UNIVERSITY OF EXETER

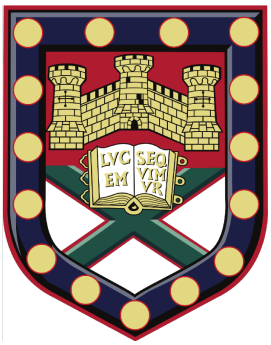

DOCTORAL THESIS

---

# Control of White Light Emission for Illumination and Imaging

---


*Author:*
Harry Penketh

*Supervisors:*
Prof. Jacopo Bertolotti
Prof. William L. Barnes


*Submitted by Harry Penketh to the University of Exeter as a thesis for the degree of Doctor of Philosophy in Physics, September 2020*

(Signed) H.Penketh

# Control of White Light Emission for Illumination and Imaging

This thesis is available for Library use on the understanding that it is copyright material and that no quotation from the thesis may be published without proper acknowledgement.

I certify that all material in this thesis which is not my own work has been identified and that no material has previously been submitted and approved for the award of a degree by this or any other University.

This work was funded by Dyson Technology Limited. I acknowledge support from the Engineering and Physical Sciences Research Council (EPSRC) of the United Kingdom, via the EPSRC Centre for Doctoral Training in Metamaterials (Grant No. EP/L015331/1).

(Signed) H.Penketh

# Abstract

The bulk of the work in this thesis addresses a well defined problem in the field of white solid-state lighting. The problem was defined by this project's industry sponsors and consequentially this thesis is to be covered by an embargo period in which it is not publicly accessible, to protect commercial sensitivities and intellectual property. The title of this thesis and this abstract are not protected by these restrictions and so must be selectively descriptive.

There are two bodies of work that can be freely described here. The first, presented in Chapter 2 concerns the performance of parabolic antennas as collimators when their size is reduced to of the order of the wavelength of emission. Using finite-element modelling we calculate the dependence of the amount of power directed into a 20° half-angle cone on the emitter's position within the paraboloid and compare with results obtained using geometrical optics.

It is demonstrated that due to variations in the local density of optical states, changing the characteristic size of the reflector within the range from 0.5 to 1.5 times the emission wavelength has a strong bearing on the optimum emitter position, a position that does not in general coincide with the paraboloid's focus. We calculate that the optimal antenna size and emitter position allow for the maximum directed power to exceed that obtained in the geometrical optics regime.

In the other section that can be described here, Chapter 6, we present a novel technique in the field of ghost imaging. In our method we show that one can modify the illumination basis patterns to perform common image post-processing steps directly during the reconstruction process. This can remove the amplification of detector noise caused by post-processing with a filter and yields favourable noise spatial statistics. This technique can be applied to any operation which could be written as a matrix multiplication with the image, which includes the application of image filters by convolution.

# Acknowledgments

I'm going to try and keep these acknowledgements brief, so I will not be presenting an exhaustive list of all those I have to thank for managing to get to this stage of my life / education. This brevity is motivated primarily by the imminent submission deadline.

Firstly I wish to thank my supervisors Jacopo and Bill. To this day I cannot think of a better combination of supervision styles and personalities that could have been applied to this slightly unusual research project. I am sincerely grateful for your time, patience and wisdom over these last four years.

Of my 'peers' special thanks needs to go to Kieran Walsh and Craig Tollerton. Kieran has been an asset as my consultant chemist, always willing to take a quick look at some strange samples on a piece of equipment on which I wasn't trained (all of them), and has been invaluable as a friend. Craig, watching you one year ahead of me in this 'process' deal with the many challenges your project presented, with varying degrees of success, has provided some twisted combination of motivation and entertainment over the years.

To my partner Carys, I can think of no one I would rather spend six months locked up with (due to COVID-19) whilst struggling to write a thesis. You will always have my thanks for everything you do for me.

# Contents


**1 Introduction** **1**

1.1 Aim and Outline of this Thesis . . . . . 1
1.2 What is White? (Some Useful Colourimetry Concepts) . . . . . 3
1.3 Solid-state Lighting . . . . . 6
1.3.1 Light-Emitting Diodes (LEDs) . . . . . 6
1.3.2 Laser diodes (LDs) . . . . . 7
1.3.3 White Light-emitting Diodes / Laser Diodes . . . . . 8
1.3.4 YAG:$Ce^{3+}$ . . . . . 10
1.4 The Problem in Detail . . . . . 12
1.4.1 Functional Requirements . . . . . 13
1.5 Ghost Imaging . . . . . 16
1.6 Summary . . . . . 17

**2 Wavelength-scale Parabolic Reflectors** **18**

2.1 Motivation . . . . . 18
2.2 The Geometrical Optics Parabolic Reflector . . . . . 20
2.3 The Oscillating Electric Dipole . . . . . 22
2.4 Fluorescence Near a Metallic Interface . . . . . 23
2.5 The Wavelength-scale Parabolic Reflector . . . . . 28
2.5.1 In the Context of Other Work . . . . . 33
2.5.2 Conclusions . . . . . 33
2.6 Methods . . . . . 34
2.6.1 A Note on Meshes . . . . . 34
2.6.2 Dipole Above a Mirror COMSOL Model . . . . . 35
2.6.3 Geometrical Optics Parabolic Reflector Method . . . . . 36
2.6.4 Wavelength-scale Parabolic Reflector Method . . . . . 37
2.6.5 A Note on Radiation 'Through' the Metal . . . . . 40

**3 Choosing the Right Approach** **43**

3.1 Assessing Strategy Feasibility . . . . . 43
3.2 Minimum Amount of Phosphor Calculation . . . . . 46
3.3 Phosphor Pump Absorption . . . . . 48
3.4 Conclusions . . . . . 49

**4 Fluorescent Fibres** **51**
4.1 Fundamentals and Motivation . . . 51
4.1.1 Fluorescent Fibres vs the Functional Requirements . . . 53
4.2 The Trapping Fraction . . . 55
4.2.1 3D Fibre Trapping Fraction . . . 56
4.2.2 Square Cross-section Trapping Fraction . . . 60
4.2.3 Discussion . . . 61
4.3 Outcoupling and Directing Fibre Emission . . . 62
4.3.1 Surface roughness . . . 63
4.3.2 Hemispherical Lenses . . . 64
4.3.3 Ball Lenses . . . 66
4.3.4 Modified Ball Lenses . . . 68
4.4 Simulating Losses in Fluorescent Guides . . . 71
4.5 Wavelength-scale Fibres . . . 73
4.5.1 Motivation . . . 73
4.5.2 Method - Modified Born Series . . . 74
4.5.3 Results and Discussion . . . 76
4.5.4 Simulation Details . . . 80
4.6 Prototype YAG:$Ce^{3+}$ Light Guides . . . 81
4.6.1 Evolution of Prototypes, Sample Descriptions . . . 81
4.7 Demonstrating a Radially Dependent Trapping Fraction Experimentally . . . 86
4.8 Prototype Efficiency and CCT Results . . . 94
4.8.1 Efficiency and CCT Calculation Methods . . . 94
4.8.2 Results and Discussion . . . 96
4.8.3 Square vs Circular Cross-section . . . 99
4.9 Conclusions . . . 100
4.9.1 Evaluation Against Functional Requirements . . . 100
4.9.2 In the Context of Other Work . . . 104
4.9.3 Summary . . . 105

**5 Fibre Simulation Details** **106**
5.1 A 3D Fibre Simulation . . . 106
5.2 3D Square Cross-section Guide Simulation . . . 109
5.3 Increasing Path Lengths Simulation . . . 110
5.4 A Lossy 3D Fibre Simulation . . . 113
5.5 2D Fibre Cross-section Simulation . . . 116

5.6 2D Hemispherical Lenses Simulation . . . . . 118
5.7 3D Hemispherical Lenses Simulation . . . . . 121
5.8 Modified Ball Lens Simulation . . . . . 123

**6 Ghost Imaging** **128**
6.1 Conventional Ghost Imaging . . . . . 128
6.1.1 Ghost Imaging vs Single Pixel Cameras . . . . . 130
6.1.2 Choice of Basis . . . . . 130
6.1.3 Experimental Method . . . . . 133
6.2 Ghost Imaging in Alternative Bases . . . . . 135
6.2.1 Noise Spatial Correlations . . . . . 147
6.2.2 Compressed Sensing . . . . . 149
6.3 In the Context of Other Work . . . . . 152
6.4 Conclusions . . . . . 153

**7 Summary and Outlook** **154**

**8 Publications and Intentions** **156**

**References** **156**

# 1 Introduction

## 1.1 Aim and Outline of this Thesis

This PhD project was sponsored by Dyson Technology Limited and as a result had a well defined primary objective since its inception. The essence of the task is captured in the following brief objective statement:

**Objective:** *Develop a highly efficient phosphor component that is capable of converting blue pump-light into white-light emission whilst retaining the incident spatial intensity pattern of the pump-beam and emitting into a narrow solid-angle.*

To those who are familiar with the methods and challenges associated with modern solid-state white lighting, particularly for high power and projection applications, this statement may be self-explanatory. We hope having this description in mind will in any case prove beneficial whilst we discuss some of the background concepts necessary to understand the problem in its full glory. This is the subject of this first chapter, defining the problem in full, and introducing the associated concepts at a level that enables understanding of the remainder of this thesis, without becoming a thesis in itself. Chapters 2 - 5 are on the subject of our research to meet this objective, with Chapter 5 containing the details of the simulations used in Chapter 4. Chapter 6 focuses on one possible application of the developed solution, presenting a novel technique in the field of computational ghost imaging.

In Chapter 2 we explore through finite element method simulations the ability of parabolic reflectors to provide directional light emission when their size has been reduced to a level comparable with the emission wavelength. The motivation for this work was exploring the minimum element size limits for a parabolic antenna array to meet the project brief, but has broader application in the area of controlling nanoscale sources of light. We explore a range of paraboloid sizes for which the optimal position for an emitter can be significantly displaced from the geometrical focus.

In Chapter 3 we briefly discuss some of the rationale behind our assessment of the likelihood of various techniques being able to meet the project brief. After calculating that the pump light absorption of typical phosphors deployed on the nanoscale would be inadequate, we conclude that our most promising candidate solution is the fluorescent optical fibre.

In Chapter 4 we evaluate our chosen strategy, fluorescent optical fibres. The optical effi-

ciency of this approach depends heavily on the trapping fraction of fluorescence generated within the fibre. We find that a simplified analytical treatment neglecting skew rays could cause a dramatic underestimate of the trapping fraction. Through simulations we quantify the effect of skew rays on pristine YAG:$Ce^{3+}$ fibres and explore modifications to enhance the efficiency and directivity of the white-light output. We also assess the performance of wavelength-scale fibres through simulation and conclude that there is no clear advantage to warrant design at this scale.

Experimentally we focus on demonstrating the impact of skew rays through the pump position dependent trapping fraction predicted theoretically. From our measurements we infer significant variations in the trapping fraction in cylindrical monocrystalline YAG:$Ce^{3+}$ rods of approximately half the magnitude predicted for pristine fibres. Finally, we quantify the performance of our phosphor rod prototypes and evaluate the strengths of fluorescent fibres as a solution to the project objective.

Chapter 5 contains details of all of the numerical (COMSOL) simulations that were used in Chapter 4.

In Chapter 6, having satisfied the project objective in Chapter 4, we move on to computational ghost imaging techniques. These imaging methods, which rely on structured illumination of the object to be imaged, can be considered a candidate application of the novel white light source developed in the preceding work, even though they are demonstrated here with traditional sources. We present a novel 'trick' for computational ghost imaging, in which the illumination basis can be modified to perform image post-processing operations directly. This technique is valid for any operation which can be written as a convolution of the ghost image with another matrix.

We choose a simple edge-detection kernel as an example, but provide the basic tools needed to evaluate the merits of this approach for an arbitrary kernel. Finally, we test the performance of this technique when used in conjunction with compressed sensing.

## 1.2 What is White? (Some Useful Colourimetry Concepts)

In simplest terms, white is the lightest version of the absence of colour, with pure black forming its darkest counterpart. This definition enables one to describe a perfect white object with relative ease as an object that diffusely reflects all visible light incident upon it, with zero absorption. Zero absorption means that no spectrally varying absorption has imparted a colour, or hue to the light reflected from the object. The requirement for diffuse reflection distinguishes a white object from a perfect mirror.

The corresponding simple definition for a white light source is dictated by the human visual system and the perception of colour [1]. A white light is one that stimulates the three types of cone cells found in a typical human eye by equal amounts. The three cone cell types are the basis for colour vision, having differing spectral sensitivities that loosely correspond to the colours blue, green and red. We will use this opportunity to translate this definition into a more convenient spatial representation of colour and whiteness, the widely adopted CIE 1931 chromaticity space.

The perceived colour of a light source, which is remarkably consistent among those possessing normal colour vision, may be described mathematically by a weighted combination of three primary colours needed to match it in appearance [2]. These so-called primary colours, traditionally red, green and blue (RGB), must satisfy the condition that none of them can be described as a combination of the other two [1]. One problem which arises from this choice of primaries is that some colours cannot experimentally be matched by positive contributions from each, and instead require a negative contribution, realised experimentally by adding some of the primary light to the colour one is trying to match. This generates some mathematical and conceptual inconveniences.

In 1931 the Commission Internationale de l'éclairage (CIE) introduced their XYZ colour space, which uses 'imaginary' primaries to provide a more convenient basis for colourimetry. The CIE XYZ standard observer colour matching functions can be seen in Fig.1.1. These functions derived from the original RGB experimental data offer two main quality of life improvements. First, they require no negative values, any real colour can be matched by positive amounts of these primaries. Second, the amount of primary $\bar{y}(\lambda)$ needed to achieve colour-matching, denoted $Y$, now specifies the luminosity (lightness) of the light on its own [2]. $\bar{y}(\lambda)$ is the photopic luminosity curve describing the spectral sensitivity of perceived brightness for the average human observer. 'Photopic' here distinguishes from the modified scotopic sensitivity curve experienced at low light levels where vision is dominated by rods not cones [1].

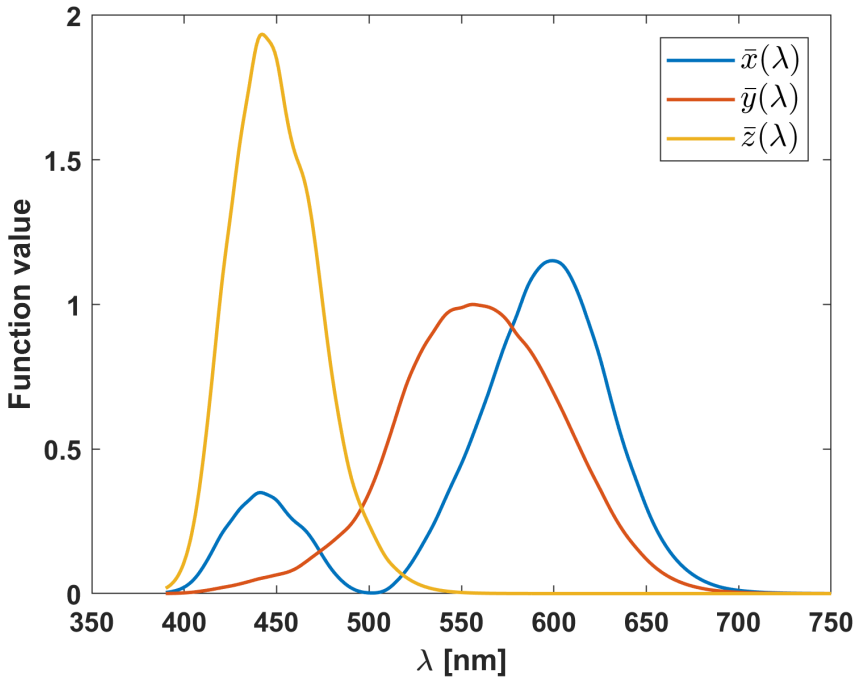


**Figure 1.1:** CIE XYZ standard observer color matching functions. Data from the Judd-Vos modified CIE 2-deg photopic luminosity curve (1978) [3]

The tristimulus values in the CIE XYZ system as defined for a given spectral power distribution ($P(\lambda)$) are defined as

$$X = \int_{380\text{ nm}}^{750\text{ nm}} P(\lambda)\bar{x}(\lambda)\mathrm{d}\lambda, \tag{1.1}$$

$$Y = \int_{380\text{ nm}}^{750\text{ nm}} P(\lambda)\bar{y}(\lambda)\mathrm{d}\lambda, \tag{1.2}$$

$$Z = \int_{380\text{ nm}}^{750\text{ nm}} P(\lambda)\bar{z}(\lambda)\mathrm{d}\lambda, \tag{1.3}$$

where $\bar{x}(\lambda), \bar{y}(\lambda)$ and $\bar{z}(\lambda)$ are as shown in Fig.1.1 and available from [3]. From these tristimulus values a convenient representation arises from the definition of normalised tristimulus coefficients, also known as chromaticity coordinates

$$x = \frac{X}{X+Y+Z}, \tag{1.4}$$

$$y = \frac{Y}{X+Y+Z}, \tag{1.5}$$

$$z = \frac{Z}{X+Y+Z}. \tag{1.6}$$

One can uniquely define the colour of a light with two chromaticity coordinates (e.g. $x, y$) and the $Y$ tristimulus value providing the luminosity. Here we have described the CIE $xyY$ colour space. There is no need to specify the $z$ coordinate as $x + y + z = 1$. Finally, if one is concerned only with colour, the $Y$ value may be excluded and the light is defined by a position in a 2D colour space, as in Fig.1.2, which is one realisation of the standard chromaticity

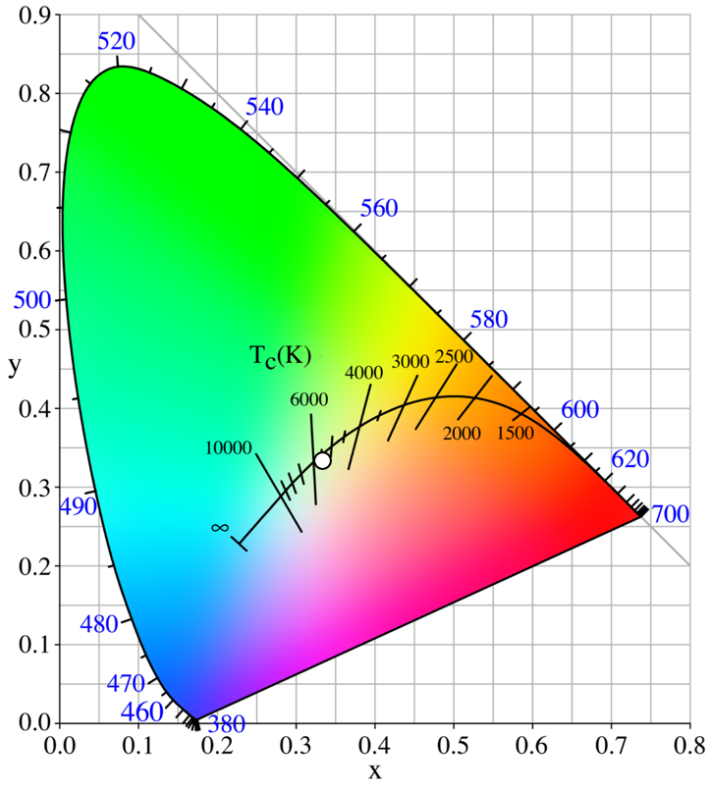


**Figure 1.2:** CIE chromaticity diagram adapted from [4]. Included are also the blackbody locus with corresponding temperatures, the wavelengths in nm of the pure spectral colours and the theoretical white point (white circle), or equal energy illuminant at $[x, y] = [\frac{1}{3}, \frac{1}{3}]$. Displayed colours are guides only.

diagram derived from the CIE XYZ colour space. It contains within it all the colours visible to a typical human. The pure spectral colours are located on the perimeter of a horseshoe shape, connected at its base by the so-called 'line of purples', which have no monochromatic realisation. If one has two different colour lamps, then the colours accessible by tuning their relative contributions to a mixed light can be found on a straight line connecting the two lamp colours in chromaticity space. If one has three lights, colours within the area of the formed triangle are accessible.

