f_{\text{bi}}(\lambda_i, \lambda_o)$ reflectances in order to then generate a *re-radiance map* that describes the fluorescent behaviour. This map can later be used in rendering.

Merge close peaks

Before estimating the reflectances, we merge uni-spectral and bispectral peaks that are close under the same illumination. If their means (μ_1 , μ_2) are sufficiently close (in our case 30 nm), they are combined ($a_2 = a_1 + a_2$, $\mu_2 = (\mu_1 + \mu_2)/2$, $\sigma_2 = \sqrt{\sigma_1^2 + \sigma_2^2}/2$) and the amplitude a_1 is zeroed. This is done in an attempt to keep the uni-spectral and bi-spectral effects separated.

Uni-spectral Reflectance

To estimate the uni-spectral reflectance $f_{\text{uni}}(\lambda_o)$, we find the piece-wise linear function that satisfies,

$$f_{\text{uni}}(\mu_1^m) = a_1^m \quad \text{for each illumination } m.$$

In other words, $f_{\text{uni}}(\lambda_o)$ is estimated by linear interpolation of the uni-spectral Gaussians' amplitudes a_1 and mean wavelengths μ_1 .

Bi-spectral Reflectance

Our parameters for Gaussian 2 (μ_2, σ_2, a_2) estimate the bi-spectral radiance $r_{\text{bi}}(\lambda_o)$ from the material for illumination m (Equation 4.3). For fluorescence, this radiance has the same spectral shape regardless of the in-going wavelengths (predicted by ‘‘Kasha’s rule’’). However, we do see this spectrum change in scale (not shape) with the illumination m . We wish to convert this scaling behaviour with respect to m to a function of incident wavelength λ_i .

Under narrow-band Gaussian-like illumination, we assume that the scaling integral s_m , samples the absorption spectra in a delta-like manner, i.e.

$$s_m = \int a(\lambda_i) l_m(\lambda_i) d\lambda_i \approx a(\lambda_m). \quad (4.6)$$

Our estimated scaling function $s(\lambda_i)$ is also found by linear interpolation, i.e the piece-wise linear function that satisfies,

$$s(\mu_1^m) = a_2^m \quad \text{for each illumination } m.$$

To determine the shape of the fluorescent emission, we take the median of all the estimated bi-spectral Gaussians' wavelength $\bar{\mu}_2$ and variance $\bar{\sigma}_2$. These two values will determine the central wavelength and spread of the fluorescent outgoing radiance. This means that we have the same outgoing bi-spectral spectra up to scale for all incident wavelengths. This gives us the following estimate for the bispectral component,

$$f_{\text{bi}}(\lambda_i, \lambda_o) = s(\lambda_i) \exp\left(-\frac{(\lambda_o - \bar{\mu}_2)^2}{2\bar{\sigma}_2^2}\right). \quad (4.7)$$

Re-radiance Matrix

Combining these two components, we get the full re-radiance map $R(\lambda_i, \lambda_o)$ that describes the radiance for in-going and outgoing wavelength that includes uni-spectral response as well as the bi-spectral part from the fluorescence,

$$R(\lambda_i, \lambda_o) = f_{\text{uni}}(\lambda_o)\delta(\lambda_i - \lambda_o) + f_{\text{bi}}(\lambda_i, \lambda_o), \quad (4.8)$$

where the delta function $\delta(\lambda_i - \lambda_o)$ ensures that the uni-spectral component only contributes when $\lambda_i = \lambda_o$.

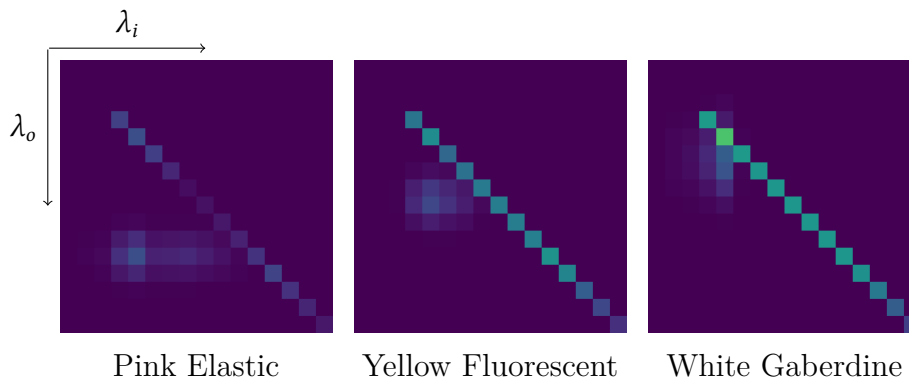


Figure 4.7: Estimated re-radiance matrices for the three sample fabrics.

The rendering system introduced later in this thesis renders using 16 spectral samples in the range 300 nm – 750 nm. There are rendering systems that allow for continuous sampling of spectra, however, in our case, these spectra must be made discrete. The Re-Radiance map (4.8) is directly sampled for each of 16 wavelengths on incident λ_i and outgoing λ_o wavelength to form a re-radiance matrix.

The estimated re-radiance matrices for the three fluorescent fabric samples are shown in Figure 4.7.

4.4 Evaluation

To evaluate the performance of the estimation method presented in this chapter, the resulting re-radiance matrices will be used to simulate the surface spectral response of the three samples under Xenon illumination (see Figure 4.8 for this light source’s illumination spectrum). The resulting response can then be compared to actual measured responses.

For evaluation, three types of measurements are used for comparison:

- *LED RGB Camera Estimate.* The estimated re-radiance matrices resulting from the method described in this chapter (Figure 4.7) are used to simulate the surface spectral response under Xenon illumination.
- *Measured Spectra.* A spectrometer is used to capture the actual spectral response from the sample fabric under an actual Xenon illumination. This serves as the *ground truth* measurement to which our estimated spectra should fit as closely as possible.
- *Direct Setup.* Here a re-radiance matrix is generated by a more traditional *direct setup*, consisting of the wide-spectrum Xenon light source, a set of nine optical band-pass filters and a spectrometer, bypassing the need of estimating the reflected uni-spectral and bi-spectral radiances via RGB values. The band-pass filters pass through light near 340, 370, 390, 400, 450, 500, 550, 600 and 650 nm. The wavelength spacing of these band-pass filters is more evenly

distributed than the distribution of the LEDs used in the LED setup, meaning that the area around 500 nm will be sampled better. Comparing with data from this direct method helps evaluate if the RGB Gaussian fitting or our particular distribution of LEDs limit our results.