At the center of the diagram, at coordinates $[x, y, z] = [\frac{1}{3}, \frac{1}{3}, \frac{1}{3}]$ lies a point of no discernible

colour, the pure white point. This is the 'colour' of a spectral power distribution that is perfectly flat across the visible region. This is however, by the principles of additive colour mixing discussed, only one of an infinite number of ways of obtaining light that appears to be white.

The black curve atop the colour space in Fig.1.2 tracks the colour of emission from an ideal black body as its temperature is increased ($T_c(K)$). The colour of a Planckian radiator (black body) provides a convenient metric for defining varying 'warmths' of white for consumer applications. A lamp with chromaticity coordinates that lie near the black body locus (limits of applicability given roughly by the extent of the black lines in Fig.1.2), may be said to have a correlated colour temperature (CCT) value equivalent to the temperature of the black body whose colour matches it most closely. If one needs a lamp for an office space, the crisp white of a CCT around 6000 K would be recommended, whereas a lamp with CCT around 3000 K will have a much warmer (not in the black body sense), yellower colour better suited to a relaxed environment.

## 1.3 Solid-state Lighting

### 1.3.1 Light-Emitting Diodes (LEDs)

In reciprocity with the manner in which a p-n junction can convert photons into current in photovoltaics, a semiconductor diode can become light-emitting with application of a suitable voltage. A p-n junction is formed at the interface between two counter doped regions of semiconductor. The P-type semiconductor having been doped to contain a positive majority of charge carriers and the N-type a negative majority. Under an applied forward voltage bias the charge carriers in each region drift towards the junction and recombine with each other. For direct band-gap semiconductors around room temperature, the probability of radiative recombination is high. The wavelength of emission from the LED is dictated by the size of the energy band-gap of the semiconductor. An LED emission spectrum is however broadened by the effect of temperature on the distribution of charge carriers within the energy bands, as governed by the Fermi-Dirac distribution. For room temperature LEDs of visible emission wavelengths, an emission full-width at half-maximum of the order 30 - 40 nm is typical, see Fig.1.3 [5, 6].

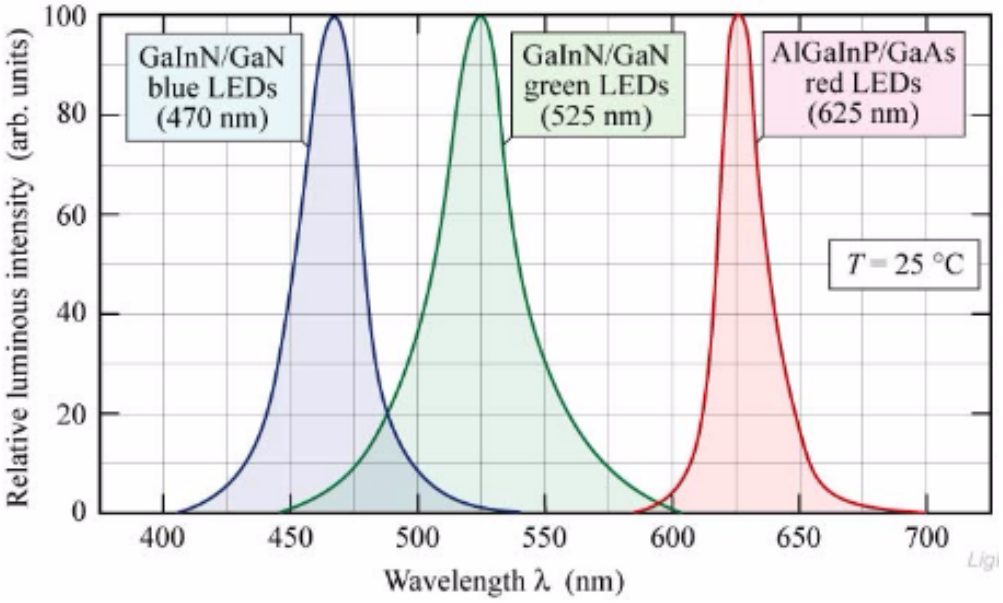


**Figure 1.3:** Figure from [5] showing typical emission spectra for common red green and blue light emitting diodes at room temperature.

### 1.3.2 Laser diodes (LDs)

A simple p-n junction can satisfy the conditions for lasing when subjected to very high levels of electrical pumping. In such a homojunction diode laser, the spontaneous emission is only weakly confined to the junction region by the slightly elevated refractive index. This index change is caused by the additional charge carriers present in the thin (1-3$\mu m$) active region [6]. In semiconductor diode lasers positive optical feedback may be achieved by cleaving the high index semiconductors along their natural crystal cleavage planes, forming highly reflective surfaces. Due to the weak optical and carrier confinement to the active region in homojunction diode lasers, such vigorous pumping is required to achieve population inversion that they are usually restricted to operation in pulsed mode [6].

The shortcomings of the homojunction diode laser are alleviated by the heterojunction diode laser, which could for example feature an additional compound at the junction centre with narrower energy gap and higher refractive index than its sandwiching compounds. In this way the efficiency of carrier and optical confinement can be vastly improved and threshold current lowered accordingly. Heterojunction diode lasers form the basis for the majority of modern semiconductor lasers and come in a variety of realisations [6].

For the purpose of this work the terms light-emitting diode and laser diode are largely interchangeable. We are concerned only with engineering the phosphor component that

comes after the pump light is generated. The only properties of the pump light source that are of any real consequence for the phosphor are the emission angular distribution and the spectral composition. However, it should be noted that laser diodes (used in conjunction with colour-converting phosphors) are predicted to dominate future high power applications [7]. This is due to a phenomenon known as the efficiency droop which exists for light-emitting diodes. The droop presents itself as a reduction in internal quantum efficiency with increasing current density that is often attributed to either Auger recombination or carrier spill-over [7, 8, 9]. In laser diodes these effects are less evident due to the reduced timescales on which stimulated emission occurs [7].

The emission from a semiconductor diode laser is divergent due to diffraction, and diverges at different rates in the directions parallel and perpendicular to the narrow confinement region. This astigmatism can be corrected with a cylindrical lens and in general the smaller emission areas of laser diodes offer much better angular and spatial control of their emission than their LED counterparts. The dramatically reduced spectral width of the typical blue LD vs its LED counterpart still coincides with strong absorption from YAG:$Ce^{3+}$, so the primary consequence of the spectrally narrow LD pump is reduced colour rendering in the blue spectral region [10]. Finally it is worth mentioning that the coherence of the LD can in some cases present unwanted interference patterns, known as speckle patterns, in the residual pump beam [11].

### 1.3.3 White Light-emitting Diodes / Laser Diodes

**The RGB Approach**

It is clear from our understanding of white light (Section 1.2) and the emission properties of LEDs (Section 1.3.1) that a single LED / LD is incapable of producing white light, such a source falls near the perimeter of the chromaticity diagram of Fig.1.2. We have seen however that a combination of coloured lights may be used to access a wide range of chromaticity coordinates. A combination of red, green and blue (RGB) LEDs such as those in Fig.1.3, allow the formation of the majority of visible colours, the 'colour' white included.

Whilst an RGB LED array may be conceptually the simplest means of producing white light from semiconductor diodes, in practice its popularity falls short of that of the alternative phosphor-converted white LED / LD for general illumination. One of the reasons for this is the poor conversion efficiency from electrical input to luminous output of green LEDs [12]. This is typically less than half that of their red and blue counterparts. Another issue for commercial applications is the cost and complexity of implementing an RGB LED

device, as each colour of LED requires its own drive electronics. Further contributing to this issue is that each colour of LED degrades at a different rate over its perhaps as much as 20 - 30 year lifespan. Whilst there are permitted drops in efficiency over the lifetime of a lighting product, the requirement to maintain the same perceived output colour further detracts from the appeal of the RGB approach in this context.

**Phosphor-Converted White LEDs (PCWLEDs)**

Phosphor is the general name given to materials which exhibit luminescence (emit light for reasons other than temperature). In this case, a phosphor layer provides partial down-conversion of a blue light emitting diode source by means of fluorescence. The light that emerges from the phosphor layer is composed of both yellow fluorescence emission and unconverted blue pump light. An emission spectrum from a typical PCWLED can be seen in Fig.1.4. The colours blue and yellow are known as complementary, meaning that a line connecting them on the chromaticity diagram of Fig.1.2 passes through the white point. As such, combined in the appropriate proportions they appear as white light to the human eye [12].

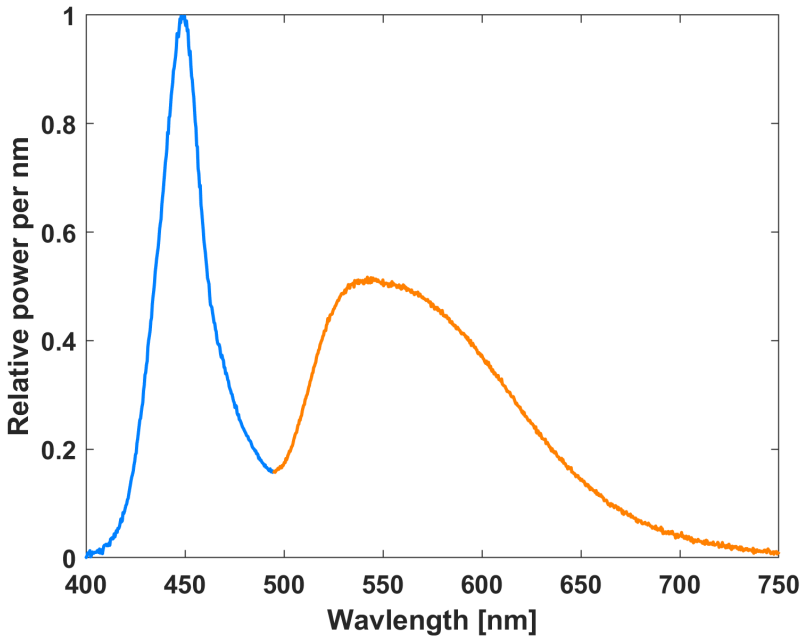


**Figure 1.4:** Emission spectrum of a YAG:$Ce^{3+}$ phosphor illuminated with a blue LED source. Phosphor illuminated with an Ocean Optics LLS-455 LED and spectrum measured in an integrating sphere. The contributions directly from the LED and from fluorescence are loosely separated by the blue and orange line colours. The specifics of the phosphor can be found in in the later Section 4.7 in Table 4.1 under the designation MonoCryst disc.

### 1.3.4 YAG:$Ce^{3+}$

The most commonly used phosphor for producing high brightness white lights is a cerium doped Yttrium Aluminium Garnet (YAG:$Ce^{3+}$). This is the phosphor considered throughout this work, although the principles could be transferred to other phosphors of compatible optical properties.

YAG:$Ce^{3+}$ earns its popularity for its favourable absorption and emission characteristics and low fluorescence lifetime of ~ 70 ns, enabling high luminous output [12]. YAG:$Ce^{3+}$ posses a strong absorption band around 460 nm (see Fig.1.5), aligning with the peak emission from efficient GaInN/GaN blue LEDs. The corresponding absorption cross-section $\sigma = 3 \times 10^{-18}$ $cm^2$. Its fluorescence emission spectrum is broad, and covers the wavelength range from ~ 500 - 700 nm, as shown in Fig.1.4. A GaInN/GaN blue LED pumped YAG:$Ce^{3+}$ phosphor can efficiently produce a bright white light.

Perhaps the only significant drawback of YAG:$Ce^{3+}$ is its lack of emission in the red portion of spectrum, resulting in a low colour rendering index (CRI) and a 'cool white' appearance. The CRI of a light source is a quantitative measure of its ability to faithfully render colours with respect to ideal daylight or a black body radiator, with 100 representing ideal rendering. For YAG:$Ce^{3+}$ pumped with blue light the maximum attainable CRI is around 70-80 [13]. This poor colour rendering is typically overcome by introducing additional dopants, such as $Eu^{3+}$ which emit around 620 nm and produce a more pleasing 'warm white' light [14].

Fig.1.5 shows a Jablonski (energy level) diagram and accompanying absorption and emission spectra for the process of fluorescence from a YAG:$Ce^{3+}$ phosphor. First a high energy photon is absorbed by a cerium ion, promoting an electron from a $4f$ energy level into an excited $5d$ level (upwards solid arrow). This electron then undergoes some rapid non-radiative decay events to reach the bottom of the lowest $5d$ energy state (wavy arrows), in which energy is lost to the surrounding crystal as heat. From this meta-stable level a radiative decay event occurs, back to the ground state. In YAG:$Ce^{3+}$ the ground state is split (primarily) by spin-orbit coupling into two $^2F_{5/2}$ and $^2F_{7/2}$ states, resulting in a broad 500 - 700 nm emission spectrum composed of two Gaussian curves [16].

The location of the favourable absorption and emission bands in YAG:$Ce^{3+}$, dictated by the $4f$ to $5d$ energy separation, is the result of the interaction of the host crystal with the dopant Cerium ion. Unlike for the free Cerium ion, which has 4f - 5d separation of ~ 50,000 $cm^{-1}$ [17], within YAG this separation is substantially reduced to ~ 22,000 $cm^{-1}$. This reduction is a consequence of the strong interaction between the YAG crystal field and the

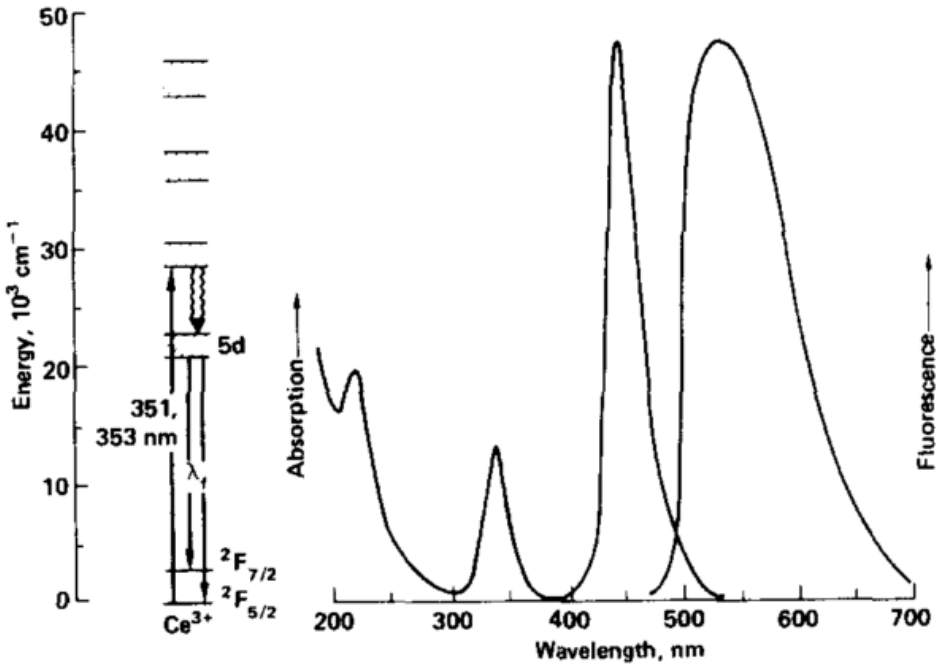


**Figure 1.5:** Figure from [15], showing the energy level diagram and absorption / fluorescence spectrum of YAG:Ce at room temperature.

radially extensive 5d electron wavefunction, which shifts and separates the 5d levels [17].

The internal quantum efficiency of a YAG:$Ce^{3+}$ phosphor is excellent with maximum around 94% [18, 19]. The Cerium dopants introduced into the cubic YAG crystal occupy the Yttrium sites by substitution and optimal doping concentrations for traditional powder implementations are around 1-2 mol.% [18]. In monocrystalline YAG:$Ce^{3+}$ the minimum achievable inter-dopant spacing is around 2.4 nm [18]. It has a refractive index of 1.83 and is most often embedded as a powder in a lower index resin host of index around 1.5 to exhibit high levels of scattering. In isotropic lighting applications this scattering provides good colour mixing of the pump light with the fluorescence and enables sufficient pump absorption within a thinner phosphor layer.

For high power applications the standard technique of embedding the phosphor in a silicon host proves ill-suited. The host materials used have poor thermal conductivity (~0.2 $Wm^{-1}K^{-1}$ [20]) and are thus ineffective at dissipating the heat generated by the Stokes shift of the emission. In addition, it is common for the host materials to discolour under prolonged operation at temperatures approaching 100°C [20]. There is therefore a wide body of research dedicated to either replacing the host materials or removing them entirely.

Low melting temperature glasses are a popular choice for replacement of the epoxy resin host as they minimize damage to the phosphor during fabrication [21, 22, 23, 24, 25]. We favour the alternative host-free approach as at present it offers higher quantum efficiencies and is readily available commercially [19]. For the duration of this work we will focus on monocrystalline (also called single-crystal) YAG:$Ce^{3+}$, which is scatter-free, as our phosphor of choice. The uninterrupted crystal lattice provides vastly superior thermal conductivity (13 $Wm^{-1}K^{-1}$ [19]) and optical transparency when compared to epoxy resin based approaches.

## 1.4 The Problem in Detail

With some basic background covered we are now better positioned to understand the full scope of the problem. The objective is reproduced here for convenience.

**Objective:** *Develop a highly efficient phosphor component that is capable of converting blue pump-light into white-light emission whilst retaining the incident spatial intensity pattern of the pump-beam and emitting into a narrow solid-angle.*

The problem that needs to be addressed is the angular and spatial distribution of the fluorescence generated by the phosphor component. The orientation of the transition dipole moments of the activator ions in the phosphor are random at the time of emission, resulting in isotropic fluorescence for the ensemble average [26]. This can be quickly confirmed with two linear polarisers and a phosphor sample that is free from multiple scattering. Fig.1.6 shows two emission spectra for polarised blue pumping of such a phosphor. The spectrum with analyser in the same orientation as the source polariser ($\theta = 0°$) shows almost exactly the same amount of fluorescence as when the analyser is rotated 90°. From this lack of change we infer that the fluorescence is essentially unpolarised and that the fluorescent ions retain no memory of the orientation of the electric field that excited them.

Fig.1.7 illustrates the problem traditional phosphor solutions pose for directional lighting applications. The coupling efficiency of isotropically emitted fluorescence into a narrow solid angle is unsurprisingly very poor. To overcome this, one would typically look to a combination of reflectors and lenses to redirect the emission into the target directions. For high power applications where the phosphor area may need to be large, this could require bulky and expensive optical components. In any case, redirecting the light must come at some increase in the area of the optical component following the conservation of

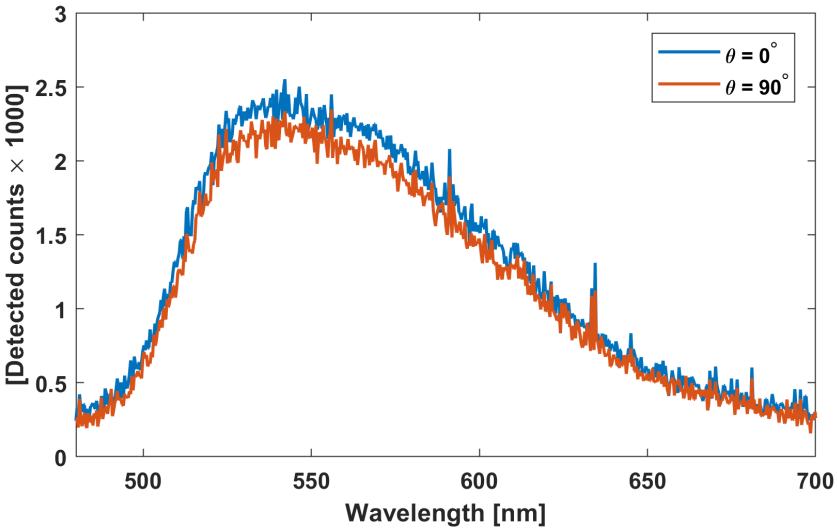


**Figure 1.6:** Unpolarised fluorescence spectrum from a blue LED pumped moncrystalline YAG:$Ce^{3+}$ disc (see Table 4.1 in Section 4.7 for full sample details). $\theta$ is the angle between the source and analyser polarisers. These spectra are not calibrated for instrument response factors.

etendue. The term etendue from the french étendue géométrique meaning 'geometrical extent', quantifies how spread out a light source or ray bundle is in a material as the product of its area and emission solid angle [27].

The problem is complicated further by the need to retain spatial patterning of the pump light in the generated fluorescence. This restricts the use of single large optical components and tends to favour arrays of optically isolated elements (or pixels). This is the essence of both the wavelength-scale parabolic reflector and fluorescent fibre approaches explored in Sections 2 and 4.

### 1.4.1 Functional Requirements

We conclude this more detailed look at the problem and objectives of this work with a breakdown of the 6 main categories or functional requirements by which a solution to the project objective may be assessed.

1. **Emission Cone**

   Emission into a narrow solid angle is the highest priority requirement of the desired phosphor component. The precise target emission cone is not a fixed parameter, as this may need to be balanced with the optical efficiency of coupling into the target

cone to maintain an overall minimum device efficiency when implemented into a larger optical system. We consider emission angles less than 20° from the target direction, i.e. into a 20° half-angle cone, to be the upper limit on 'narrow solid angle'.

2. **Imaging Capability**

   The phosphor component needs to be able to retain the incident spatial intensity pattern of the pump beam in its output. This means if an image were to be projected in blue light onto the phosphor, the image is maintained in the spatial distribution of the fluorescence and thus in the mixed white light output. Again no hard limits on the fidelity of the preserved pattern are imposed. However, it is useful to point out that if the proposed solution relies on an array of distinct optical elements acting as pixels, a lower resolution limit of the order tens by tens of pixels is suggested.

3. **Optical Efficiency**

   For the phosphor component to be used in interior lighting applications, the product it is implemented in will need to meet some overall regulatory minimum efficiency. The target efficiency for the phosphor component is to produce 200 lumens of direc-

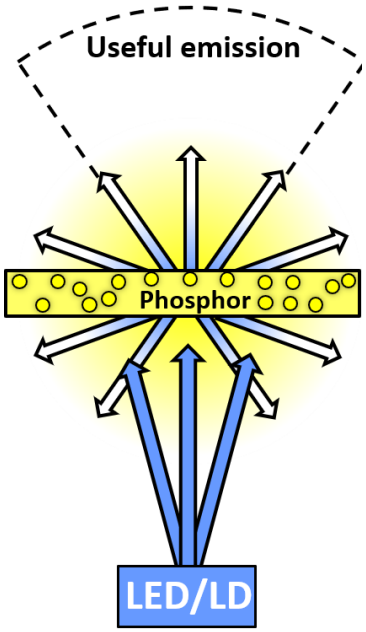


**Figure 1.7:** The problem presented by traditional phosphors for directional lighting applications: due to the isotropy of the fluorescence generated within the phosphor and the added scattering, the optical coupling of emission into a narrow solid angle is very poor.

tional white output per watt of blue light it is pumped with.

4. **Brightness / Power Handling**

   Here we find the first of the requirements not explicitly stated in the project brief. The phosphor component needs to be capable of producing in the region of 10,000 to 20,000 lumens of directional white output within a package size no larger than your average shoe box (a very approximate upper size limit). This high power handling has implications for the thermal management of the phosphor, discussed briefly in the Thermal Conductivity section below. In addition, this imposes some limits on the total amount and spatial distribution of the phosphor.

5. **Colour**

   To be considered a high quality white light suitable for interior lighting, the output from the phosphor component is desired to meet certain minimum colour standards. The first is an achievable correlated colour temperature (CCT) that covers the range 3000 - 7000 K, with the option for dynamic tuning of the CCT within a product being a desirable bonus.
   Second is the colour rendering index, which for the final phosphor solution is desired to be greater than 90. However, this high CRI target is not imposed on this work if the proposed solution has the potential to be adapted to a higher CRI. This flexibility acknowledges that the preferred phosphor candidate, YAG:$Ce^{3+}$, exhibits poor colour rendering in the red portion of its emission spectrum, but that with the inclusion of additional dopants or secondary phosphors this may be rectified [28, 29, 30, 31, 32].
   The final key consideration is the angular and spatial uniformity of the colours exiting the phosphor. Whilst no official uniformity standards were imposed, it is natural to favour methods of achieving directional emission which do not exhibit unnecessary spectral discrimination.

6. **Thermal Performance**

   In a phosphor, heat generation is an inevitable consequence of the Stokes shift between the absorbed and emitted light, i.e. the change in wavelength from pump to fluorescence. Operating at high powers dissipation of this generated heat becomes critical. The existing industry standard phosphor implementation of a polycrystalline YAG:$Ce^{3+}$ powder embedded typically in a silicone resin is ill-equipped to handle

high power lighting applications. The host materials used offer poor thermal conductivity and may themselves discolour under prolonged operation at temperatures approaching 100 ° [20]. Choosing an alternative phosphor implementation with improved thermal management was therefore necessary.

## 1.5 Ghost Imaging

In the final chapter of this thesis, Chapter 6, having presented our solution to the posed lighting problem, we shift our focus to the potential applications of a directional phosphor component. The potential applications of a phosphor component meeting the functional requirements outlined in Section 1.4.1 are numerous. Even without meeting the imaging capability requirement, an efficient, bright and directional phosphor could find applications in medical and scientific imaging, automotive headlights, architectural and general lighting to name just a few examples.

Once one includes the ability to maintain spatial patterning of the pump light in the directional white emission, the applications of such a source become more specialised. The majority of such applications will naturally fall under the category of 'projection applications'. These could include conventional projectors for display technology, augmented reality displays, multi-channel visible light communication and computational imaging devices. It is in the last category that we find ghost imaging, the subject of the final chapter of this thesis.

In ghost imaging the information regarding where light from the scene originated is not provided by a multi-pixel detector. Instead the total transmission / reflection is measured by a bucket detector without spatial resolution and this measurement is combined with knowledge of the spatially patterned wavefront illuminating the scene, to enable reconstruction of an image [33]. The illumination wavefront is usually either generated in some chaotic way (e.g. with random speckle patterns) and measured with a second detector, or is user-specified by a device such as an array of digital micromirrors. The latter scenario is referred to as computational ghost imaging [34], and can be considered a candidate application of the desired directional phosphor with imaging capabilities.

In Chapter 6 we will cover the principles of ghost imaging in greater detail, and will introduce a method for incorporating image post-processing steps directly into the illumination basis.

## 1.6 Summary

In this introductory chapter we've covered the basics of solid-state lighting and colour science, to a degree that has allowed us to understand the problem presented by conventional phosphor solutions for application in directional light sources. The requirements of the desired phosphor component have been outlined in the Functional Requirements Section 1.4.1, these form an important part of the assessment of candidate phosphor solutions and will be referred to throughout this work. Finally we introduced ghost imaging as a potential application for such a phosphor, as it will be the topic of the final body of work in this thesis (Chapter 6).

We now begin the results chapters of this thesis, starting with an assessment of the lower size limits at which the traditional parabolic reflector can still provide highly directional emission.

# 2 Wavelength-scale Parabolic Reflectors

## 2.1 Motivation

The majority of phosphor converted white LED products with directional light output exploit a reflective component between the phosphor and unwanted emission directions, often in conjunction with a simple lens. An example of standard low-cost directional white LED implementation is shown in Fig.2.1.

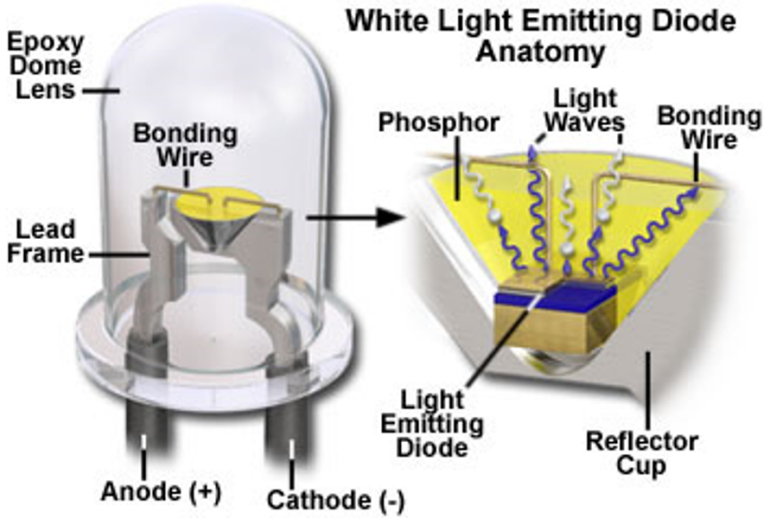


**Figure 2.1:** Figure from [35] showing typical low-cost directional white LED design. Here the reflector cup is of conical shape and the dome lens also serves as the encapsulant.

The complexity of the reflective component design can vary considerably depending on factors such as the emission cone requirements, manufacturing cost and product dimensions. Reflectors may be conical as in Fig.2.1, hemispherical or planar reflective substrates, but the mathematically ideal shape is the 3D parabolic reflector (paraboloid of revolution).

A paraboloid of revolution has the remarkable property that all incident rays parallel to the central axis will be reflected to the focus, making it the ideal shape to collect a collimated beam, or to generate one starting from a point source. Fig.2.2 shows a simplified scheme of such a parabolic antenna. To achieve the perfect collimation depicted in Fig.2.2, the source positioned at the focus must be infinitesimally small and either the paraboloid infinite or the source emission anisotropic (direct emission out of the opening of a finite paraboloid will not be collimated). Nevertheless, for a spatially extended source in a finite reflector, one can usually balance the emission angle or efficiency requirements with the size of the reflector to achieve highly directional illumination.

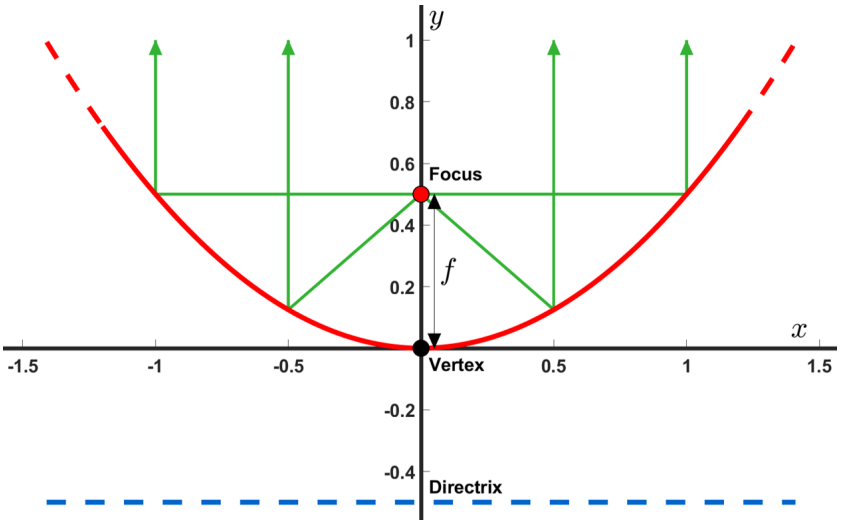


**Figure 2.2:** A section of an infinite parabola (red) on Cartesian axes, defined by being equidistant from both the focus (red dot) and the directrix (blue, dashed) at all points. The distance from vertex to the focal point is the focal length $f$. Emission from the focal point (green lines), is collimated by specular reflection from a parabolic surface.

A single parabolic reflector is unable to meet the imaging capability requirement of the desired directional phosphor component outlined in Section 1.4.1. Spatial patterning of the input beam cannot be preserved in the generated fluorescence in any meaningful way if the fluorescence is generated within a single homogeneous phosphor lump at the parabola's focus. A logical solution is then to discretise the phosphor and use an array of parabolic reflectors, each acting as a spatially separated directional light source. Each reflector may then be illuminated as an individual pixel, thereby preserving the patterning of the input beam in the directional fluorescence output.

The minimum number of pixels required has been suggested in the Functional Requirements Section 1.4.1 as a few tens of pixels by tens of pixels. To meet the phosphor area limitations defined therein an upper limit on the width of an individual (square) pixel is of the order of 1 cm. This is of course just a limit and the option to use smaller elements with higher system resolution or reduced overall size (easier optical manipulation) will always be attractive. As such, an interesting question arises - how small can one make a parabolic reflector before it stops working effectively? Additionally, in the vicinity of this limit does the geometrical optics notion of a focal point remain the optimum position for an emitter? These are the central questions explored in this chapter.

For a reflector of size significantly greater than the wavelength of emission from within it ($s >> \lambda$), the geometrical optics (ray optics) approximation is an excellent model for

light propagation in the reflector system. This is the scenario depicted by the green rays in Fig.2.2. The geometrical optics approximation neglects the effects of diffraction and interference, which play an increasingly important role as the reflector size decreases towards that of the emission wavelength. It is at the wavelength scale ($s \sim \lambda$), where the geometrical optics approximation dramatically breaks down that we explore the answers to our posed questions. In the following sections, we will first use finite-element modelling to assess the finite parabolic reflector in the geometrical optics regime and then compare the results to those obtained for wavelength-scale parabolic reflectors.

## 2.2 The Geometrical Optics Parabolic Reflector

To compare the performance of a reflector for a given emitter position, we first define a figure of merit which will quantify its utility for a directional lighting (or by reciprocity, measurement) system. The 'go-to' parameter for an antenna system might be its directivity ($D$), which is a measure of the how concentrated its radiated power is into a single direction. Specifically, it is the maximum value of the radiation intensity in a given direction ($U(\theta,\phi)$), with respect to the average power radiated in all directions by the antenna [36]. $\theta$ and $\phi$ are the polar and azimuthal angle in the spherical polar coordinate system. In the absence of a specified direction for which the directivity is being provided ($D(\theta,\phi)$), the direction of maximum radiation intensity is implied ($D_{max}$) [36]

$$D_{max} = \max\left(\frac{U(\theta,\phi)}{P_{tot}/4\pi}\right), \tag{2.1}$$

where $P_{tot}$ is the total power radiated by the antenna. As the directivity is a maximum value and could correspond to peak power radiated in any given direction, we do not consider it a natural choice for comparing neighbouring emitter positions where only radiation into a common emission direction is desired. This could be overcome by instead using the directive gain $U(\theta,\phi)/\left(P_{tot}/4\pi\right)$ in a chosen direction (e.g. along the central axis of the paraboloid), however this does not address the most significant limitation of these parameters for this application. For assessment of an optical component of this type the total power radiated into a range of acceptable angles, or a target emission cone is of far greater value than just the peak directional power value somewhere within that cone.

For the reasons outlined above, we have chosen the power directed into a 20° half-angle right circular cone in the far field, as a useful measure of the performance of the reflector. The cone and paraboloid share a common central axis and vertex. We have cho-

sen the paraboloid defined in Cartesian coordinates by $z = (x^2 + y^2)/2$, with finite limits $x \in (-\sqrt{2}, \sqrt{2})$, $y \in (-\sqrt{2}, \sqrt{2})$, and $z \in (0, 1)$, to present our results. This is the finite section of a parabola shown in Fig.2.2, with the $\hat{z}$ axis in place of the $\hat{y}$ axis, and rotational symmetry about this central axis.

Fig.2.3 shows how the directed power (as defined above) varies with emitter position, for one quadrant of the paraboloid. The solution has rotational symmetry about the central axis but this quadrant view will prove beneficial when comparing to the non-rotationally symmetric results for individual dipole emitters in Section2.5. The directed power values were obtained through finite-element modelling for a regular 2D grid of emitter positions within the paraboloid, in a plane intersecting its central axis. As the simulated emitters are isotropic, a 2D sweep on one side of the central axis is sufficient to obtain values for the entire volume of the reflector, exploiting the rotational symmetry present. Full details of the COMSOL model used and post-processing can be found in Section 2.6.3.

As expected, in Fig.2.3 a region of efficient coupling into the target cone is found when

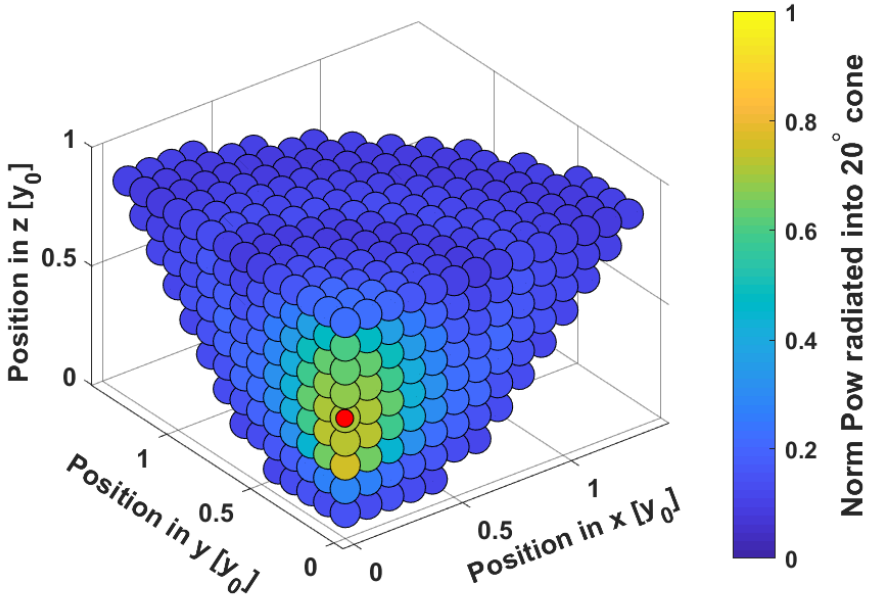


**Figure 2.3:** Geometrical optics solution for power directed into a 20° half-angle cone, for an isotropic emitter at varying positions within a finite parabolic reflector. As the solution is independent of the size of the reflector in the geometrical optics regime, the dimensions are given in terms of the vertex to opening distance $y_0$. Power values are normalised to the total emitted power. The red dot marks the focus of the paraboloid. Simulation details in Section 2.6.3.

the source is near the focus of the reflector. This efficiency remains below 100% as the reflector is of finite extent. It is also noteworthy that the focus is not the optimum position for a source even within the geometrical optics approximation, this is a consequence of both the finite extent of the paraboloid and the finite collection angle. Whilst a source at the focus still optimises emission on axis, emission into the cone is optimised when the source lies below the focal point, as seen in Fig.2.3.