To evaluate the performance of matching perceived colour, the estimated and actual measured spectral measurements are integrated using the CIE XYZ colour matching functions (Section 2.1).

The comparisons are shown in Figure 4.10 and Figure 4.9 and are discussed in the following subsections.

Simulated Radiance/Appearance Accuracy How well can the estimated re-radiance matrices simulate the fluorescent behaviour of the fabric under a certain illumination? For this evaluation, the estimated re-radiance matrix is multiplied by a vector representing the illumination spectrum of a Xenon light source (its illumination spectrum can be seen in Figure 4.8). For comparison, spectrometer measurements, from the actual fluorescent fabric under a real Xenon light source, is also shown.

The estimated and measured spectra for “Pink Elastic” and “Fluorescent Yellow” are shown in Figure 4.10 and in Figure 4.9 the same is shown for “White Gaberdine”. From the resulting spectra for all three fabrics, we see that the re-radiance matrix (estimated by the LED RGB camera setup) gives a spectrum very similar to the one measured on the physical sample. A promising result.

To evaluate how the difference in spectra affects perceived colour, sRGB representations have been calculated by integrating the spectra with the CIE XYZ 1931 Colour Matching functions and then converted to linear RGB. Finally, the RGB values are normalised and then converted to sRGB. These sRGB values are used to colour the rectangles in the figures. For best colour reproduction, it is recommended to read these figures on a display capable of sRGB. A digital colour-picker can also be used to retrieve the numerical values of the colours.

Colour wise, the match is not as promising. For “Pink Elastic” and “Fluorescent Yellow” in Figure 4.10, the sRGB representation of the actual measured spectra are more saturated than the colour retrieved from the estimates. Specifically for “Fluorescent Yellow”, the sRGB of the actual measurement is more yellow than the lime green from both the LED RGB estimate and the estimate from the direct setup. The fact that both matrices give colours that are off in a similar way suggests that the error is not due to a lack of LEDs near 500 nm in the LED RGB setup, since the direct setup has a more even distribution of band-pass filters. More discussion of this in Section 4.5.




Figure 4.8: Spectrum of light sources used when capturing photographs and rendering.

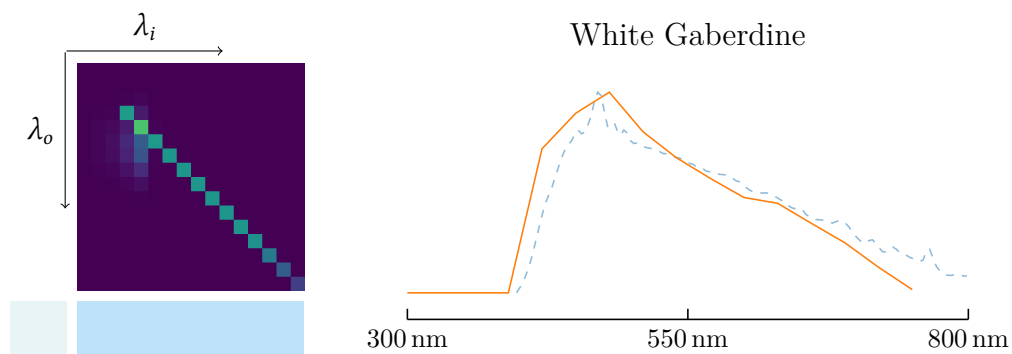


Figure 4.9: The estimated re-radiance matrix for “White Gaberdine” multiplied by a 16 sample approximation of the Xenon light source (see Figure 4.8) giving a simulated radiance under such illumination. The colour patch to the left is an sRGB representation of the actual measured spectral response (---). The right colour patch is a sRGB representation of the simulated response (—). The spectrometer measuring the radiance from this sample has a range of 400-1000 nm, causing the spectrum to be cut off for wavelengths lower than 400 nm. Therefore, for this sample, simulated radiance is zeroed for wavelengths lower than 400 nm.