## 2.3 The Oscillating Electric Dipole

When considering the interactions of an ensemble of fluorescent emitters with a structure, it is often a necessary (if not at least enlightening) step to first understand the interactions of a single emitter with the structure. The nature of an individual emitter becomes particularly important when length scales approach that of the wavelength of emission, as we shall see. The majority of sources of spontaneous emission (fluorescence from YAG:$Ce^{3+}$ included) can be described by an oscillating electric dipole [37, 38, 39]. Starting from an electric dipole moment ($\boldsymbol{p}$) oscillating at angular frequency $\omega$ in the $\hat{\boldsymbol{z}}$ direction

$$\boldsymbol{p}(t) = p_0 \cos(\omega t)\, \hat{\boldsymbol{z}}, \tag{2.2}$$

it can be shown ([40]) that the radiative fields at angle $\theta$ from $\hat{\boldsymbol{z}}$ and distance $r$ far from the emitter are

$$\mathbf{E} = -\frac{\mu_0 p_0 \omega^2}{4\pi}\left(\frac{\sin\theta}{r}\right)\cos[\omega(t - r/c)]\, \hat{\boldsymbol{\theta}}, \tag{2.3}$$

and

$$\mathbf{B} = -\frac{\mu_0 p_0 \omega^2}{4\pi c}\left(\frac{\sin\theta}{r}\right)\cos[\omega(t - r/c)]\, \hat{\boldsymbol{\phi}}, \tag{2.4}$$

where the angle $\theta$ is defined as starting along the axis of the dipole moment $p_0$ and all other symbols have their conventional meanings. The approximations used rely on the dipole length $l$ being small and the field evaluation distance $r$ large, such that $l \ll \lambda \ll r$, where $\lambda$ is the emission wavelength. In particular, as only radiative fields are of interest, conservation of the integral of the Poynting vector ($\mathbf{S}$) over the surface of a sphere at increasing radii requires that only fields following a $1/r$ dependence are retained [40]. The Poynting vector is obtained from the electromagnetic fields via

$$\mathbf{S} = \frac{1}{\mu_0}(\mathbf{E} \times \mathbf{B}). \tag{2.5}$$

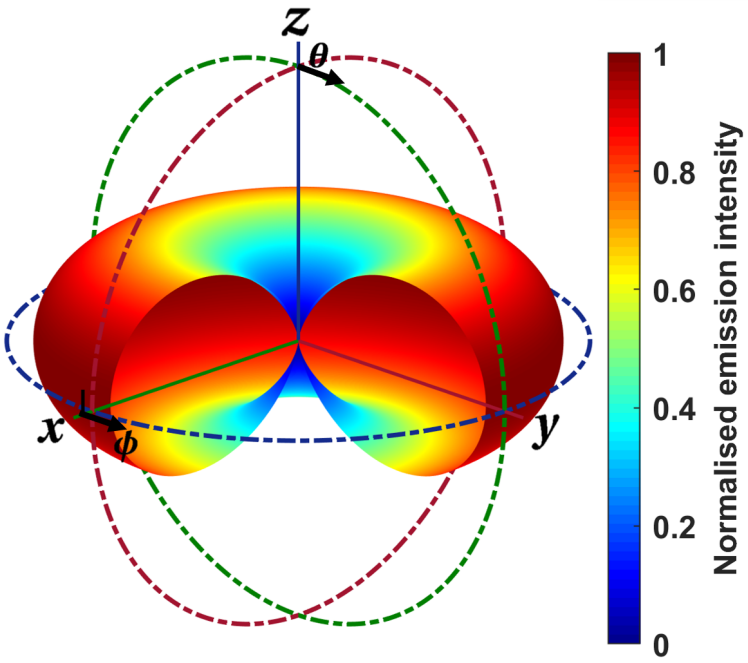


**Figure 2.4:** Free space radiation pattern of an electric point dipole oscillating in the $\hat{z}$ direction. Data for $0 < \phi < \pi/2$ have been omitted for better visual clarity.

Finally, the intensity is given by the time averaged Poynting vector

$$\langle \mathbf{S} \rangle = \left( \frac{\mu_0 p_0^2 \omega^4}{32\pi^2 c} \right) \frac{\sin^2 \theta}{r^2} \, \hat{\mathbf{r}}. \tag{2.6}$$

The resulting doughnut-shaped radiation pattern can be seen in Fig.2.4. Crucially, the oscillating electric dipole does not radiate along the axis of its dipole moment and so interactions with structures or external electromagnetic fields may be highly dependent on dipole moment orientation.

## 2.4 Fluorescence Near a Metallic Interface

The simplest scenario in which an emitter interacts with a reflector is that of the planar metallic interface, that is, an emitter some distance above an infinite flat mirror. In the geometrical optics regime this primitive reflector simply redirects all rays radiated towards the mirror (half of the total) back into the upper half-space preserving their angles relative to the mirror normal. This results in double the amount of light, with the same angular distribution in the upper half-space and is a principle applied to increase the efficiency of

phosphor discs in projector applications. It has been assumed here that the mirror is perfect (i.e. lossless) and this assumption will be applied for the remainder of this chapter for simplicity.

If the emitter is separated from the mirror by a distance ($d$) that is not significantly larger than the emission wavelength, the geometrical optics approximation is no longer valid and the effects of interference must be considered. The electromagnetic fields surrounding the emitter and radiated into the far-field are now the result of the combination of what would be the free space fields and those reflected by the mirror. The reflected fields may be conveniently handled using the concept of image charges. The total fields of the emitter - mirror system can be calculated as the sum the original fields with an out of phase image dipole, located a distance $d$ below the mirror. The oscillating dipoles are in antiphase due to the $\pi$ change in phase of the electric field reflected from the mirror. This notion is not valid within the mirror, where the fields are zero.

Fig.2.5 compares the radiation pattern of an emitter one emission wavelength above a perfect mirror with that radiated into the upper half-space by an equivalent free space emitter. The angular distribution of emission is dramatically modified by the interference between the direct and reflected fields.

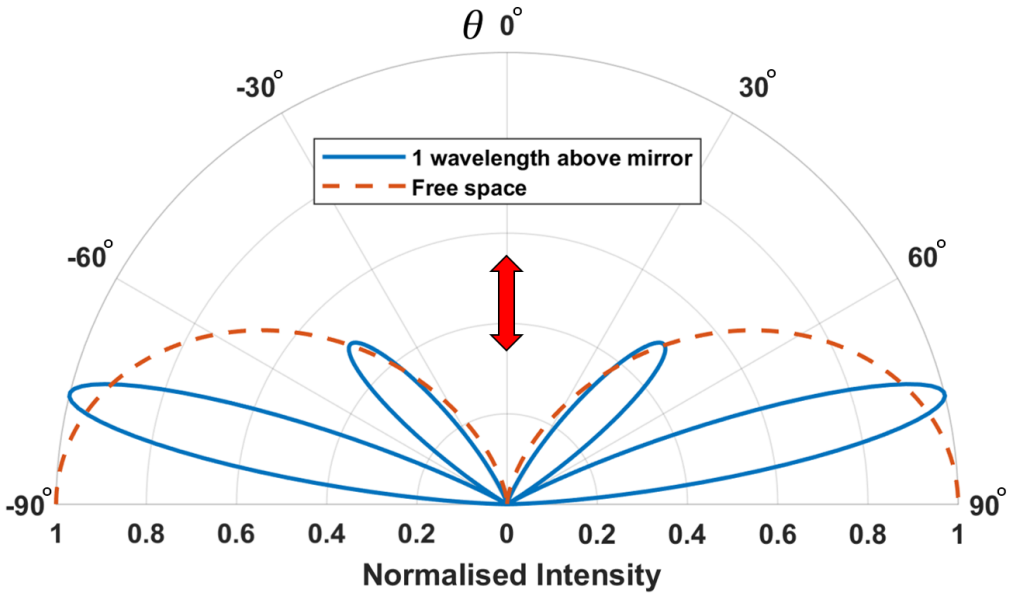


**Figure 2.5:** Polar plot of the upward radiation pattern of a free space oscillating electric dipole with one located one emission wavelength above a perfect infinite mirror. The dipole moment is oscillating in the $\theta = 0$ direction, as indicated by the red arrow, in both cases. Calculation details in the accompanying text.

The radiation pattern of the emitter near the mirror in Fig.2.5 has been calculated from the net electric field far from the original and image oscillating dipoles using eq.2.3, as $|\boldsymbol{E}|^2$. The image dipole fields include a phase shift of $\pi$ inside the Cosine function and the distance from each source is given by

$$r_{\pm} = \sqrt{r^2 + d^2 \pm 2rd\cos(\pi - \theta)}, \tag{2.7}$$

where $r_+$ and $r_-$ are the distances to the original and image dipole respectively and $d$ is their separation from the mirror surface. The resulting patterns can be verified with the following more compact intensity expressions for dipoles parallel and perpendicular to a perfect mirror [41]

$$I_{\parallel}(\theta) = \cos^2(\theta)[2 - 2\cos(\phi)], \tag{2.8}$$

$$I_{\perp}(\theta) = \sin^2(\theta)[2 + 2\cos(\phi - \pi)], \tag{2.9}$$

where

$$\phi = \frac{4\pi d}{\lambda}\cos(\theta). \tag{2.10}$$

The modified radiation pattern seen in Fig.2.5 can also be understood in terms of emission into allowed and forbidden electromagnetic modes. Due to the destructive interference between the source and image dipole at $\sim 60°$ for example, there is no longer an electromagnetic mode available for the emitter to radiate into and thus no radiation occurs. Conversely at $\sim 75°$ the availability of radiative modes is at its highest. It may be tempting to remark that the number of available modes here is elevated with respect to free space, but equations 2.8 amd 2.9 tells us nothing of the total radiated power, which indeed may be modified considerably with respect to free space.

K. H. Drexhage was the first to demonstrate that the total radiated power of an emitter may vary with position when in close proximity to a mirror [42]. There the changing power was probed by measuring the variable emission lifetime, which in the absence of non-radiative decay mechanisms shares an inverse proportionality with the total radiated power and the emission rate $\Gamma$. The connection between these properties of an emitter and the availability of electromagnetic modes can be understood using Fermi's golden rule [43]

$$\Gamma_{i,j} \propto |M_{i,j}|^2 \rho(E_j), \tag{2.11}$$

where $\Gamma_{i,j}$ is the transition rate from excited state $i$ to lower state $j$, $|M_{i,j}|$ is the matrix element connecting states $i$ and $j$ and corresponds to the dipole moment of the source and $\rho(E_j)$ is the density of final states [44]. One can translate this quantum terminology to the language of electromagnetism by realising the words state and mode are synonymous and the density of final states $\rho(E_j)$ is no different from the density of available electromagnetic modes for the emitter, at a given frequency $\rho(\omega)$.

The quantity probed experimentally by Drexhage was the spatially varying relative of the Density Of States (DOS) $\rho(\omega)$, the so-called Local Density of States (LDOS) $\rho(\boldsymbol{r},\omega)$, where $\boldsymbol{r}$ is the spatial coordinate [42]. When the LDOS is further separated into its components from a given orientation of dipole moment, it is termed the Partial Local Density of States (PLDOS) [37]. In the case of our infinite mirror example, the LDOS is constructed from the PLDOS for parallel and perpendicular dipole moments as the orientational average $\rho(\boldsymbol{r},\omega) = (1/3\rho_{\perp} + 2/3\rho_{\parallel})$. Note that the PLDOS is conventionally defined not to require the 1/3 pre-factor in the orientational average, but we choose to use it here [37].

Fig.2.6 shows the variation in total radiated power, or equivalently the PLDOS, for dipolar emitters in close proximity to a perfect mirror. The PLDOS in each case has been determined by numerical integration of following expressions [37]

$$\rho_{\parallel}(\boldsymbol{r},\omega) = \frac{\omega^2}{\pi^2 c^3}\int_0^{\frac{\pi}{2}} \Big[\sin(\theta) - \frac{1}{2}\sin^3(\theta)\Big]\sin^2\Big(\frac{\omega}{c}\cos(\theta)d\Big)\mathrm{d}\theta \tag{2.12}$$

and

$$\rho_{\perp}(\boldsymbol{r},\omega) = \frac{\omega^2}{\pi^2 c^3}\int_0^{\frac{\pi}{2}} \sin^3(\theta)\cos^2\Big(\frac{\omega}{c}\cos(\theta)d\Big)\mathrm{d}\theta. \tag{2.13}$$

It should be noted these expressions are derived from the full electromagnetic fields of the emitters, not the far field approximations of eqs.2.3 and 2.4. Fig.2.6 also contains the same parameters evaluated in 3D COMSOL simulations, see Section 2.6.2 for details, which show good agreement with the analytical expressions.

The PLDOS is shown in Fig.2.6 to be highly sensitive to the orientation of the emitter and oscillates periodically with distance from the mirror. As the distance from the mirror increases the amplitude of the fluctuations decreases, with the PLDOS tending to its free space value for all orientations. These observations are easily understood by considering the emitter to be a classical dipole antenna with a radiation efficiency that is sensitive to the amplitude and phase of the reflected fields at the antenna location, which may interfere constructively or destructively with the charge oscillations in the antenna. The PLDOS thus fluctuates with the phase of the reflected field, has a period of approximately half the emission wavelength and fluctuates less with distance from the mirror. The amplitude of

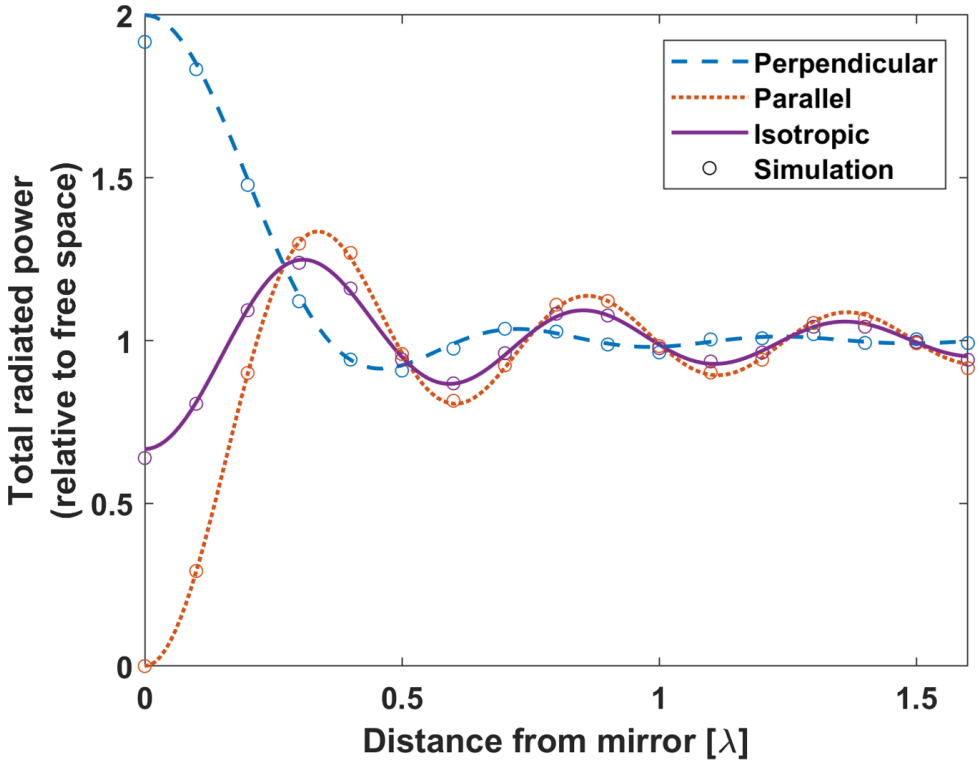


**Figure 2.6:** Variation in the total radiated power of dipolar emitters with distance from a perfect mirror. The orientation specific analytical dependencies (lines) are those of eq.2.12 and eq.2.13. The circular data points mark the results of COMSOL simulations, details of which can be found in Section 2.6.2.

these fluctuations is much larger for the parallel emitter, which radiates preferentially towards the mirror.

The final and most striking observation from Fig.2.6 is the PLDOS values at zero emitter - mirror separation. Once again the image charge concept provides the most intuitive explanation. In Fig.2.7 we can see the impact of the induced image charges on the effective dipole moment of the emitter as it approaches the mirror. For a dipole oscillating parallel to the mirror surface, at zero separation, the image charges cancel those of the emitter. This results in no net oscillation of charge, which means no radiation and a PLDOS of zero. For the perpendicular dipole the result is a dipole with twice the length and hence twice the dipole moment. This would provide a factor of 4 enhancement in radiated power, following eq.2.6, for twice the dipole moment in free space. However, as the radiative fields are only non-zero in the upper half-space, the result is a factor of 2 enhancement in the PLDOS.

It is important to note that for a real metal of finite conductivity one would observe

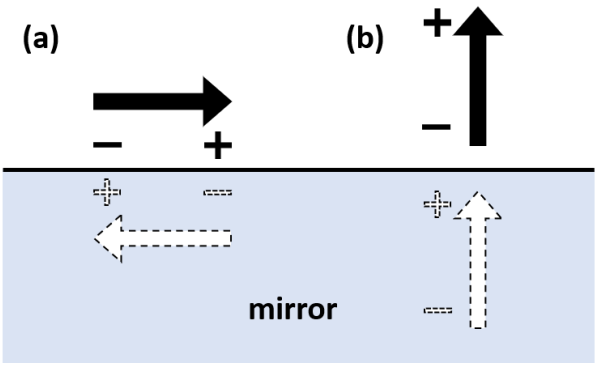


**Figure 2.7:** As a dipole emitter approaches zero distance above a perfect mirror the image dipole induced in the metal may serve to cancel ((a) - parallel orientation) or double ((b) - perpendicular orientation) the dipole moment of the emitter [44].

quenching of the emission at very small emitter - mirror separations. This is due to near field coupling to surface plasmon polaritons, hybrid light-matter states that may be supported by a metal dielectric interface [44]. This is typically only of consequence for emitter - mirror separations less than ~ 20 nm at visible wavelengths [44]. We have restricted our investigations to perfect mirrors of infinite conductivity in the following simulations.

One can expect the effects of a fluctuating LDOS to play a critical role in the performance of reflectors whose size is close to the emission wavelength.

## 2.5 The Wavelength-scale Parabolic Reflector

With the macroscale solution established in Fig.2.3 as a reference, we now consider the corresponding case for wavelength-scale paraboloids. We have chosen a paraboloid with vertex to opening height equal to the emission wavelength to examine in detail first and then we shall see how the observed behaviour varies with emission wavelength.

The method is explained in detail in Section 2.6.4 but in short is as follows: for a given emitter orientation and position within the paraboloid, use COMSOL to determine the electromagnetic fields around the reflector and perform a near to far field transformation on the simulation boundary. An example of both the spatial and angular field distribution for a given emitter can be seen in Fig.2.8. Next, numerically integrate the resulting radiation pattern to determine the power coupled into a 20° half-angle cone in the far field, our chosen performance metric. Repeat for the source positions and orientations needed to map the entire volume of the reflector, exploiting present symmetries.