Comparison of Re-radiance matrices under Xenon Illumination

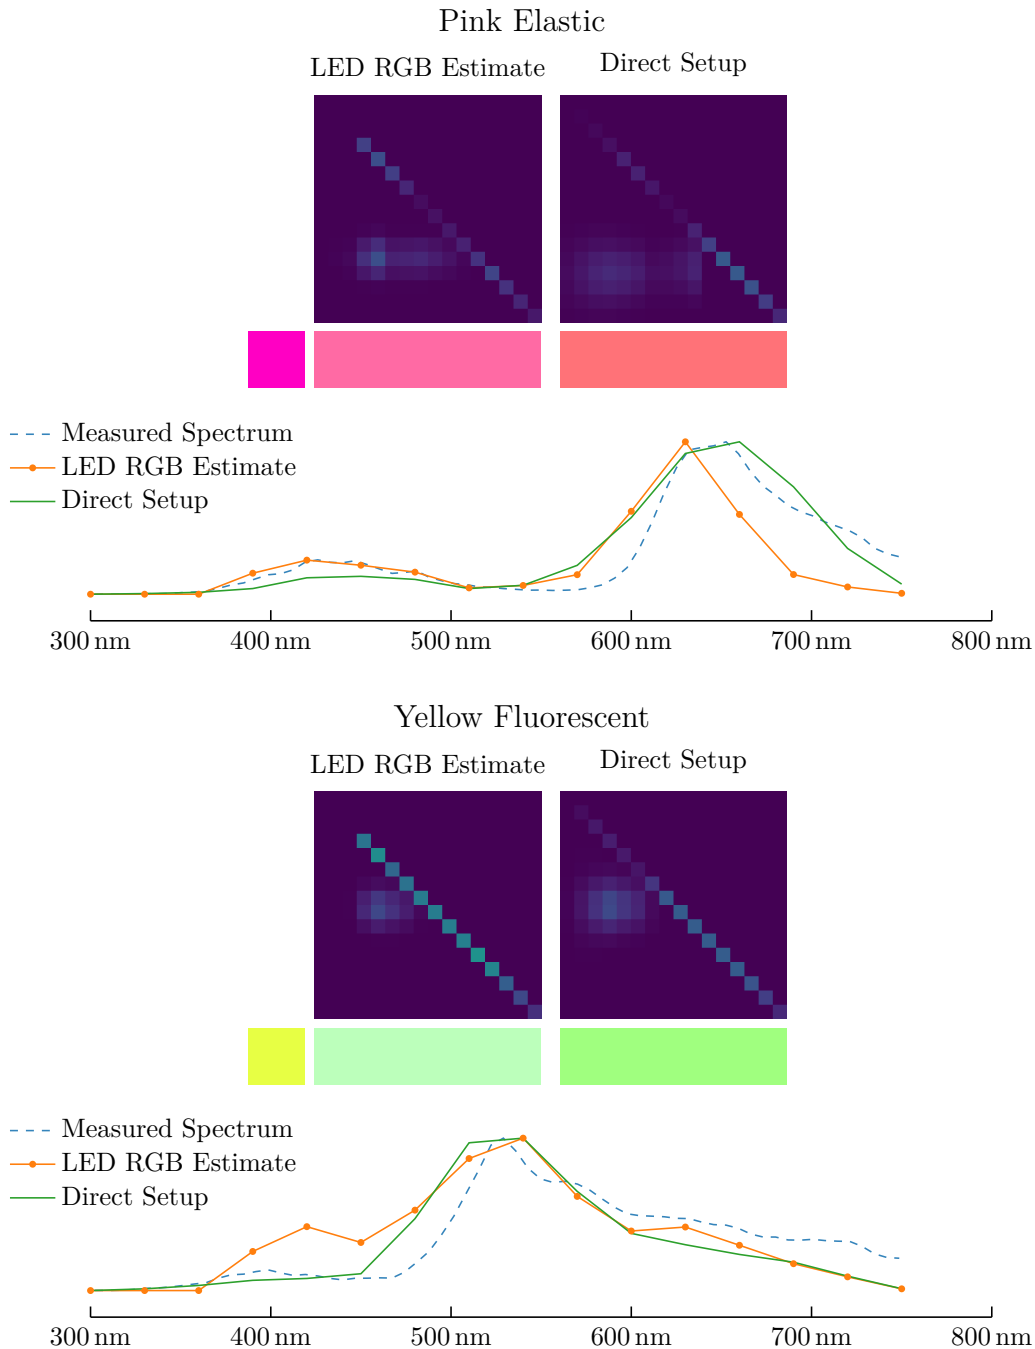


Figure 4.10: A comparison of the re-radiance matrices. The *LED RGB estimate* matrix has been estimated using the LED setup and an RGB camera. The *Direct Setup* matrix has been generated using optical filters and instead fits Gaussians directly to the measurements from the spectrometer, thereby skipping RGB fitting. The graph shows the simulated radiance under Xenon illumination resulting from each matrix. Underneath each matrix is also a colour patch illustrating the sRGB representation of the corresponding spectra. The leftmost patch is the colour of actual measured fabric radiance spectrum under xenon illumination. This figure helps evaluate the overall accuracy of the setup and the accuracy of the RGB estimator.

RGB fitting performance As a final evaluation, the RGB fitting performance is demonstrated. Estimated reflectance spectrum (consisting of two Gaussian functions) for each illumination is compared to the validation spectrometer data that was captured at the same time as the images were taken during the estimation process. This comparison is shown in Figure 4.11.




Figure 4.11: The plots show the mix of at most two Gaussian functions (—) that were found by minimising the RGB difference between the average RGB value from the HDR images and the estimated RGB value from inputting the Gaussian functions into our camera model (4.2). Above each plot is a circle with a colour representing the average RGB value from the HDR image. It is only this RGB value that is used to find the Gaussian functions. The measurements from the spectrometer (---) are only there for evaluation.

4.5 Conclusion

In this chapter, a setup than can capture the spectral behaviour of a fluorescent fabric was proposed. Based on the results in Chapter 3, we saw no special angular dependence that our setup would have to capture. The setup consisted of a flat fabric sample being illuminated by nine narrow-band LEDs and an RGB-camera. By fitting a model of the radiance to the RGB-values from the camera, the material’s *uni-spectral* and *bi-spectral* response to each illumination could be found and then used to generate a re-radiance matrix.

When using the re-radiance matrices to simulate the response under xenon light and comparing the spectrum to measurements we see a close match with small discrepancies. However, when interpreting these spectra as colour, the discrepancies result in a significant change in colour.

We also looked at matrices that come from a “Direct setup”, that circumvents the need to fit uni-spectral and bi-spectral responses via RGB values and instead measures directly using a spectrometer. Even with such a setup, the colour mismatch is significant. This is initially unexpected, since much of the potential error in the estimation is due to trying to fit those responses via RGB. There could be a number of reasons why. Perhaps the way the gaussian approximation and the re-radiance matrices are being generated limit the accuracy. Or perhaps it is more a matter of wavelength spacing of the different LEDs in the LED RGB Camera setup and the optical filters in the Direct Setup. Another potential issue is that the specular reflections, which we do not filter out in our setup (Section 4.3), affect our estimation of a diffuse effect.

The estimation method performs and shows promise, however for appearance capture it is not good enough, the colour mismatch is too significant. Further work is needed.