Figs.2.9 (a) - (c) show the position-dependent power radiated into a 20° half-angle cone

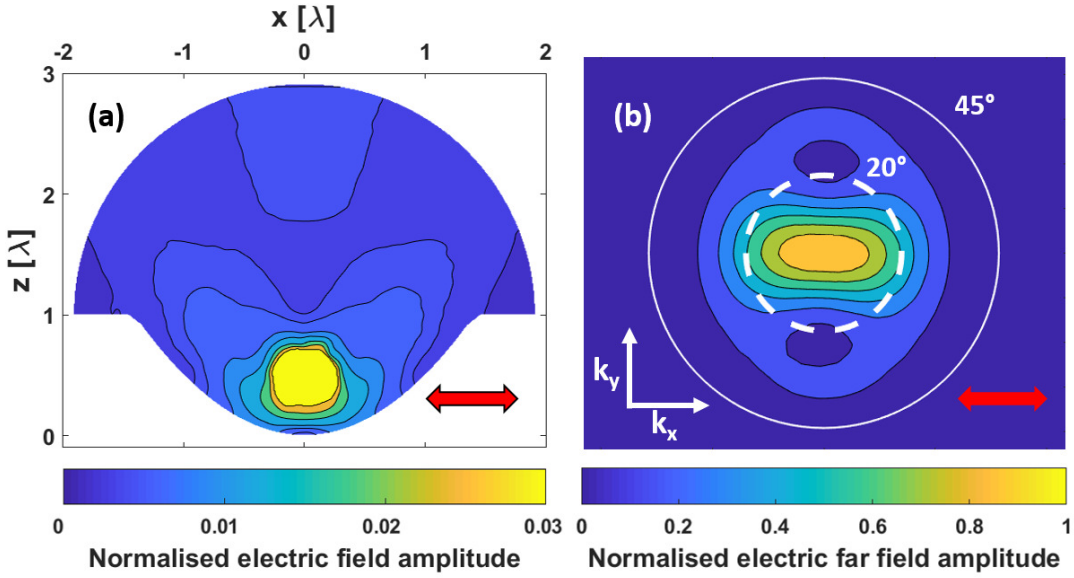


**Figure 2.8:** (a) Contour plot of the calculated electric field amplitude in the y = 0 plane, surrounding an oscillating electric point dipole within a wavelength-scale parabolic reflector. The emission wavelength $\lambda$ = 600 nm. The dipole is located at the parabola's focus (on axis at z = 0.5$\lambda$) and is oscillating in the $\hat{\boldsymbol{x}}$ direction as indicated by the red arrow. The amplitude of the field has been normalised to unity. (b) Corresponding far field contour plot, in wavevector $(k_x, k_y)$ space. The amplitude of the far field has been normalised to unity. The white circles mark emission inclination angles as labelled, with the dashed circle corresponding to the 20° half-angle cone used to measure the performance of the reflector.

for dipole moments oscillating in the vertical $\hat{z}$ (a) and horizontal $\hat{x}$ (b) directions. (c) shows the geometric mean of the solutions for the three basis dipole moment orientations ($\frac{1}{3}\hat{x} + \frac{1}{3}\hat{y} + \frac{1}{3}\hat{z}$ ). It can be seen in (c) that the maximum power radiated into the target cone is approximately 50% of the total power radiated by an equivalent emitter in free space. As one would expect, this value falls short of the ~ 75% achieved in the geometrical optics regime, which ignores diffraction. It can also be seen that whilst in the macroscale solution of Fig.2.3 only a small improvement is gained by re-positioning sources below the focal point, for the geometry chosen in Fig.2.9 the geometrical focus is one of the worst possible choices.

In the COMSOL simulations we observe that for some emitter positions as much as 10% of the total power radiated was into directions below the reflecting surface. The origin of this backwards propagating power, which cannot have passed through the PEC boundary condition is discussed in Section 2.6.5.

The maxima and minima present in Fig.2.9 (a) - (c) can be understood by examining the spatial variation in the density of optical states, shown via the changes in total emitted

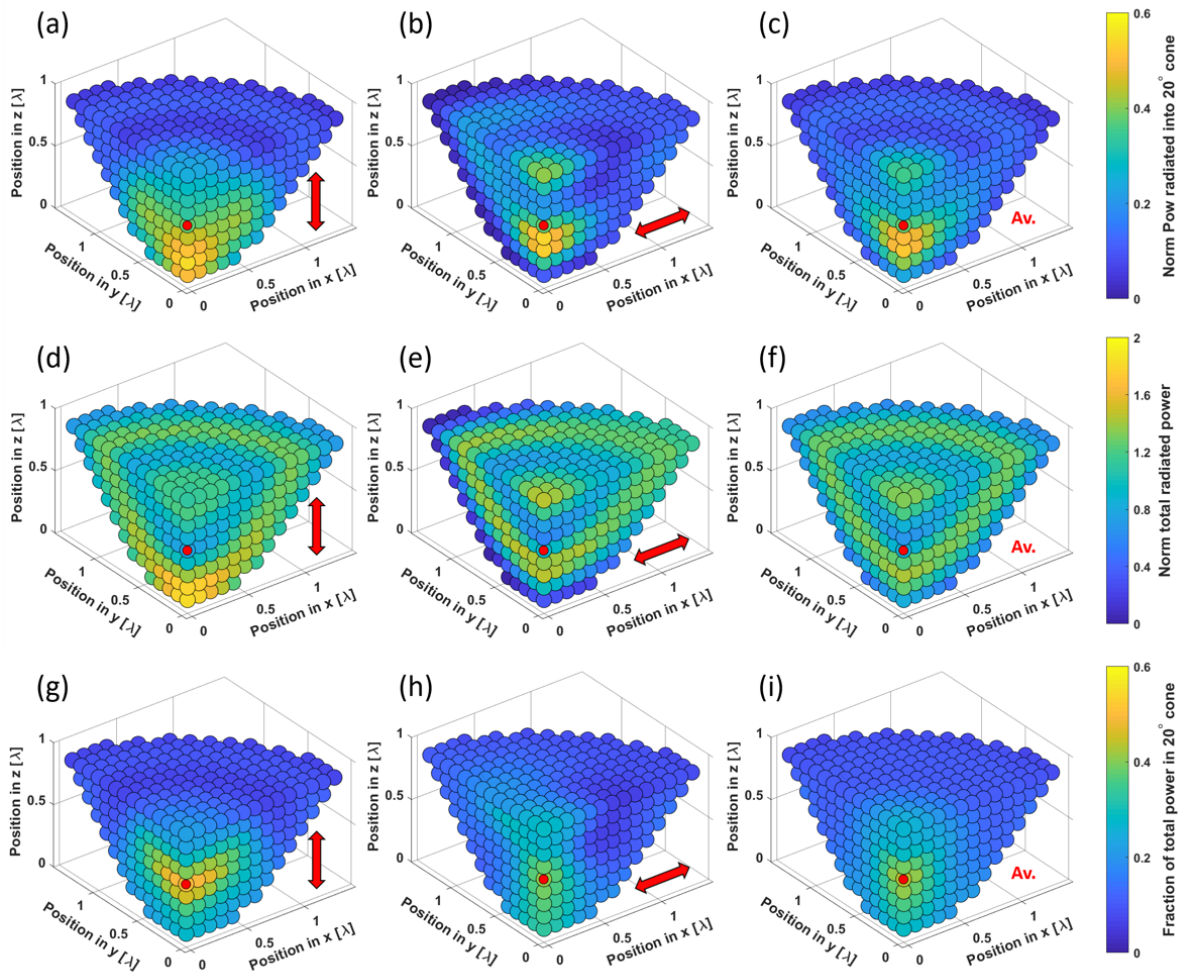


**Figure 2.9:** Directed power for dipole emitters as a function of their position within a wavelength-scale parabolic reflector (top row), corresponding total radiated power (middle row) and fraction of total radiated power within the target cone (bottom row). The emission wavelength $\lambda = 600$ nm. The dipole orientation is given by the red arrows, where Av. denotes the geometric mean of the three orthogonal dipole orientations (where the dipole oscillation in $y$ case is obtained from the oscillation in $x$ solution shown). In (a) - (c) the power directed into a 20° half-angle cone is normalised to the total power emitted in the absence of any reflector, as is the total radiated power (d) - (f). Results for dipoles orientated in the $\hat{z}$ direction and the geometric mean possess rotational symmetry about the $z$ axis whilst perpendicular orientations have symmetry about the $x = 0$ and $y = 0$ planes. Red dots indicate the geometrical focal point of the paraboloids.

power. Fig.2.9 (d) - (f) are equivalent to their top row counterparts, but display the total power emitted for each source position and orientation. It can be seen that the poor performance around the focal position in Fig.2.9 (a) - (c) can be attributed to a considerable reduction in the density of states here, compared both with its free space value and more noticeably with the elevated density of states above and below the focal position. It can also be seen in Fig.2.9 (d) and (e) that the variations in the density of states follow intuitively from the case of an oscillating dipole above an infinite ground plane [42, 44, 45, 46]. In this simplified case, consideration of the image charges induced in the mirror by the dipole provides an intuitive explanation. For a dipole oscillating immediately above and parallel to a mirror the induced image charges effectively cancel out the dipole moment, whereas for a perpendicular dipole its effective moment is increased. This behaviour is clearly seen for emitters close to the reflector in Fig.2.9 (d) and (e).

Fig.2.9 (g) - (i) show the fraction of total emitted power which is coupled into the target cone. This is the results of the top row for directed power, normalised to the total radiated power of the middle row, as opposed to the power radiated in free space. This allows one to independently see the effects of the changes in the density of states and the 'focusing ability' of the reflector, which combine to yield the directed power behaviour seen in Fig.2.9 (a) - (c). As one might expect, the focusing behaviour seen in Fig.2.9 (i) is less pronounced and less spatially localised than that seen in the geometrical optics regime of Fig.2.3.

Fig.2.9 shows that at this scale the local density of states has a strong bearing on both the optimum emitter position and the corresponding value for the directed power. As the size of the reflector is varied, so too will be the local density of states. Fig.2.10a shows the amount of power directed into a 20° half-angle cone for emitter positions along the central axis of the reflector, for a range of emission wavelengths, effectively varying the reflector size. The geometrical optics solution is present at the lambda = 0 position, for comparison. It can be seen that as the effective size of the antenna decreases, maxima in the directed power broaden and decrease in number. The expected trend of a general decrease in directed power with decreasing effective size is also observed, but has its exceptions as it competes with the changing location of maxima and minima in the local density of states. Fig.2.10b shows the variation of the total emitted power.

The impact of the density of states on the directed power can clearly be seen by comparison of Fig.2.10b and Fig.2.10a. For a $\lambda : \lambda_0$ ratio of 4/6, the directed power can be seen to exceed the optimal value obtained in the geometrical optics regime. This occurs for an emitter positioned 7/10 of the way from paraboloid vertex to the opening, over half an emission wavelength above the optimum inferred from geometrical optics.

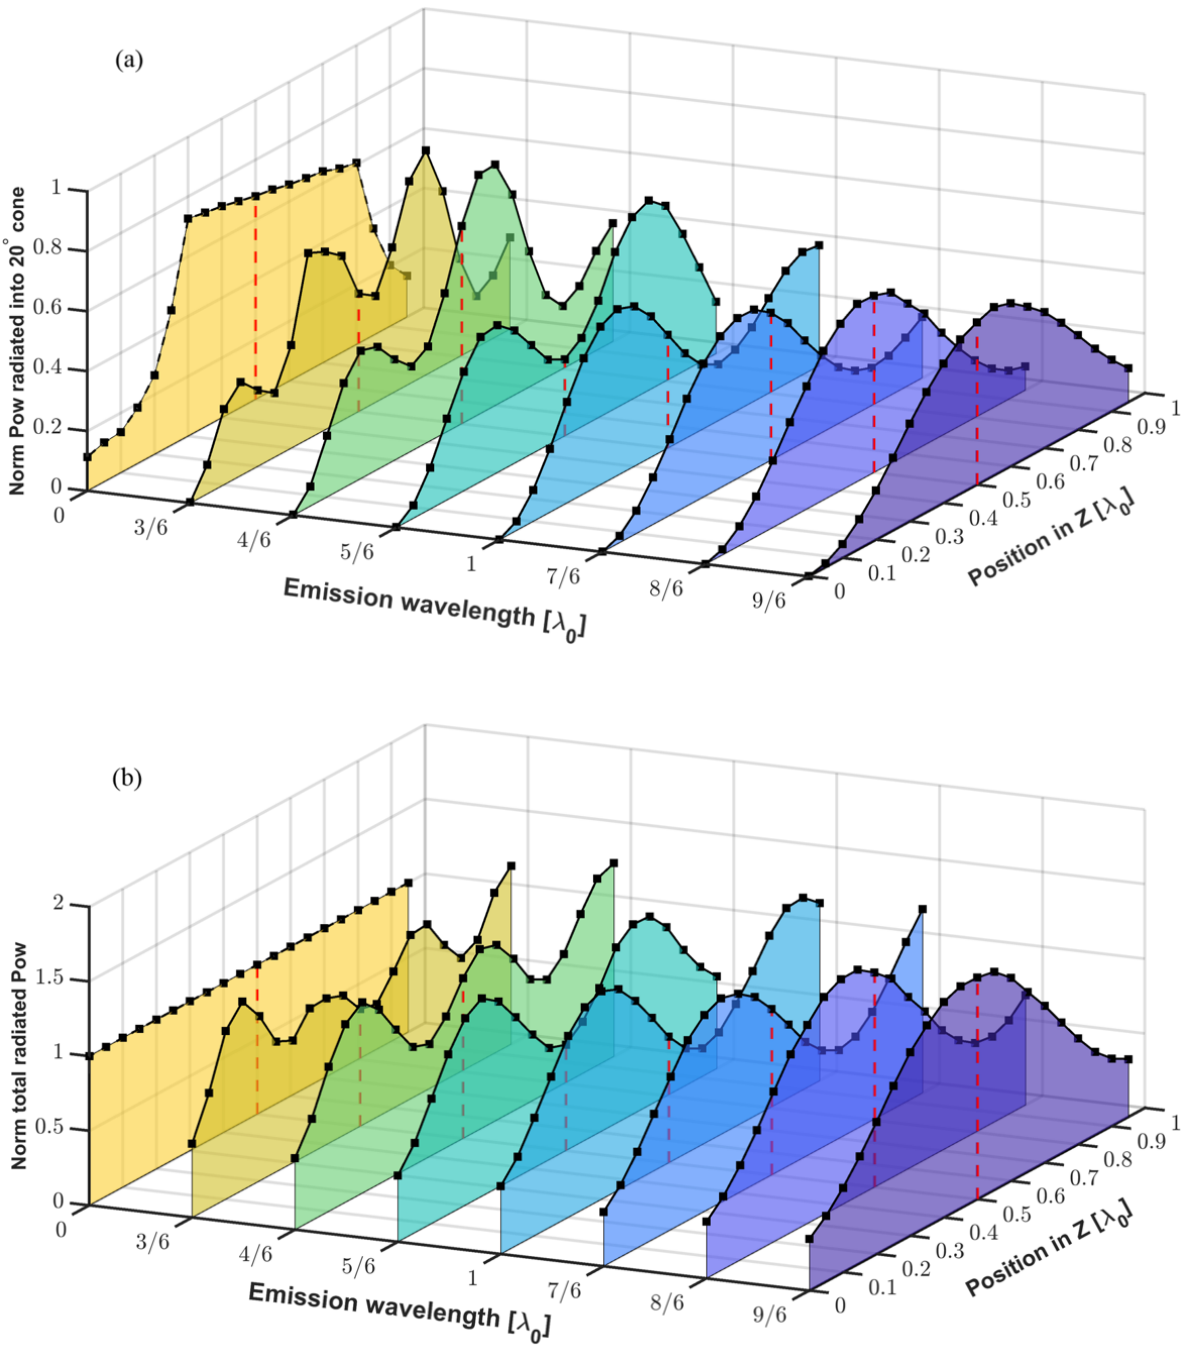


**Figure 2.10:** (a) Variation in the power directed into a 20° half-angle cone with emitter position along the reflector's central axis and with emission wavelength. The emission wavelength is given in units of the design wavelength $\lambda_0 = 600$ nm, which determines the (fixed) size of the reflector. The size and shape of the reflector remains as in Fig.2.9, with a vertex to opening distance of $\lambda_0 = 600$ nm. Directed power values are normalised to the total power radiated by an equivalent emitter (at the specified emission wavelength) in free space. The black squares indicate values obtained by simulation and the connecting black lines are visual guides. The dashed red lines indicate the position of the geometrical focus. The '$\lambda = 0$' values represent the solution in the small wavelength limit in which the geometrical optics approximation may accurately be applied and the solution does not depend on wavelength. (b) As in (a), but showing the total radiated power.

### 2.5.1 In the Context of Other Work

The application of parabolic reflectors to nanoscale sources of light is far from unique to this work. Some examples of which would be the substitution of the conventional objective lens with a parabolic mirror in confocal microscopy [47, 48, 49] and for enhancing the collection efficiency from a single fluorescent molecule or ion [50, 51, 52]. These are examples of geometrical optics regime reflectors applied to nanoscale sources of light, but there have also been several recent works where the reflector size itself has been reduced to of the order of the emission wavelength.

In [53] the challenge of fabricating a micro / wavelength-scale paraboloid is circumvented by using using plasmonics to allow a parabolic antenna to be realised from a planar elliptical nanometallic structure. Fabrication of 3D paraboloids of wavelength-scale dimensions has been demonstrated by reactive ion etching in diamond [54] and also by exploiting parabolic iso-intensity contours of a Gaussian beam focused near a photoresist [55, 56]. In the latter works, which include simulation of a range of paraboloid sizes, the sizes of interest are typically just off the larger end of those considered here, with respect to the emission wavelength in its host medium.

### 2.5.2 Conclusions

It has been shown through numerical calculation that for parabolic antennas whose characteristic size is approximately the emission wavelength, the optimum emitter position deviates considerably from the paraboloid's focus. An informed choice on placing the emitter allows the power emitted into a cone to exceed the optimum value for the geometrical optics solution. In contrast, choosing to position the emitter at the focus may result in less than 50% of the directed power expected from geometrical optics. Fig.2.10a shows that the minimum size for such a reflector at which one can expect directed power comparable to the geometrical optics solution, is just above one wavelength (vertex to opening distance). Below this size diffraction causes the reflector to be of decreasing utility. With approximate size limits established, other factors impacting the applicability of wavelength-scale parabolic antennas to directional phosphor components are examined in Chapter 3.

If wavelength-scale parabolic antennas were concluded to warrant further investigation for application in directional phosphors, one would need to include the phosphor - air dielectric interface and reflective properties of an imperfect real metal in the simulations. Including the refractive index of the phosphor would change the effective size of the reflectors in proportion to its index and perturb the emission patterns and PLDOS variations accord-

ing to the geometry of the new dielectric interface. Antenna efficiency would be reduced by the lossy reflection from a real metal and quenching of emission would be expected in close proximity to the reflector surface [44].

## 2.6 Methods

### 2.6.1 A Note on Meshes

An auto-generated 'Physics-controlled mesh' ([57]) has been used for both the geometrical optics and wavelength-scale parabola models. The Physics controlled mesh is described as follows in the COMSOL reference manual:

*"A physics-controlled meshing sequence examines the physics to automatically determine size attributes and sequence operations needed to create a mesh adapted to the problem. The physics-controlled sequence is based on heuristics and knowledge built-in by application experts." - [57]*

As this is unlikely to be of much utility to someone without access to the COMSOL software, we have included additional information on the generated meshes. For each model two visualisations of the mesh are included, one colour-coded to emphasize the element sizes and the other the element quality.

The element size is defined as the length of the longest edge of the element [57]. The element quality $Q$ is quantified using the default measure of the 'skewness' of a mesh element with respect to a perfectly regular element.

$$Q = \min\left(1 - \max\left(\frac{\theta - \theta_e}{180 - \theta_e}, \frac{\theta_e - \theta}{\theta_e}\right)\right) \tag{2.14}$$

"Where $\theta$ is the angle over a vertex (2D) or edge (3D) in the element, $\theta_e$ is the angle of the corresponding edge or vertex in an ideal element, and the minimum is taken over all vertices (2D) or edges (3D) of the element" - [57]. The mesh element quality ranges from 0 to 1 where 0 represents a degenerated element and 1 the ideal regular element in terms of 'skewness' [57].