Further work

Narrow-band delta assumption In the estimation process, we assumed that the illumination band was so narrow that we could treat it as a delta function (Equation 4.6). This might be too great of an approximation and should instead be scaled by the shape of the illumination. The approximation then becomes,

$$s_m = \int a(\lambda_i)l_m(\lambda_i) d\lambda_i \approx \frac{a(\lambda_m)}{\int l_m(\lambda_i) d\lambda_i}.$$


LED wavelength spacing The nine LEDs used for the setup are spaced at different points in the visible spectrum (see Figure 4.2). Perhaps more LEDs are needed or that their spacing on the visible spectrum needs to be changed. A synthetic study could be performed that simulates the estimation process with a different number of

LEDs and different spectral illumination ranges. Such a study would indicate how many LEDs are needed and how good of a result we can expect from this.

Specular filtering for fabrics As mentioned in Section 4.3, in the proposed setup, no explicit filtering is done to prevent specularly reflected light from the fabric sample to reach the camera. This means that our uni-spectral estimate is due to both diffuse and specular contributions, which can cause a wrong estimate. Currently, we try to reduce specular reflections by making sure that the camera is not viewing the sample relative to the light in the ideal angle of reflection. To reduce the risk of this impacting the results, a better method for filtering out specular reflections should be considered. One approach would be to use a linear polariser in front of the light source and a 90° rotated polariser in front of the camera, thus filtering out light that simply reflects directly off the surface (specular reflections). However, the fact that we are dealing with fabrics complicates things, since the twisting nature of the fabric's yarns can cause specularly reflected light to change polarisation. Also, when dealing with low powered LEDs, the polarisation method reduces light intensity and the light into the camera, making the capture process more susceptible to noise and stray light.

Prior-knowledge of Fluorescence There are databases with many hundreds of fluorophores together with their emission and absorption spectra. This data could be used to inform the optimisation and estimation. Currently, we linearly interpolate the measured values to predict the scaling with incoming and outgoing wavelengths. With prior knowledge, the shape of this scaling can be fitted using basis functions via singular value decomposition (similar to [22]) or via Bayesian inference (as suggested in [19]).

Colour optimisation Small mismatches in spectra seem to be quite unforgiving when representing them as colours. If final colour accuracy is our goal, perhaps we should reformulate our optimisation problem to solve that instead. That is, instead of a bottom-up approach (trying to map the physical behaviour, by measuring directly), we instead approach the problem from the other end, optimise to get the correct colour under different sources of illumination and hope to find the correct matrix based on that.



5

Simulating Fluorescence

This chapter will discuss changes needed for a standard RGB path-tracing renderer to allow it to simulate effects such as fluorescence. Also, the measurements from the previous chapter are incorporated into the renderer. Simulations using these measurements result in rendered images that are compared to photographs of sample fabrics.

These goals necessitated changes to the rendering engine used during this project. The engine used is a proprietary RGB path-tracer similar to the open-source *PBRT*¹ renderer as its described in the textbook *Physically Based Rendering: From Theory to Implementation* [23] by Pharr, Jakob, and Humphreys. At the beginning of this project, the proprietary renderer did not support simulating effects such as fluorescence. The changes discussed in this chapter should be directly applicable to the official PBRT implementation.

5.1 Background

Rendering engines are designed to simulate the way light traverses a scene and how it interacts with objects. In general, the representation used for light in such rendering engines is a single ray of light with a set of three scalar values for red, green and blue. The reasoning for using only three scalars for this is given by the trichromatic theory, which says that we can create any colour perceivable by humans using a combination of three primary colours (see section 2.1). Using a trichromatic representation has the benefit of mapping well to digital image formats and display devices which all use RGB colour primaries or similar.

With *spectral rendering*, instead of having the ray of light and object's colours being

¹github.com/mmp/pbrt-v3

represented by three primary colours, light and surface albedo are represented by spectra. This means that all interactions are performed in terms of spectra instead of sets of RGB values. Simulating in spectral terms allows rendering systems to achieve a more accurate simulation. If the goal is to capture the appearance of an object and be able to correctly render it in any new setting, with any camera or illumination, spectral capture and rendering are very promising. In order to accurately simulate effects that are dependent on the light’s wavelength, such as diffraction, dispersion and Rayleigh scattering, spectral rendering is necessary. And for our intents, the ability to simulate in spectral terms is also a crucial step for correctly simulating bi-spectral effects such as fluorescence.

However, spectral rendering does come with practical challenges. It introduces a heavy cost in the additional computation and data manipulation needed during the simulations. There is also a lack of data to base the spectral simulations off. Authoring spectral content (such as textures) is difficult, since serious software for creating spectral images or textures is pretty much non-existent. See [25] for a deeper discussion on the barriers for the adoption of spectral rendering (and a proposed technique to overcome some of the drawbacks when using an RGB representation). To address some of these concerns, researchers have introduced performance efficient spectral sampling techniques [28] and methods to upsample RGB textures to equivalent estimated spectra [26, 16]. Because of these barriers, most professionally used rendering software has opted to stay with a trichromatic representation for the light that they simulate, and spectral rendering is mostly seen in non-production renderers aimed at research. Notable exceptions are *Maxwell* by Next Limit and more recently *Manuka* by Weta Digital [8].

Accurate simulation of bi-spectral effects has not seen adoption in many rendering systems and it is something that also is still mostly seen in research-oriented rendering systems. For instance, Manuka, the production rendering system with support for spectral rendering, does not support bi-spectral effects such as fluorescence. The reason for this being that “it is unclear how to make bi-directional path construction efficient” [8, 31:8], i.e it is hard to support bi-spectral effects while also allowing for bi-directional path-tracing, a set of techniques that gives other important benefits for efficiency and other effects [23, p. 16.3].