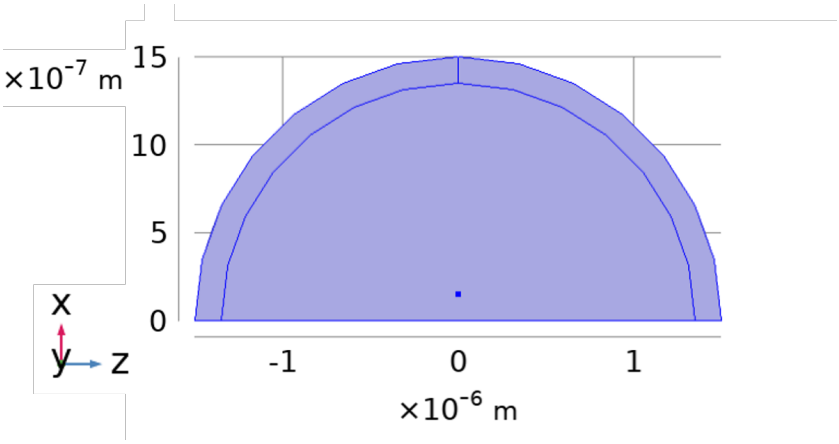


**Figure 2.11:** Cross-section in the $xz$ plane of the simulation geometry for the dipole in air above mirror results of Fig.2.6. The simulation wavelength was 300 nm and here the dipole (black dot) is displaced $0.5\lambda$ from the mirror.

### 2.6.2 Dipole Above a Mirror COMSOL Model

A frequency domain study within COMSOL's Radiofrequency Electromagnetic Waves, Frequency Domain module was used to produce the result in Fig.2.6. The geometry for this simulation is shown in Fig.2.11. The size of the finite mirror was maximised by simulating only a hemispherical domain with a large radius of 5 × the emission wavelength $\lambda$. With reference to Fig.2.11, the horizontal surface at $x = 0$ was defined with the Perfect Electric Conductor (PEC) boundary condition, which simulates a perfect mirror. The PEC condition imposes that the tangential component of the electric field ($\boldsymbol{E}$) at the boundary must be zero, $\boldsymbol{n} \times \boldsymbol{E} = 0$, where $\boldsymbol{n}$ is the surface normal. In addition Perfectly Matched Layers (PMLs) were defined between the inner (at radius $4.5\lambda$) and outer curved surfaces as an absorbing boundary condition. The outer curved surface was thus also given the PEC condition.

The oscillating point dipole, with orientation parallel or perpendicular to the mirror was swept from 0 to $1.6\lambda$ in increments of $0.1\lambda$. A small offset of $0.0001\lambda$ in the positive $\hat{\boldsymbol{x}}$ direction was applied to allow a solution at 'zero' displacement from the mirror. The total radiated power was determined by integrating the time averaged power outflow, by variable 'emw.nPoav' over the surface of the inner hemisphere.

A physics controlled 'fine' mesh was used, as shown in Fig.2.12. Note the anisotropic mesh applied within the PMLs by the physics controlled mesh, in line with the anisotropic material properties of the PMLs.

### 2.6.3 Geometrical Optics Parabolic Reflector Method

For the case of the macroscale paraboloids (size >> wavelength), radiation patterns were obtained using COMSOL's Geometrical Optics Module (Ray Tracing study). In a three-dimensional simulation, isotropic emission (1000 rays) was released from a point within a finite-height reflecting paraboloid defined by the equation $z = (x^2 + y^2)/2$, with $x \in (-\sqrt{2}, \sqrt{2})$, $y \in (-\sqrt{2}, \sqrt{2})$, and $z \in (0, 1)$, with a perfect reflection boundary condition. The absolute size of the described shape is of no real consequence in the geometrical optics regime, here the vertex to opening height ($y_0$ in Fig.2.3) was set to 1 m as can be seen in Fig.2.13. The time stepping used for the solver was set to 0 to 50 ns in 0.5 ns increments.

The simulation domain was closed with a hemispherical boundary over the open portion of the paraboloid. The 'freeze' boundary condition was applied here to terminate rays propagating into the far field. Wave vectors from rays frozen on this boundary were then used to obtain a far field radiation pattern. Exploiting the radial symmetry present, this

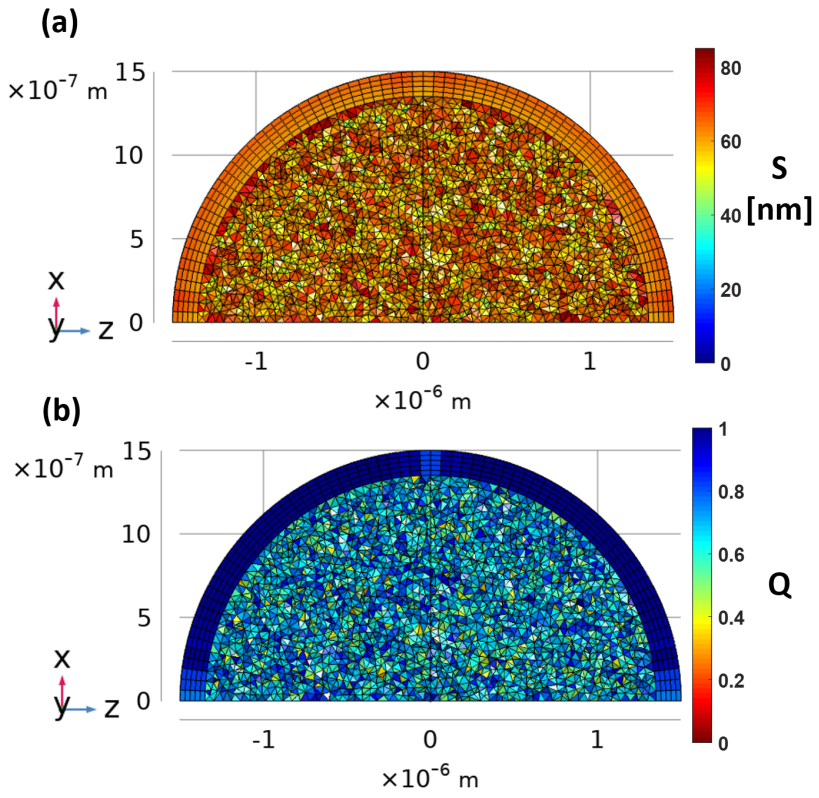


**Figure 2.12:** Cross-sectional view of the fine physics controlled mesh used for the results of Fig.2.6. The parameters **S** and **Q** show the element size and quality as defined in Section 2.6.1. The elements shown are volume mesh elements satisfying the condition $y > 0$.

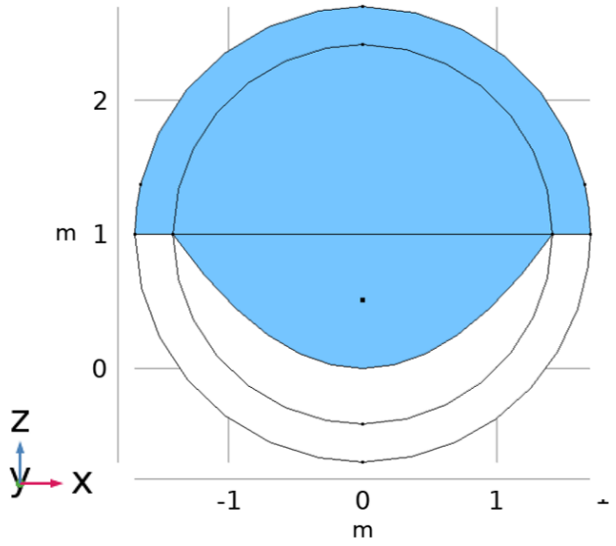


**Figure 2.13:** Cross-section of the 3D reflector geometry, in the $y = 0$ plane. The shaded blue regions show those that are in the simulation domain and filled with air ($n = 1$), the white regions lie below the reflection boundary condition and need not be simulated.

process was repeated for a regular two-dimensional grid of source positions, sampling the area within the parabola on one side of its central axis of symmetry. Emitter positions were shifted in the positive $\hat{z}$ direction by 0.01 m from their natural integer multiples of 0.1 m to avoid the variable source positions coinciding with any 'naturally occurring' points in the COMSOL geometry. In the event of such a definition clash, for example at a height of 0 or 1 m, COMSOL can loose track of the point which should be defined as the dipole.

A 'Normal' physics-controlled mesh, generated internally by COMSOL was used throughout. An example of the meshes used can be seen in Fig.2.14, where both the element size and quality are as defined in Section 2.6.1.

Post-processing of the obtained ray angular distributions was performed in MATLAB. The number of rays satisfying the 20° half-angle emission cone was divided by the total number to obtain the directed power.

### 2.6.4 Wavelength-scale Parabolic Reflector Method

For the wavelength-scale paraboloids a frequency domain study within COMSOL's Radiofrequency Electromagnetic Waves, Frequency Domain module was used. The position and orientation of an oscillating electric point dipole was swept within a reflective paraboloid of revolution with the same basic shape as in the macroscale case. For a dipole moment perpendicular to the central axis of the reflector, a three dimensional sweep of one quad-

rant of the reflector is required, whilst for the case of a dipole moment oscillating parallel to this axis, a two-dimensional sweep over $r = \sqrt{x^2 + y^2}$ and $z$ is sufficient. In Fig.2.9 the $z$ displacements cover the range 0.11 to $0.91\lambda$ in increments of $0.1\lambda$, whereas for the fine central axis sweeps of Figs.2.10a and 2.10b the range was enhanced to 0.01 to $0.96\lambda_0$ in $0.05\lambda_0$ increments.

The paraboloid was designed to have a height of $\lambda_0$ and a radius at its opening of $\sqrt{2}\lambda_0$, with a design wavelength $\lambda_0 = 600$ nm. To produce the results of Fig.2.10a the emission wavelength $\lambda$ was varied between $0.5\lambda_0$ and $1.5\lambda_0$ . As COMSOL employs the Stratton-Chu method [58, 59, 60, 61] to transform from the near to the far-field, it is necessary to enclose the whole system in a closed surface with uniform properties. The reflecting parabolic surface was continued by a horizontal annulus of width $0.2\lambda_0$, which formed a closed surface with a hemisphere below the paraboloid. The perfect electric conductor boundary condition was applied to this entire surface. This allows an additional spherical air layer, of annular radius $0.3\lambda_0$ to act both as the homogeneous boundary for near to far field transformation, and a second order scattering boundary condition (which effectively releases radiation from the simulation domain) on its outer surface. A 'normal' physics-controlled mesh was used throughout, an example of which can be seen in Fig.2.16. More advanced meshing strategies were explored, such as adaptive mesh refinement, but the marginal improvements observed were far outweighed by computational cost of running several iterations of refinement for each dipole position.

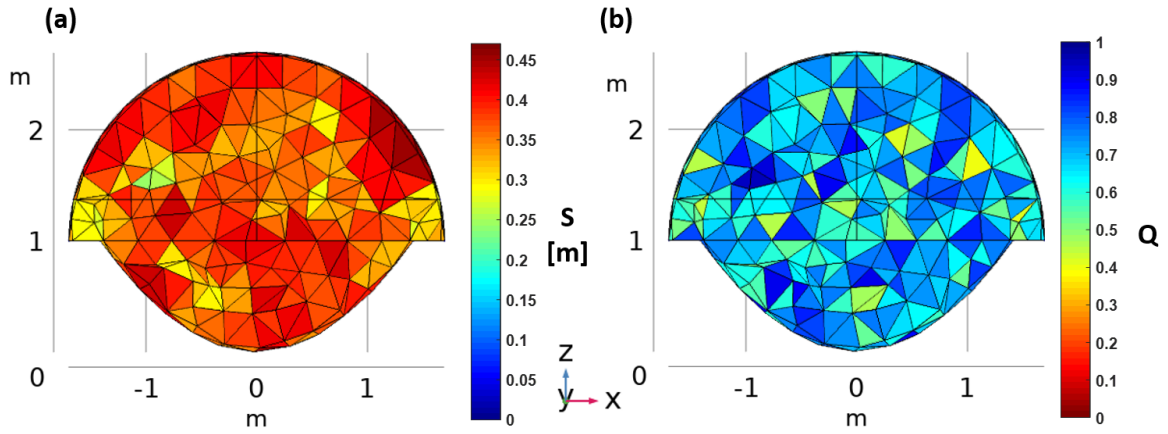


**Figure 2.14:** Cross-sectional view of the physics controlled mesh used for the results of Fig.2.3, when the point source of emission was located at 0.5 m. The parameters **S** and **Q** show the element size and quality as defined in Section 2.6.1. The elements shown are volume mesh elements satisfying the condition $y > -0.15$ m.

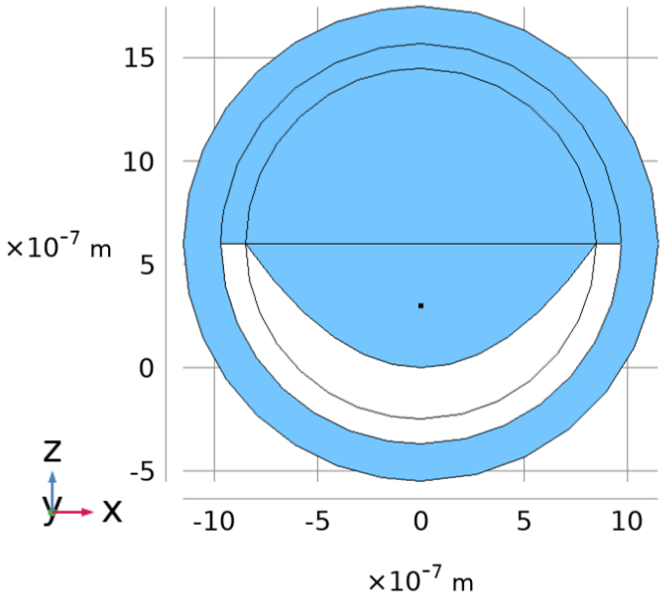


**Figure 2.15:** Cross-section of the 3D wavelength-scale reflector geometry used in Fig.2.9, where $\lambda_0 = \lambda = 600$ nm, in the $y = 0$ plane. The shaded blue regions show those that are in the simulation domain and filled with air ($n = 1$), which here must fully enclose the emitter. The white regions are excluded from the simulation and modelled with the Perfect Electric Conductor boundary condition.

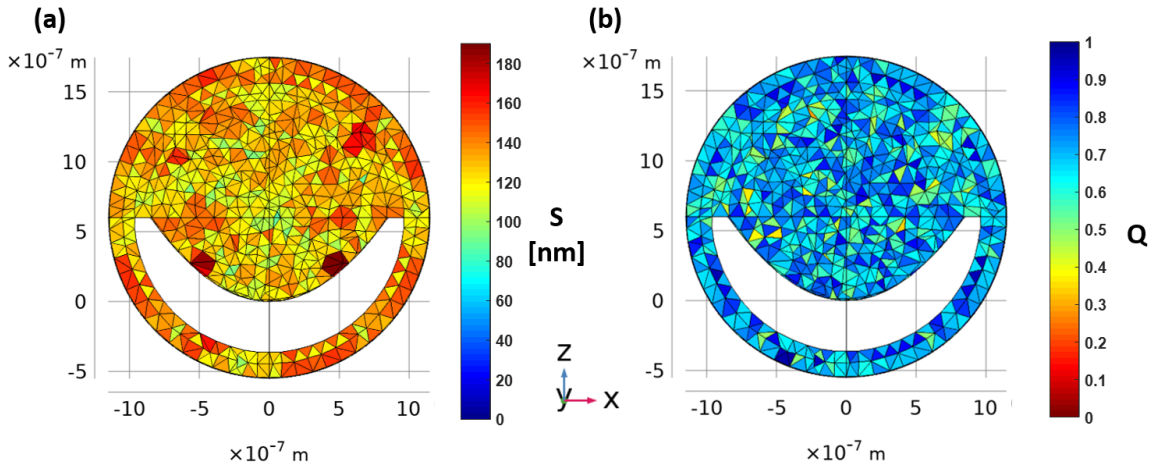


**Figure 2.16:** Cross-sectional view of the physics controlled mesh used for the results of Fig.2.9, when the point dipole was located at $0.5\lambda$. $\lambda_0 = \lambda = 600$ nm. The parameters **S** and **Q** show the element size and quality as defined in Section 2.6.1. The elements shown are volume mesh elements satisfying the condition $y > -0.1\lambda$.

3D Far-field plots were calculated within COMSOL with angular resolution settings - 60 points in elevation $\theta$ and 40 points in azimuth $\phi$, covering the full emission sphere of $4\pi$ steradians. The variable evaluated within COMSOL was the square of the electric far field norm 'emw.normEfar$^2$', representing the power density. Far-field data is exported from COMSOL as a set of Cartesian coordinates. These are then converted to spherical polar coordinates. The number of points per unique $\theta$, $\phi$ combination (excluding the $\theta = 0$ singularity where 40 points are expected) is frequently and seemingly randomly greater than one, for reasons unclear. To resolve this the spherical polar dataset is interpolated using the MATLAB function 'scattered interpolant' with default settings, this removes the duplicate data points.

The power within a given angular range is calculated as the numerical integration of the radial coordinate (which was already squared within COMSOL), weighted by the sine of the $\theta$ value, for all data points within the specified $\theta$ range. Due to the typical representation of Far-field amplitude/intensity as the radial coordinate in a classic Far-field plot, care must be taken when performing the integration. In reality one is integrating a function (the far-field amplitude/intensity) over the surface of large (in the far-field) sphere, of constant and irrelevant radius at all points. One could erroneously include the far-field amplitude/intensity as the $r^2$ factor in the spherical polar surface area element of integration $r^2 \sin\theta \mathrm{d}\theta \mathrm{d}\phi$. This process was repeated for the results of Fig.2.10a and Fig.2.10b with varying calculation wavelength $\lambda$ and fixed size dictated by the design wavelength $\lambda_0$. For each value of $\lambda$ the total emitted power in free space was determined in a simulation as above but in which the metal regions were included and modelled as air.

### 2.6.5 A Note on Radiation 'Through' the Metal

Fig.2.17 shows an example of a radiation pattern obtained for an emitter within the parabolic reflector defined in Fig.2.15. A small portion of the power can be seen to be radiating in directions that would correspond to emission through (or around) the perfectly reflecting boundary. One could be forgiven for writing off this minor contribution as some negligible error, however, they would have fallen victim to this deceptive yet widely adopted way of representing a radiation pattern. The seemingly minor back lobe(s) in Fig.2.17 actually correspond to $\sim 7\%$ of the total power radiated by this emitter. One should instead realise those low far field values correspond to an entire hemisphere of emission angles and their integrated contribution is larger than instinct would have you believe.

The origin of this 'backwards power' was the focus of our investigations for some time.

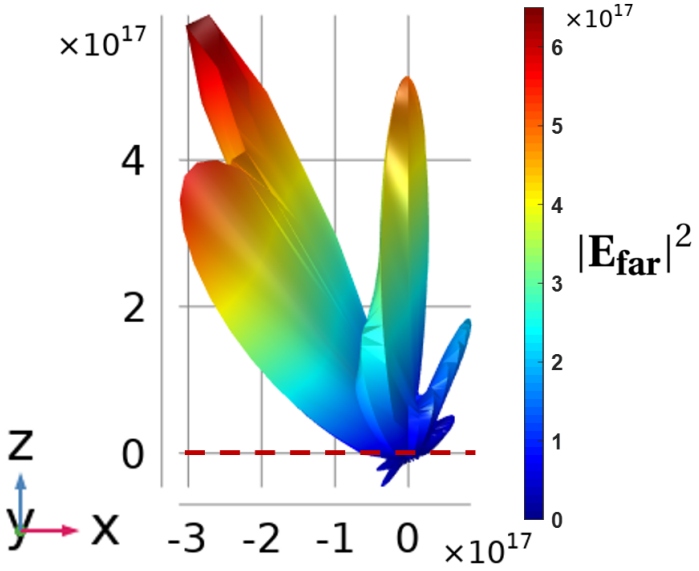


**Figure 2.17:** Far field radiation pattern for a dipole oscillating in the $\hat{\boldsymbol{x}}$ direction, positioned at $x = 1.3\lambda$, $y = 0$ and $z = 0.95\lambda$ inside the reflector of Figs.2.15 and 2.9. In this visualisation of the radiation pattern, the squared amplitude of the electric far field in a given direction is provided by the distance from the origin. The dashed line at $z = 0$ separates the expected forward emission of the antenna from that which appears to be radiated through the perfect mirror.