Even though direct simulation of bi-spectral effects are not present or possible in production rendering systems, a handful of research works have presented models and methods to simulate fluorescence. In 1995 Glassner [10] first extended the rendering model of Kajiya [17] (the Rendering equation) to allow bi-spectral effects. Introducing a new integral for the incoming wavelengths, it allowed the simulation of materials that depend on both incident and outgoing wavelengths, thus enabling fluorescence. Recently, a more accurate analytic appearance model by Jung et al. was presented [15, 1]. With such a model, the absorption and emission spectra can be simulated very accurately. However, when the goal is to recreate the appearance of a physical sample, it might be difficult to know what the absorption and emission spectra are for the fluorescent component. For such a use case, simulating from captured data might be the better approach. The only prevalent work regarding simulating fluorescence from captured sample data is [13] by Hullin et al., where

high-resolution BRRDFs are captured for fluorescent spherical homogeneous samples. The simulation is performed using re-radiance matrices, the same as those that have been introduced in Chapter 4. As we will see, such matrices are straight forward to implement into an existing rendering system and we will opt to follow this approach during this chapter.

However, before starting to implementing re-radiance into the renderer, spectral rendering needs to be supported.

5.2 Spectral Rendering

As stated previously, spectral rendering is necessary to accurately simulate bi-spectral effects such as fluorescence. How can this be introduced into a standard RGB path-tracer?

In Section 2.3, Chapter 2 the theoretical rendering equation, the problem we wish to solve each time we simulate a light ray in a rendering system, was given by

$$L_o(\omega_o, \lambda_o) = L_e(\omega_o, \lambda_o) + \int_{\Omega} f_r(\omega_i, \omega_o, \lambda_o) L_i(\omega_i, \lambda_o) (\omega_i \cdot \mathbf{n}) d\omega_i. \quad (5.1)$$

However, in practice, when the rendering system represents light using a trichromatic RGB representation, the actual equation being solved does not have a direct dependence on the outgoing wavelength λ_o , instead this dependence is expressed via the channel c ,

$$L_{o,c}(\omega_o) = L_{e,c}(\omega_o) + \int_{\Omega} f_{r,c}(\omega_i, \omega_o) L_i(\omega_i) (\omega_i \cdot \mathbf{n}) d\omega_i \text{ where } c \in \{R, G, B\}. \quad (5.2)$$

There are two common approaches to introduce spectral rendering.

Approach One: Per Wavelength Sampling

Emulating the original spectral version of the rendering equation (5.1) requires that BRDF f_r and light L_i models are able to be sampled using a wavelength λ_o . The framework in the rendering system needs changes to allow for this. Firstly, when determining a light-path to trace, the system must also determine the wavelength of that light-path. This means that the system that previously determined the light-path, by using the position of the pixel to be evaluated and the virtual camera properties, now also has to determine what wavelength this light-path should sample. Naively, the wavelength can be determined by randomly selecting one in the visible range and hope that the colour will converge when averaging enough samples. There exist more sophisticated wavelength sampling strategies that can significantly improve performance [28], but the general idea is the same.

The system will now simulate light-transport in spectral terms. But now during simulation, for each pixel, we will receive multiple samples for different wavelengths.

How should the system handle these samples? The samples will have to be converted to an RGB representation for a display. Thus, samples for pixel i and wavelength k will need to be integrated using an appropriate colour matching function (such as the CIE XYZ CMF introduced in section 2.1). This will result in a intensity for channel c that can be averaged with the contribution for other wavelengths. Mathematically we can write the intensity for channel c and pixel i for N spectral samples as,

$$I_{i,c} = \frac{1}{N} \sum_{k=0}^N \int L_o(\lambda_k) p_c(\lambda) d\lambda,$$

where $p_c(\lambda)$ is the spectral sensitivity function for channel c . Structuring it this way has the benefit of high spectral resolution, allowing accurate simulation of effects where this is important. The draw-back is that there is an extra dimension for the estimation method used (such as monte-carlo sampling), which will manifest itself as colour noise. Rendering a good image will require many more paths to be traced. Clever sampling strategies can help mitigate this.

Approach Two: Wavelength Channels

Another option to achieve spectral rendering is to keep the system evaluation structure that was used when using an RGB-triplet and thereby not introducing a new dimension into the estimation process. Here the RGB channels are replaced with N number of channels where each channel now represents a wavelength in the visible spectrum. This is a similar approach to what Pharr, Jakob, and Humphreys suggest in the PBRT renderer [23, Chapter 5.1]. A straight-forward approach is to evenly divide the visible spectrum N times and assign the wavelengths to the channels in that way. E.g. for channel k the corresponding wavelength is $\lambda_k = \lambda_{\min} + k/N(\lambda_{\max} - \lambda_{\min})$. In this way, the rendering system no longer has to determine a wavelength to assign each light-path to be traced, and so reduces colour noise. However, this can cause increased computing time and/or limit spectral resolution.

As with the first approach, the spectral samples will need to be integrated with a colour matching function before their contributions can be added to the image buffer. This is done in a similar manner, but can here be done in bulk, since each sample carries with it N channels. So, the intensity in image buffer for pixel i and RGB-channel $c \in \{R, G, B\}$ is given by,

$$I_{i,c} = \sum_{k=0}^N L_{k,o} p_c(\lambda_k).$$

It is this second approach that was used to introduce spectral rendering into the rendering system used during this thesis. Here $N = 16$ channels for spectral samples were used, resulting in increases in rendering times. Listing 5.1 shows how this method can be implemented in code. Instances of the class `SampledSpectrum` replace the RGB colours in the render system's data structures.

Listing 5.1: Implementation of a `Spectrum` class with coefficients for wavelengths in the visible range in the style suggested in PBRT [23, Chapter 5.1].

```
1 static const int sampledLambdaStart = 400;
2 static const int sampledLambdaEnd = 700;
3 static const int nSpectralSamples = 16;
4
5 template <int nSpectrumSamples>
6 class CoefficientSpectrum {
7     ...
8
9 protected:
10     float c[nSpectrumSamples];
11 };
12
13 class SampledSpectrum : public CoefficientSpectrum<nSpectralSamples> {
14     ...
15 };
```

5.3 Bi-spectral Reflection

With support for spectral rendering, the next task is to simulate re-radiance effects. The measurements of the fabrics in the previous chapter (Chapter 4), resulted in re-radiance matrices. Using matrices and combining them with the `SampledSpectrum` implementation discussed in the previous section allows for fluorescent simulation. The interaction can then be modelled by a straight-forward and efficient vector-matrix multiplication.

To facilitate this better, the fundamental light interaction system in the renderer has and been extended to support two new concepts (see Listing 5.2).

- `LightSpectrum` – 16 samples to represent a spectrum for describing the light of the light-rays being simulated and for describing the emission spectrum of light sources. An identical and inherited class to the `SampledSpectrum` class.
- `ReflectanceSpectrum` – Instances of this class describe how material surfaces affect incoming light. It can do this either by a normal 16 element vector that describes a traditional reflectance spectrum or with a 16 by 16 matrix describing a re-radiance matrix. Interaction between a reflectance-spectrum and a light-spectrum results in another light-spectrum and can be realised as a matrix-vector multiplication or as a component-wise vector multiplication, depending on if the interaction is bi-spectral or not.

With these changes in place, the rendering system is capable of simulating bi-spectral effects such as fluorescence, provided that we have re-radiance matrices. In Chapter 4, a method to capture such data from fabrics was introduced, which resulted in a set of re-radiance matrix. These matrices can now be directly used by the rendering system in a fundamental way.

Listing 5.2: Two sub-classes where the first represents spectra used for light-spectra and light-transport. The second represents surfaces' reflectance, which does not necessarily have re-radiance properties.

```

1 class LightSpectrum : public SampledSpectrum {
2     ...
3 };
4
5 class ReflectanceSpectrum {
6     LightSpectrum interact(LightSpectrum &in_spectrum);
7
8     protected:
9         MatrixNxN<nSpectrumSamples, nSpectrumSamples> m_reradiance_matrix;
10 };

```




Figure 5.1: Test image of a spectral render system capable of simulating re-radiance developed during this thesis. It shows a sphere with a fluorescent diffuse surface BRDF under three different illuminations. Under each image is an additional render of a Xrite ColorChecker Passport chart and the same fluorescent material. The colour chart was rendered using spectra for the colour patches available online. The resulting sRGB values for the rendered colour patch under D65 illumination are very close to the target sRGB values reported by Xrite, indicating that the spectral rendering works well.




Figure 5.2: The photo-setup used to capture comparison photos.

Figure 5.1 shows test images generated with an extended rendering-system capable of spectral and bi-spectral rendering.

5.4 Renders

The re-radiance matrices resulting from the estimation method introduced in Chapter 4 will now be used in the spectral rendering system.

Photographic Comparisons

A validation scene was used for photographing the samples under different illuminations, see Figure 5.2. The same scene is then digitally recreated for comparison. A Macbeth colour checker was used to colour correct and white balance the images for the D65 LED illumination, the renderer is white balanced for CIE D65 illumination.

In Figure 5.3 a mesh is rendered using a diffuse Lambertian shader with the albedo set to the estimated re-radiance matrices for the three fabrics. The mesh is rendered using illumination spectra for three light sources, a 6500K LED light, a halogen lamp and an ultra-violet light, see Figure 4.8 for a plot of their illumination spectra. The corresponding photographs and rendered images are joined so that the lower-left half shows the rendered image and the upper right half shows the photo.

From the photos, we see that the estimation for “Pink Elastic” has the best matches overall for the different illuminations. The estimates for other fabrics are close in appearance for D65 LED and Ultra-violet illumination, but is very different for the halogen bulb, where the photograph shows a much darker and brown response, suggesting that some part of the spectrum is being weighted too heavily. Without spectrometer measurements for these test cases, it is unfortunately hard to say definitively where the estimation is lacking.




Figure 5.3: The estimated re-radiance matrices used in a Lambertian surface model. The upper right half of each image is an actual photograph of the sample fabric under the same illumination. The photographs have been colour corrected and white balanced with a Macbeth chart under the D65 LED illumination. A: “Pink elastic”, B: “Yellow Fluorescent”, C: “White Gaberdine”.




Figure 5.4: Grid of meshes with a diffuse surface reflectance using the estimated re-radiance matrices for the three sample fabrics. The lower left half uses the full re-radiance matrix, the upper right half uses only the uni-spectral reflectance. The columns correspond to different illuminations. This demonstrates the important contribution that fluorescence can give to an object’s appearance.

Demonstrating the Effect of Fluorescence

In Figure 5.4 the same mesh is rendered using the same diffuse Lambertian shader, re-radiation matrices and illuminations as in our photographic comparisons. Here, however, top-half of these images the render has the bi-spectral component disabled, thereby effectively “switching off” the fluorescence. This illustrates the dramatic effect that fluorescence can have for appearance.

Isotropic Fluorescence in Fibre Scattering

Based on the experiments in Chapter 3, we saw behaviour that suggests that fluorescence is an isotropic effect, i.e. there is no apparent angular dependence on the fluorescent radiance at the yarn level and no special angular dependence that can’t be attributed to the geometry at the fibre level.

Thus, an attempt was made to naively extend the fibre scattering model by Khungurn et al. [18], to model the fluorescence effect as an additional isotropic scattering effect, using the uni-spectral component of the estimated matrix for the normal re-



Figure 5.5: The estimated re-radiance matrix for “Pink Elastic” used in a fibre scattering model, with the bispectral component as an isotropic lobe. The upper right half of each image is an actual photograph of the sample fabric under the same illumination. Here we note that for the re-radiance matrix that best fits the photographs in a Lambertian surface shader, does not work as an isotropic effect in a fibre scattering context.

flectance part of the fibre scattering model and the bi-spectral component for the new isotropic component.

This extension is implemented in a volumetric path tracer and is then used to render a garment using this model. In Figure 5.5 photographs are compared to renders using the extended model with this additional isotropic fluorescent scattering. The final appearance does not fit well to the actual look of the fabrics.

Although the capture method introduced in Chapter 4 estimates the re-radiance matrices in a way that can give close results in surface rendering (Figure 5.5), we cannot use them directly in this fibre scattering model. This discrepancy is not necessarily due to our isotropic scattering assumption. Problems with this assumption would manifest itself as discrepancies in the shading and not so much in the visual colour. The error could be due to a scaling issue in how to translate surface-level measurements to volumetric scattering parameters. Further study is needed in order to know how to apply these surface measurements into the a volumetric fibre scattering context.