Of course in the geometrical optics regime emission in this direction couldn't be permitted, even with the truncated metallic surround required in the wavelength-scale geometry (Fig.2.15), it would require rays to bend in air. The answer then appears obvious, we are not in the geometrical optics regime and the emitter of Fig.2.17 is fairly close to the outer edge of the horizontal reflecting annulus, it is entirely reasonable to expect some power to be diffracted in the 'backwards' direction. The perceived difficulty with this explanation was that the backwards radiation component still seemed too large for emitters closer to the paraboloid central axis and was diminishing too slowly with increasing width of the reflecting annulus.

The solution to this puzzle came from the discovery that we may have been unjustified in our expectation of 'fast' convergence to the expected result for an infinite reflecting surround, with increasing surround size. In [62] we see that for the related system of a monopole antenna above a finite ground plane the convergence to the result for an infinite ground plane is surprisingly slow. This slow convergence can be seen in Fig.2.18, where the peak angle of the radiation pattern tends very slowly to the infinite case of 90°, with increasing ground plane radius.

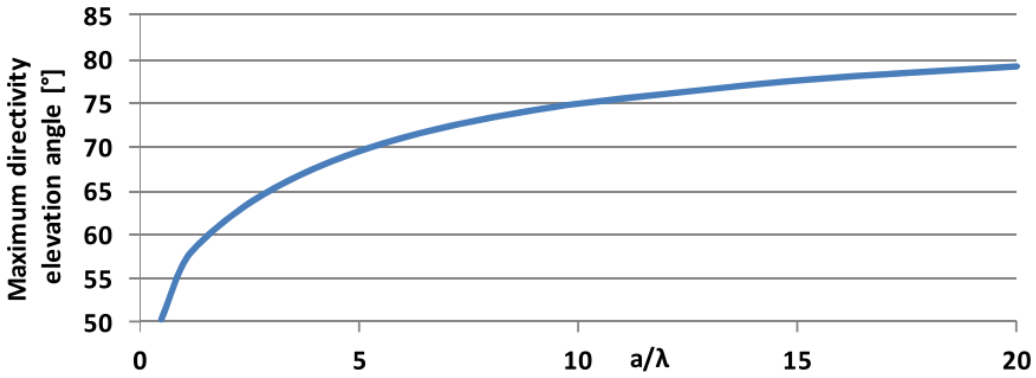


**Figure 2.18:** Figure from [62], showing slow convergence of the radiation pattern of a quarter wavelength monopole antenna above a finite ground plane of radius $a$ to that of an infinite ground plane. Convergence is tracked here via the peak emission angle, which is 90° in the perfect infinite ground plane case. $\lambda$ is the emission wavelength.

Even at a radius of 10 $\lambda$ the emission maximum is significantly raised from the horizontal and a substantial portion of the radiated power is into directions below the finite ground plane [62]. In light of this behaviour it seems reasonable to conclude the backwards radiation we observe is physical and can be explained by the finite reflecting annulus.

# 3 Choosing the Right Approach

## 3.1 Assessing Strategy Feasibility

We have seen in Section 2 that one can obtain directed power values comparable to those achieved in the geometrical optics regime by using parabolic reflectors whose characteristic size is just above the wavelength of emission. However, we have not yet evaluated the wavelength-scale reflector array approach against the full functional requirements outlined in Section 1.4.1, nor have we compared it to other strategies for producing directional fluorescence. This is the subject of this chapter, in which we explain the rationale behind our final choice of strategy to pursue in full.

To date, the strategies for generating directional fluorescence with imaging capability we have considered fall somewhere within one or more of the following categories (those in **bold purple font** indicate techniques that rely on a thin [roughly hundreds of nanometers or less] phosphor layer :

- Reflectors
  - **Wavelength-scale Parabolas**
  - Microscale Parabolas
  - **Nanoscale Pyramidal Reflectors**
  - Microscale Pyramidal Reflectors
- Refractors (lenses)
  - Microlens Arrays
- Nanoantennas
  - **Coupling to extended lattice resonances**
  - **Near-field coupling to metal / dielectric antennas**
  - **Metasurfaces (active)**
- Waveguides
  - **Nano-waveguides**
  - Diffractive (grating) outcoupling
  - Macroscale fibre optics

We shall discuss each of these briefly in turn, providing representative examples of works in each category where possible and highlighting some of the relevant characteristics of the

underlying physical mechanisms. Then, we perform some basic calculations in Sections 3.2 and 3.3 to help inform our final decision on which to proceed with, which we make in Section 3.4.

**Reflectors**

We begin with the category of reflectors, following on directly from the work in Section 2. This is the widely adopted principle for enhancing the directional efficiencies of most commercial white LEDs (see Fig.2.1). We've seen that the parabolic reflector can maintain its directional capabilities even when miniaturised to the wavelength-scale [55, 56, 63]. One consequence of designing a phosphor component at this scale is one can expect dispersion in the angular emission pattern, as the broad spectral coverage inherent to the white LED will produce differing radiation patterns for each of its component wavelengths (see Fig.2.10a). This lack of angular colour uniformity could prove problematic for the intended lighting applications.

As the reflector size increases through the macroscale these interference, diffraction and density of states induced spectral dependencies become less problematic and one can quickly find the geometrical optics picture of Fig.2.3 to be accurate. Thus there exists a wide range of reflector sizes, from a few microns up to the $\sim 1$ cm width limit, which could present a viable solution for this work. One would however still need to take care to manage diffractive effects from the effective grating that forms when the array pitch is at the lower end of microscale.

We include pyramidal reflectors of both nanometer and microscale in this category as ease of manufacturing could warrant deviation from the ideal parabolic shape. Through anisotropic etching of the crystal planes in silicon, pyramidal reliefs can be formed over large area by the intersecting crystallographic planes [64]. These structures find application in light trapping in solar cells [65, 66] but could also be employed to provide directional emission to phosphors emitting near the inverted pyramid's apex.

**Refractors**

We include this category primarily for the sake of completeness. To meet the imaging requirements outlined in Section 1.4.1, refracting elements would likely need to be deployed in an array much as for the reflectors discussed above. A microlens array placed in contact with a continuous phosphor layer could provide enhanced outcoupling and some directionality to its emission. The efficiency of emission into a narrow solid-angle is likely to be hindered by the uniform spatial distribution of the emitters, which will in general not be at

the focus of one of the lenses. Consequently a lens array might be better used in conjunction with a reflector array, where the emitting areas are more spatially localised.

**Nanoantennas**

The control of the emission direction, lifetime and spectrum can be achieved by altering the local electromagnetic environment of an emitter, as we have seen for parabolic antennas. Optical analogs of familiar radio-frequency antenna designs can be implemented by coupling emitters to metallic or dielectric nanostructures that support favourable resonant modes [67, 68]. This is typically achieved via near-field interaction with the resonant structure, which limits the range over which the emission properties can be modified [67, 68]. Efficient overlapping of the emitters with the maxima of the resonant modes enables impressive photoluminescence enhancement factors of the order $10^4$ in some cases [69].

The spatial extent of the resonant modes can be extended considerably in periodic arrays of metallic nanoparticles, making nanoantennas more appealing for solid-state lighting applications [70]. These extended surface lattice resonances enable interaction with a much greater fraction of the bulk phosphor, extending hundreds of nanometers in the direction perpendicular to the array plane ([70, 71]) and can beam fluorescence emission into narrow solid angles as each metal nanoparticle behaves as a coherent emitting element in a phased array [71]. The extended spatial character of these modes results in spectrally narrow resonances when compared to their localized counterparts [70, 71, 72]. This could prove problematic for spectrally matching enhancements to the broad emission from YAG:$Ce^{3+}$ and such nannoantenna arrays would also see the beaming direction vary with wavelength, compromising colour angular uniformity.

Metasurfaces (2D metamaterials) achieve material properties beyond those accessible to their constituent parts through the combined effect of subwavelength structures (or meta-atoms) and their positioning in space. Metasurfaces can generally be described as either passive or active [73]. Most often they are passive in the sense that they are tailored to control the propagation of waves whereas in the newer class of active metasurfaces the emission may be engineered via direct incorporation of the emitter into the metasurface itself, becoming light-emitting [73]. The greater degree of control offered by active metasurfaces makes them better suited to lighting applications than their passive counterparts.

Clearly there is substantial overlap between the between the concepts of light-emitting metasurfaces and the nanoantenna arrays discussed above, with the latter considered an example of the former [73]. Therefore the characteristics and anticipated challenges remain consistent with those already mentioned.

**Waveguides**

Due to the high refractive index of YAG:$Ce^{3+}$, waveguiding via total internal reflection occurs naturally in its scatter-free forms such as monocrystalline YAG:$Ce^{3+}$. One simple method of generating directional emission would be to allow directional outcoupling of the guided fluorescence with a grating structure. In this system the momentum needed to outcouple from the guided modes to produce radiative modes is supplied by scattering from the periodic grating. This results in a well defined emission angle which varies with the emission wavelength. Again we expect an undesirable spectral dependence on the emission angle.

If the size of the waveguide is reduced to a few hundred nanometers so that the number of modes supported by the guide becomes small, one can encourage preferential emission into the guided modes of the structure [74]. Furthermore the concepts of enhanced emission into confined guided modes and directional outcoupling through periodic structuring have been combined to produce resonantly enhanced emission from a luminescent nanostructured waveguide [75, 76]. In [75] a 200 nm thick YAG:$Ce^{3+}$ waveguide is periodically nanostructured to achieve up to a 30-fold enhancement in the emission into a well defined angle, vs an unstructured guide. The enhancements stem from a combination of the elevated local density of states and directional outcoupling. The enhancements are obtained in narrow (FWHM<20 nm), polarization dependent spectral regions and exhibit the beaming angle spectral dependence characteristic of grating outcoupling.

Finally, we consider traditional geometrical optics regime fibres emitting from their end facets. Directional emission from familiar fibre optic patch cables prompted interest in the structures, but emission properties for fluorescence generated within the fibre are unknown. If an acceptable fraction of the emission is guided to the fibre ends then the small emitting area could provide a bright, flexible and low etendue illumination component. In addition, in the geometrical optics regime one can expect reduced dispersive effects.

## 3.2 Minimum Amount of Phosphor Calculation

With the rough size constraints and luminous output target defined in the Functional Requirements Section 1.4.1, we can begin to predict the suitability of some of the techniques just discussed. The first question we sought to answer is: could a nano-film phosphor meet the 20,000 lumen white output target within an area no larger than a shoe box (~30×20 cm)? If the answer is no, half of the considered approaches in Section 3.1, those colour-coded in purple, can be discounted.

To answer this question we estimate the minimum area of a YAG:$Ce^{3+}$ layer only 100 nm thick needed to produce 20,000 lm of output, subject to a number of approximations. These simplifications reduce the complexity of the calculation to match the degree of precision needed at this stage and enable the calculation to be clearly conveyed in just a few lines of simple mathematics.

We begin with the 20,000 lm target, which we convert to the necessary pump light power with a typical white light phosphor yield of 200 lumens per watt of pump [77]. This is also the target value for the sought phosphor component. This tells us we would need to illuminate the phosphor with 100 W of blue pump.

Having removed lumens from the problem, we convert the pump power to a number of fluorescence photons emitted per second with the following two approximations for simplicity. First, we know the internal quantum efficiency of fluorescence is approaching unity [18], we consider it and the external quantum efficiency to be 100%. Second, we assume all of the fluorescence is emitted at a wavelength ($\lambda$) of 580 nm, roughly at the intensity weighted center of its spectrum (see Fig.1.4). The number of 'yellow' photons needed per second is then dictated by the change in energy from the 450 nm blue pump to the 580 nm fluorescence and the energy of a single yellow photon. Both values governed by the photon energy at a given wavelength

$$E = \frac{hc}{\lambda}, \tag{3.1}$$

where $h$ is the Planck constant and $c$ the speed of light in vacuum. From eq.3.1 we estimate the necessary fluorescence yield to be 78 W or $2.3\times10^{20}$ yellow photons per second.

Now we make our most significant approximation by neglecting the excitation statistics and assuming the fluorescent ions spend no time in the ground state, in other words that they are immediately re-excited after emission. This is the assumption that makes this a minimum area needed calculation. With a fluorescence lifetime of just 70 ns [12], each ion can emit up to $1.4 \times 10^7$ emitters per second and therefore the minimum number of emitters needed in the phosphor is $1.6 \times 10^{13}$.

Finally, we take a value of 2.4 nm as the lowest inter-dopant spacing ($d$) achievable in monocystalline YAG:$Ce^{3+}$ [18]. The dopant density $\rho = 1/d^3$, which with the number of emitters and chosen film thickness of 100 nm, determines the minimum area of phosphor to be just 2.2 $mm^2$. As this is several orders of magnitude below 'shoe box size', we are thus not able to rule out techniques using a nanoscale phosphor layer based upon this simplified calculation.

To obtain a more accurate prediction of the required phosphor area, one could relax the assumption that the emitters are constantly re-excited. This is expected to yield a much

higher minimum phosphor volume requirement, as the pump power ($P$) necessary to reach the saturation condition assumed in the calculations above is estimated to be ~ 45 kW. This is determined from the pump photon energy ($E$), calculated minimum phosphor area ($A$), fluorescence lifetime ($\tau$) and absorption cross-section ($\sigma$) as $P = \frac{AE}{\tau\sigma}$. At several orders of magnitude higher than the necessary 78 W pump power needed in this application, factoring in the excitation dynamics could rule out a nanoscale approach. However, as one begins to consider the probability of exciting an activator ion to an excited state, it becomes clear that the fractional pump absorption by the layer could be a limiting factor in achieving white output. This is the topic of following section.

## 3.3 Phosphor Pump Absorption

It has been seen in the preceding section that the potential to use a nanoscale phosphor layer approach has not been ruled out based on the maximum light output from the number of emitters that could be contained within a thin (100 nm) phosphor film of practical cross sectional area. However, the probability of absorption of pump light by an activator ion in the ground state occurring has not yet been accounted for. The absorption of light by a material of uniform concentration is related to the path length of light through the material by the Beer-Lambert Law [78, 79]

$$\frac{I}{I_0} = e^{-\alpha x} \tag{3.2}$$

where $I/I_0$ is the ratio of light intensity at distance $x$ along the optical path to its starting intensity and $\alpha$ is the absorption (or more generally attenuation) coefficient.

We consider the monocrystalline disc which produced the cool-white (CCT ~ 7500 K) emission spectrum shown previously in Fig.1.4 as a point of reference. Its absorption coefficient $\alpha$ =75 cm$^{-1}$ for a pump wavelength of 460 nm, as provided by the manufacturer (Crytur spol.s.r.o. [80]). Its dopant concentration is estimated to be just under 0.2 at.% based on the experimentally determined absorption cross-section of $\sigma = 3 \times 10^{-18}$ cm$^2$ for the (460 nm) $^2F_{5/2} \rightarrow ^2 D_{3/2}$ transition in [18]. The % pump absorption achieved within a given thickness of this phosphor is shown in Fig.3.1.

The thickness of the monocrystalline disc that produced the white output spectrum of Fig.1.4 under blue LED illumination was ~ 180 $\mu$m. The corresponding pump absorption following eq.3.2 is $\simeq$ 70%. A 100 nm thick layer of the same phosphor would achieve just 0.07%, three orders of magnitude short of that needed produce white output in a single

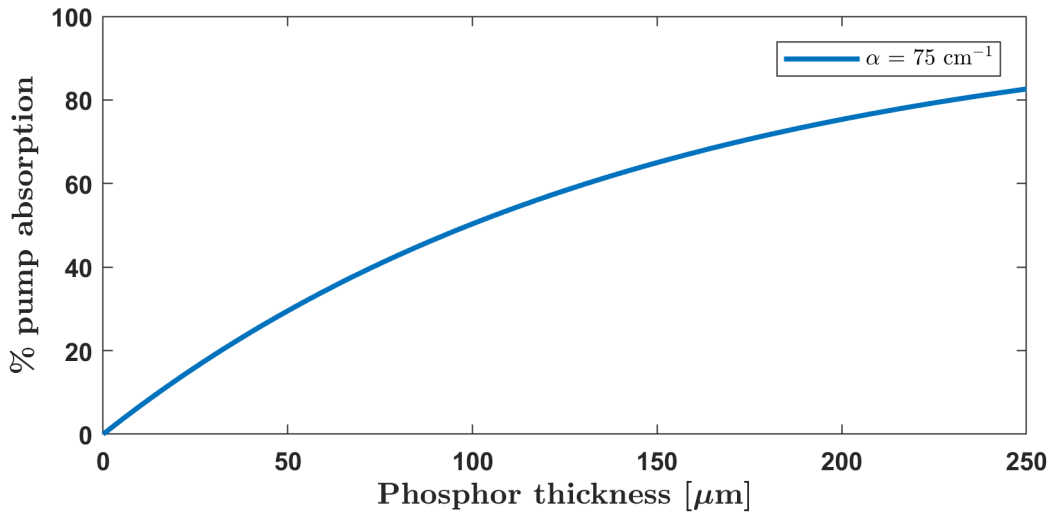


**Figure 3.1:** Percentage pump absorption by a given thickness of phosphor, monocrystalline disc samples from Crytur. The disc that produced the cool white spectrum of Fig.1.4 had a thickness of ~ 180 $\mu$m. An equivalent 100 nm thick disc would achieve pump absorption of just 0.07%. The absorption coefficient $\alpha$ =75 cm$^{-1}$ comes from correspondence with Crytur.

pass.

It is important to note that the reference phosphor we have considered does not exhibit the maximum achievable pump absorption with a YAG:Ce$^{3+}$ phosphor. At just under 0.2 at.% concentration it is around $\frac{1}{3} \rightarrow \frac{1}{2}$ of the maximum concentration imposed by the solubility limit in monocrystalline samples [18]. In polycrystalline samples the cerium incorporation can be an order of magnitude higher [18], but this does not alter the conclusion that just obtaining the necessary pump absorption in a nanoscale phosphor layer would present a significant challenge.

## 3.4 Conclusions

We have seen in Section 3.2 that the ability of a nanoscale phosphor layer to meet the luminous output and size targets has not been ruled out based upon a much simplified treatment of the phosphor excitation dynamics. When factoring in the probability of emitter excitation, it quickly becomes apparent that the pump absorption in a nanoscale layer is orders of magnitude below that needed to meet the output colour functional requirement. This shortfall is covered in the previous Section 3.3. It is important to note that we do not conclude achieving sufficient absorption to be impossible, after all concentrated electro-

magnetic fields can lead to substantial increases in absorption through changes in the local density of states [81]. In addition, one could consider designing the component such that the pump light makes multiple passes through the thin phosphor layer, perhaps in some sort of optical cavity. Pump absorption is however a sufficiently large obstacle to be managed alongside the other requirements that it sends nanoscale solutions to the bottom of the list of preferred strategies.

Of the remaining strategies, we deemed macroscale fluorescent optical fibres to be the most promising avenue for further investigation. This decision is aided by the fact that the next step in assessing this approach, calculation of the trapping fraction of fluorescence within the fibres, is a fairly straightforward exercise. Indeed we investigated the suitability of fibres further and were not sufficiently discouraged by what we found. The assessment of fluorescent optical fibres as a platform for meeting the project objectives is the topic of the next chapter, in which we present them as our proposed solution.