5.5 Conclusion

In this chapter, we outlined how a rendering system can be extended to support spectral rendering and to support bi-spectral reflectance, two necessary steps for simulating fluorescence.

The estimated re-radiance matrices from Chapter 4 for the sample fabrics were then used for rendering. We saw that using the re-radiance matrices in a surface material (diffuse Lambertian) gave results that are near in appearance of photographs of the same samples under different illuminations, except for Halogen.

Finally, a fibre scattering model was extended to include fluorescence as an isotropic scattering effect, motivated by the findings in Chapter 3, where results suggest that

fluorescence is an isotropic effect. However, the renders using the extended scattering model show significant differences in the colour from the photos. We can also note that the renders using the surface material, are much closer to the photographs. I suspect that this discrepancy is not due to a false assumption about isotropic fluorescence but instead is a problem with the scaling between the scattering of the uni-spectral and bi-spectral components and that scaling is different when compared with surface rendering.

Matching the appearance of fabrics using such a scattering model is more difficult. Focusing only on achieving a closer match colour wise, I suspect that better results can be had if the model parameters are optimised to match images of the fabrics under multiple illuminations, a similar conclusion was reached at the end of Chapter 4.



6

Summary

From early on in the work for this thesis, the main goal has been to look at the challenge of capturing and simulating fluorescence in fabrics. During the process of studying these challenges and coming up with possible solutions, many questions appeared. One question that arose from this, and has resurfaces in every part the project, is the one of potential directional dependence of fluorescence in fabrics. Knowing this helps in how to properly capture the phenomenon and later, how to properly simulate it.

A short summary of these main results and points presented in this thesis follows.

Behaviour From experiments on the fabric samples at hand, at the yarn-level and at the fabric-level, no detected angular dependence was found that can't be attributed to the geometry of the fabric. From these results the conclusion is that fluorescent light from a yarn or fabric has no special angular dependence and appears to be an isotropic volumetric effect at the yarn level and a Lambertian diffuse behaviour at the fabric-level, at least for the samples measured. A definitive answer on the directional behaviour of fluorescence in fibres and yarn can not be given from the experiments performed here. Investigation and experiments towards this end would be interesting to pursue further.

Capture The findings about the behaviour of fluorescence helped inform the capture process. It allows for a practical setup by only requiring the measurement of its bi-spectral response, directional measurements are not required. Furthermore, the setup was made practical and more accessible than other appearance capture methods for fluorescence, by using a set of nine LEDs and a common RGB-camera, thus requiring no lab equipment. The proposed setup and method successfully captures the fluorescent behaviour from fabrics. The data from this setup can be used

to simulate the spectral response of these samples under new illumination. However, when considering colour accuracy, the method leaves room for improvement. An approach for improvement that would be interesting to pursue is using prior knowledge about fluorescent compounds to improve the estimation.

Simulation Finally, given that fluorescence shows little angular dependence, a volumetric fibre scattering model was extended to support fluorescence by treating fluorescence as isotropic scattering. However, using the fluorescent measurements from the capture setup did not translate well for use in a fibre scattering model. Comparisons with photographs and renders show a big difference in colour and brightness. Using the same measurements in surface reflectance model gave better results. Perhaps there is some scaling that needs to be considered when applying the measurements made to a fibre scattering model. Previous work with fibre scattering models have used an optimisation process to match the appearance of fabrics and correctly scale the colours. It would be interesting to extend such an approach to also support fluorescence.

Conclusion

The three main areas of concern for this thesis, behaviour, capture and simulation, have each been interesting to study and offer many interesting phenomena and problems to study and solve. The work for this thesis has consisted of identifying these phenomenon and problems, in addition to the actual experiments, analysis and development towards a method to capture and simulate fluorescence.

The methods and experiments presented are an attempt at answering questions about the directional of fluorescence and to solve the problem of capture and simulation. These attempts give promising results and insight, but, there is more work that can be done: a more thorough study of the potential angular dependence of fluorescence for example. The capture method for appearance capture can be improved further or approached from another direction. There is also need for more focused measurements to more definitively determine how fluorescence behaves in a fibre-scattering context.

The work in this thesis gives insight into three related but different areas, each with their own very intriguing problems and questions. The presented experiments, I think, show encouraging results but also give new interesting paths and questions to explore further.

Bibliography

- [1] Marwan Abdellah et al. “A Physically Plausible Model for Rendering Highly Scattering Fluorescent Participating Media”. In: (June 2017).
- [2] Carlos Aliaga et al. “An Appearance Model for Textile Fibers”. In: *Computer Graphics Forum* 36.4 (2017), pp. 35–45. DOI: [10.1111/cgf.13222](https://doi.org/10.1111/cgf.13222).
- [3] Applied Sensor Group of the Institute of Analytical Chemistry at the Graz University of Technology. *Fluorophores.org*. URL: www.fluorophores.tugraz.at (visited on 09/10/2019).
- [4] Xun Cao et al. “A prism-mask system for multispectral video acquisition”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 33.12 (2011), pp. 2423–2435. DOI: [10.1109/TPAMI.2011.80](https://doi.org/10.1109/TPAMI.2011.80).
- [5] Carlos Castillo, Jorge López-Moreno, and Carlos Aliaga. “Recent advances in fabric appearance reproduction”. In: *Computers & Graphics* (July 2019). DOI: [10.1016/j.cag.2019.07.007](https://doi.org/10.1016/j.cag.2019.07.007).
- [6] Jonathan Dupuy and Wenzel Jakob. “An adaptive parameterization for efficient material acquisition and rendering”. In: *SIGGRAPH Asia 2018 Technical Papers, SIGGRAPH Asia 2018* 37.6 (2018). DOI: [10.1145/3272127.3275059](https://doi.org/10.1145/3272127.3275059).
- [7] Hugh S. Fairman, Michael H. Brill, and Henry Hemmendinger. “How the CIE 1931 color-matching functions were derived from Wright-Guild data”. In: *Color Research and Application* 22.1 (1997), pp. 11–23. DOI: [10.1002/\(SICI\)1520-6378\(199702\)22:1<11::AID-COL4>3.0.CO;2-7](https://doi.org/10.1002/(SICI)1520-6378(199702)22:1<11::AID-COL4>3.0.CO;2-7).
- [8] Luca Fascione et al. “Manuka”. In: *ACM Transactions on Graphics* 37.3 (2018), pp. 1–18. DOI: [10.1145/3182161](https://doi.org/10.1145/3182161).
- [9] Ying Fu et al. “Reflectance and Fluorescence Spectral Recovery via Actively Lit RGB Images”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 38.7 (2016), pp. 1313–1326. DOI: [10.1109/TPAMI.2015.2439270](https://doi.org/10.1109/TPAMI.2015.2439270).
- [10] Andrew S Glassner. “A Model for Fluorescence and Phosphorescence”. In: (1995).
- [11] Mayank Goel et al. “HyperCam”. In: *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing - UbiComp '15*. New York, New York, USA: ACM Press, 2015, pp. 145–156. DOI: [10.1145/2750858.2804282](https://doi.org/10.1145/2750858.2804282).
- [12] Ralf Habel, Michael Kudenov, and Michael Wimmer. “Practical spectral photography”. In: *Computer Graphics Forum* 31.2 (2012), pp. 449–458. DOI: [10.1111/j.1467-8659.2012.03024.x](https://doi.org/10.1111/j.1467-8659.2012.03024.x).
- [13] Matthias B. Hullin et al. “Acquisition and analysis of bispectral bidirectional reflectance and reradiation distribution functions”. In: *ACM Transactions on Graphics* 29.4 (2010), p. 1. DOI: [10.1145/1833351.1778834](https://doi.org/10.1145/1833351.1778834).
- [14] Xiangyang Ji et al. “Separating reflective and fluorescent components for dynamic scenes”. In: *Optics Communications* 404.711 (2017), pp. 11–17. DOI: [10.1016/j.optcom.2017.05.075](https://doi.org/10.1016/j.optcom.2017.05.075).

-
- [15] Alisa Jung et al. “A Simple Diffuse Fluorescent BBRRDF Model”. In: *Workshop on Material Appearance Modeling*. Ed. by Reinhard Klein and Holly Rushmeier. The Eurographics Association, 2018. DOI: [10.2312/mam.20181193](https://doi.org/10.2312/mam.20181193).
- [16] A Jung et al. “Wide Gamut Spectral Upsampling with Fluorescence”. In: *Computer Graphics Forum* 38.4 (July 2019), pp. 87–96. DOI: [10.1111/cgf.13773](https://doi.org/10.1111/cgf.13773).
- [17] James T. Kajiya. “The Rendering Equation”. In: *SIGGRAPH Comput. Graph.* 20.4 (Aug. 1986), pp. 143–150. DOI: [10.1145/15886.15902](https://doi.org/10.1145/15886.15902).
- [18] Pramook Khungurn et al. “Matching Real Fabrics with Micro-Appearance Models”. In: *ACM Transactions on Graphics* 35.1 (Dec. 2015), pp. 1–26. DOI: [10.1145/2818648](https://doi.org/10.1145/2818648).
- [19] Koji Koyamatsu et al. “Reflective and Fluorescent Separation under Narrow-Band Illumination”. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2019, pp. 7577–7585.
- [20] Stephen R. Marschner et al. “Light scattering from human hair fibers”. In: *ACM SIGGRAPH 2003 Papers, SIGGRAPH '03* (2003), pp. 780–791. DOI: [10.1145/1201775.882345](https://doi.org/10.1145/1201775.882345).
- [21] Wojciech Matusik et al. “A data-driven reflectance model”. In: *ACM Transactions on Graphics* (2003). DOI: [10.1145/882262.882343](https://doi.org/10.1145/882262.882343).
- [22] Jong Il Park et al. “Multispectral imaging using multiplexed illumination”. In: *Proceedings of the IEEE International Conference on Computer Vision* (2007). DOI: [10.1109/ICCV.2007.4409090](https://doi.org/10.1109/ICCV.2007.4409090).
- [23] Matt Pharr, Wenzel Jakob, and Greg Humphreys. *Physically Based Rendering: From Theory to Implementation*. 3rd. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2016.
- [24] Gaurav Sharma. *Digital Color Handbook*. CRC Press, 2003.
- [25] Greg Ward and Elena Eydelberg-Vileshin. “Picture Perfect RGB Rendering Using Spectral Prefiltering and Sharp Color Primaries”. In: *Proceedings of the 13th Eurographics Workshop on Rendering*. Pisa, Italy: Eurographics Association, 2002, pp. 117–124.
- [26] Jakob Wenzel and Johannes Hanika. “A Low-Dimensional Function Space for Efficient Spectral Upsampling”. In: *Computer Graphics Forum (Proceedings of Eurographics)* 38.2 (2019).
- [27] Alexander Wilkie et al. “A reflectance model for diffuse fluorescent surfaces”. In: *Proceedings of the 4th international conference on Computer graphics and interactive techniques in Australasia and Southeast Asia - GRAPHITE '06*. New York, New York, USA: ACM Press, 2006, p. 321. DOI: [10.1145/1174429.1174484](https://doi.org/10.1145/1174429.1174484).
- [28] A. Wilkie et al. “Hero Wavelength Spectral Sampling”. In: *Computer Graphics Forum* 33.4 (July 2014), pp. 123–131. DOI: [10.1111/cgf.12419](https://doi.org/10.1111/cgf.12419).
- [29] Cherry Zhang and Imari Sato. “Image-based separation of reflective and fluorescent components using illumination variant and invariant color”. In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35.12 (2013), pp. 2866–2877. DOI: [10.1109/TPAMI.2012.255](https://doi.org/10.1109/TPAMI.2012.255).