# 4 Fluorescent Fibres

## 4.1 Fundamentals and Motivation

Optical fibres are traditionally long cylinders of optically dense (high refractive index) material, capable of efficiently transporting light along their length via total internal reflection (TIR). If the diameter of the fibre is of the order of microns, as is usually the case, it will also be flexible. For a fibre of diameter much greater than the relevant wavelength (Geometrical Optics regime), the propagation of light is captured by the well-known Snell's Law of refraction at an interface

$$n_0 \sin\beta_0 = n_1 \sin\beta_1, \tag{4.1}$$

where $n_0$,$n_1$ and $\beta_0$,$\beta_1$ are the refractive indices and incident / exit angles (from normal) in media 0 and 1. The condition for the onset of TIR is obtained from 4.1 when the exit angle $\beta_1 = \frac{\pi}{2}$ and hence this critical angle

$$\beta_c = \arcsin\left(\frac{n_1}{n_0}\right), \tag{4.2}$$

which occurs only when light is attempting to exit a higher index medium, such as in an optical fibre. This basic principle allows the high speed optical transfer of vast amounts of information over great distances, with minimal losses.

We propose the use of fluorescent optical fibres made of scatter-free phosphor such as monocrystalline YAG:$Ce^{3+}$ as the colour converting component in a directional solid-state white light. The high energy blue pump light is coupled into the fibre at angles satisfying the conditions for TIR, after which a fraction of it is absorbed and excites fluorescence within the phosphor fibre. As the generated fluorescence is isotropic, only some of the emission will be guided by the fibre. This is illustrated in Fig.4.1. The guided fluorescence and pump light are then emitted from the small end facet(s), enabling easy collimation with external or directly coupled optical components.

Before going into the details of the calculations of the trapping efficiency of such fibres, we briefly discuss how the efficiency of the whole directional light source might be optimized. Assuming the fibre is clad in air to maximize the light trapping efficiency, one has three basic variables one can adjust. We consider these to be the fibre itself, modifications to the input facet and likewise to the output facet of the fibre. Fig.4.2 shows schematically our earliest proposal for a full directional light source, highlighting these three components and suggesting implementations.

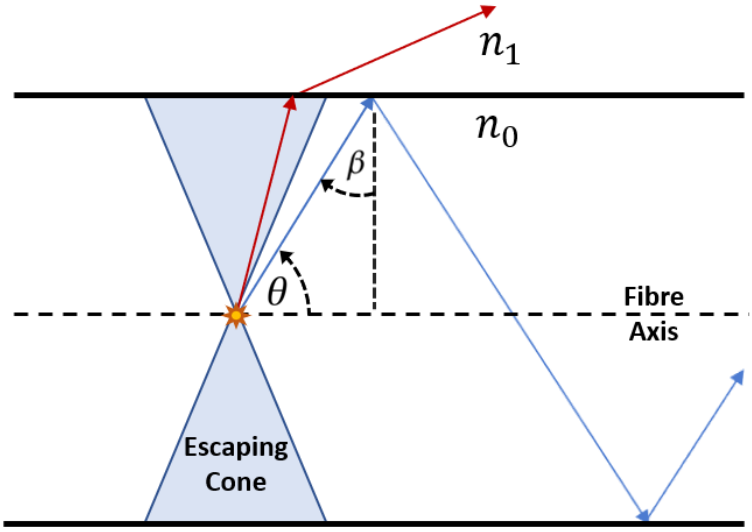


**Figure 4.1:** Illustration of light trapping in a fluorescent fibre. Emission from the central axis may either satisfy the conditions for total internal reflection and be guided (blue arrow), or be emitted into angles which cause the light to escape (red arrow). $n_0$ and $n_1$ are the refractive indices of the fibre and cladding respectively. The illustrated geometry has rotational symmetry about the fibre axis, forming a cylindrical fibre.

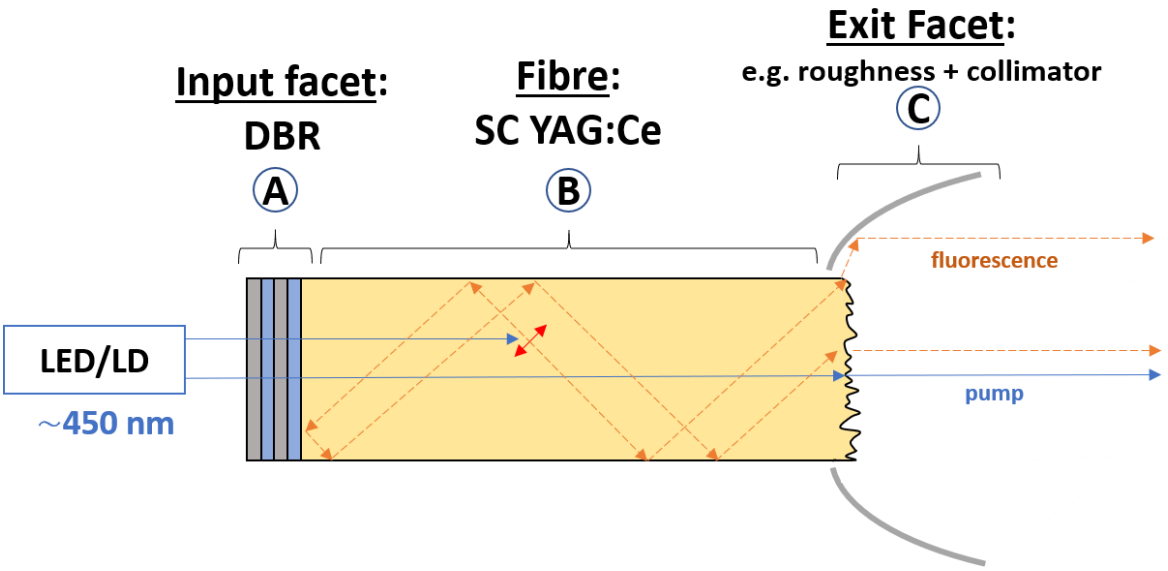


**Figure 4.2:** An early proposal for a directional fluorescent fibre based solid-state white light. The function of each of the components A-C are discussed in the text. Blue and orange lines illustrate ray trajectories (artistic) for the pump light and generated fluorescence. DBR = Distributed Bragg Reflector.

The functions of each of the components / modifications labelled A-C in Fig.4.2 are discussed below along with suggestions for achieving them with conventional optical components where possible.

- **A - Input Face: Dichroic element** 50% of the fluorescence trapped within the fibre will propagate towards the input facet. This component needs to distinguish the incident pump light from the fluorescence, to allow the pump light into the fibre whilst reflecting fluorescence back towards the exit facet. The obvious solution would be a dichroic element, which has differing reflectively at the pump and fluorescence wavelengths. For this a distributed Bragg reflector seems a natural choice.
- **B - The Fluorescent Fibre.** Cerium doped YAG in its monocrystalline form provides a scatter-free, high index (1.83) medium with the familiar and favourable luminescence properties that have seen it widely adopted for lighting applications. Alternative phosphors such as transparent ceramics ([30, 82]), glasses ([23, 83]) and glass-ceramics ([20, 24, 25, 84]) could be considered if they can provide a high refractive index, excellent surface and bulk properties and have suitable absorption and emission spectra.
- **C - Exit Face: Emission shaping / outcoupling component(s).** The two main functions that might need to be performed by modifications to the end facet of the fibre are redirecting and outcoupling emission. These could be performed by two separate components, as is depicted in Fig.4.2 with outcoupling provided by surface roughness and directionality by the external parabolic collimator, or with a multi-function modification such as shaping the fibre end into a lens. One example of an existing device that could perform both functions would be the compound parabolic reflector explored in [85].

### 4.1.1 Fluorescent Fibres vs the Functional Requirements

We now look briefly at the perceived advantages of a fluorescent fibre as a directional phosphor component, with respect to the functional requirements outlined in Section 1.4.1. These were the appealing characteristics predicted prior to embarking on a full investigation and so contain some important unknowns. At the end of this chapter (Section 4.9.1), we will look again at these requirements for a more complete characterisation.

1. **Emission Cone**

   Emission into a narrow solid angle is the highest priority requirement of the sought phosphor component. The angular distribution from an unmodified YAG:$Ce^{3+}$ fibre was at this stage yet to be determined. However, due to the ability to distribute the necessary volume of phosphor along the length of the fibre whilst maintaining a small cross-sectional area (CSA), fibres can provide low etendue sources of light that lend themselves to efficient collimation with external reflective or refractive components.

2. **Imaging Capability**

   The phosphor component needs to be able to retain the incident spatial intensity pattern of the pump beam in its output. Fluorescent fibres were chosen with imaging capability at the heart of the decision. A bundle of optically uncoupled fibres illuminated with a patterned beam will relay the spatial patterning to the fibre output with a resolution determined by the number and distribution of the individual fibres. This is the principle behind the fibrescope. Achieving a resolution that exceeds the lower limit "of the order tens by tens of pixels" within a shoebox ($\sim 20 \times 30$ cm) area is expected to present little challenge due to the small CSA of optical fibres.

3. **Optical Efficiency**

   This was at this stage undetermined and will depend heavily on the trapping efficiency of the fluorescence generated within the fibre, which is considered in the following section.

4. **Brightness / Power Handling**

   The phosphor component needs to be capable of producing in the region of 10,000 to 20,000 lumens of directional white output within an area $\sim 20 \times 30$ cm. Similar to the rationale presented in the Emission Cone section, the ability to add additional luminescent material to the length of the fibre without needing to increase its CSA makes fibres well suited to meeting this requirement. Luminescence saturation as the number of activator ions in the ground state decreases under high illuminance can in principle be compensated for simply by extending the length of the fibre, with no need to increase its CSA. In addition, YAG:$Ce^{3+}$ is noted to have a "remarkably

high saturation threshold" [86].

5. **Colour**

   The lack of absorption which presents a significant challenge for a nanoscale phosphor layer is addressed with ease in this geometry, where optical path length is available in abundance. The correlated colour temperature of the emission of a fluorescent fibre is tuneable with the fibre length or alternatively the pump angular distribution.

6. **Thermal Performance**

   Monocrystalline YAG:$Ce^{3+}$ offers vastly improved thermal properties over the traditional resin dispersed phosphor implementations, with a thermal conductivity of $13 Wm^{-1}K^{-1}$ [87] and stable performance up to operating temperatures of 300°C [18]. Again the flexibility (by which we don't mean bending) of the fibre geometry allows one to extend the overall surface area of the phosphor, aiding in heat dissipation, whilst maintain a small output CSA.

## 4.2 The Trapping Fraction

One key question facing the proposed fibre structure is; with what efficiency will fluorescence generated within the fibre couple to its guided modes? For a large number of randomly orientated dipole emitters, the generated light will be radiated isotropically into $4\pi$ steradians. Only a subset of these emission angles will satisfy the condition for total internal reflection at the fibre surface and thus be guided.

From the critical condition eq.4.2 one can integrate emission angles which are guided and determine the trapping fraction $F_T$ by comparison of guided emission $\Omega$ with the total emission of $4\pi$

$$F_T = \frac{\Omega}{4\pi} = 2 \times \int_0^{2\pi} \mathrm{d}\phi \int_0^{\frac{\pi}{2}-\beta_c} \sin(\theta)\mathrm{d}\theta = 1 - \frac{n_1}{n_0} \tag{4.3}$$

$\theta$ here is the emission angle starting from the fibre axis, as shown in Fig.4.1. This expression for the trapping fraction in a fluorescent fibre is commonplace in the assessment of fluorescent scintillating fibres for radiation detection [88, 89]. For plastic scintillating fibres, where the core is clad in a material of similar refractive index, trapping fractions around 10% are common [88, 90]. Here, due to the high refractive index of YAG:$Ce^{3+}$ and by maximising the index contrast through free space cladding, $F_T = 0.45$ or 45%.

### 4.2.1 3D Fibre Trapping Fraction

The simple analytical expression for the trapping fraction (eq.4.3) derived previously was validated in a COMSOL multi-physics (Geometrical Optics module), three dimensional simulation. Details of the simulation can be found in Section 5.1. When evaluating the trapping fraction in an air clad cylindrical YAG fibre it quickly became clear that the trapping fraction had a significant dependence on the radial position of the emitter. Fig.4.3 shows how the trapping fraction varies with radial displacement from the central fibre axis.

The results of Fig.4.3 confirm the previously calculated trapping fraction of 45% (within simulation precision), but only for sources positioned on the fibre axis, as was implicitly assumed previously. Examination of Fig.4.1 offers no insight as to why such a radial dependence should exist. However, when one considers a cross-section of the fibre, perpendicular to its axis, this dependence becomes clear. Fig.4.4 demonstrates qualitatively, via

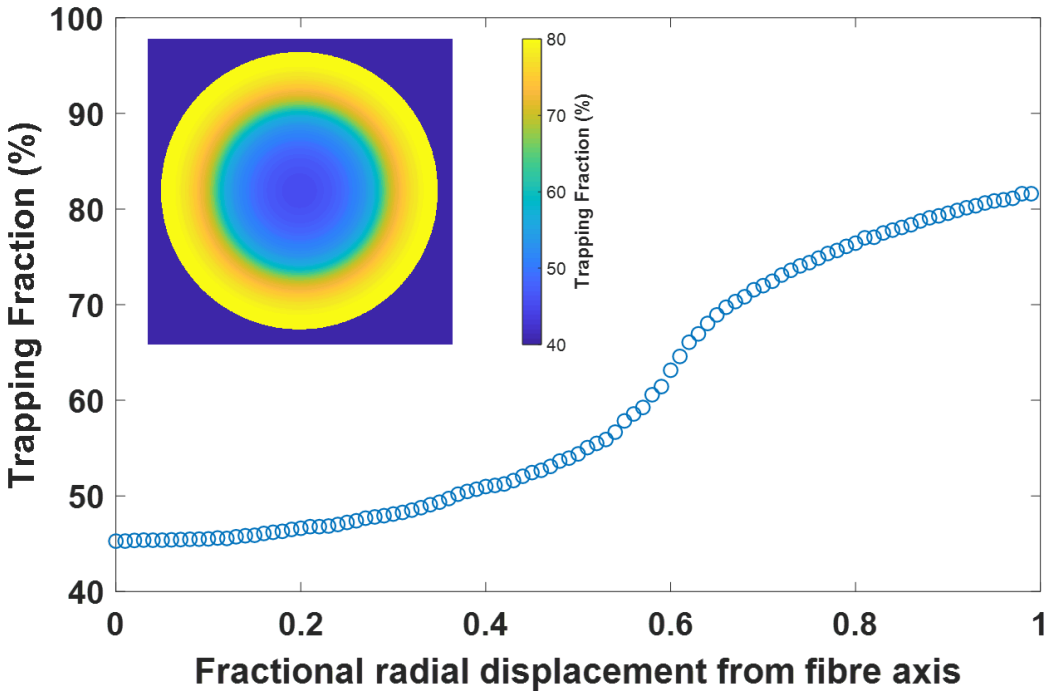


**Figure 4.3:** Variation in fraction of isotropically emitted light which is guided with radial distance of soruce from the center of a YAG fibre. Results of 3D COMSOL simulation, full details in Section 5.1. (Inset) the same radially symmetric data shown as a colour scale on a cross-sectional view of the fibre. The edge of the fibre is the transition from yellow (~ 80%) to dark blue (0 / no data).

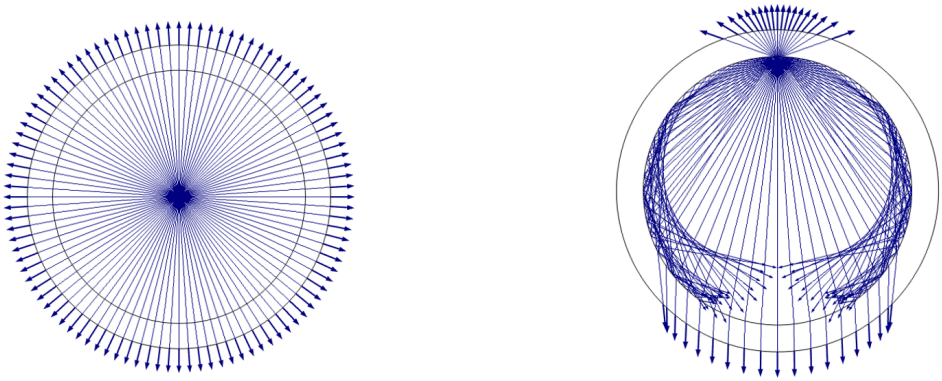

**(a)** Emission from fibre center

**(b)** Radially displaced emission

**Figure 4.4:** Ray trajectories for emission from within a YAG fibre (inner circle), from a 2D COMSOL simulation. The outer circle is the simulation boundary, upon contact with which rays (blue arrows) freeze. Improved trapping ability is observed in (b).

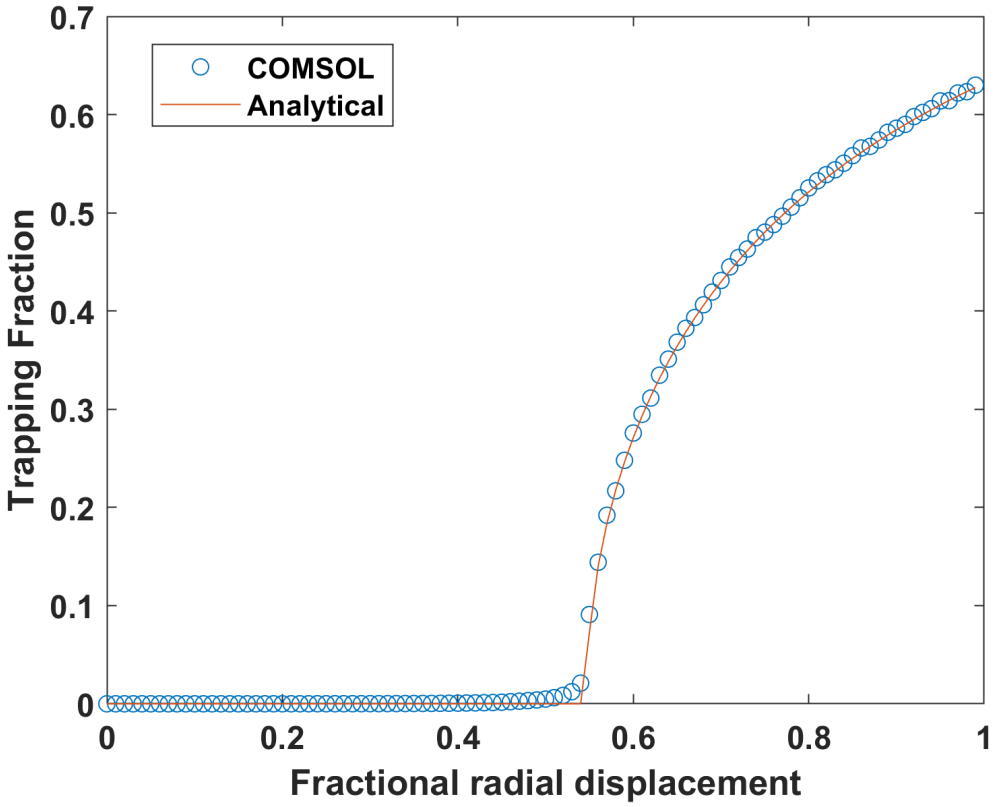


**Figure 4.5:** Trapping fraction variation with radial position for a two-dimensional YAG fibre cross-section, COMSOL simulation vs analytically derived expression (see text, eq.4.5). $r$ is the fractional radial displacement from the circle center. Simulation details in Section 5.5

plots of ray trajectories that displacing the source from the fibre axis should result in an increased trapping fraction.

Fig.4.5 quantifies this behaviour, comparing the trapping fraction calculated from the 2D COMSOL simulation with the following analytical expression for $F_T$. If one considers a source radially displaced from the center of a circle in the $\hat{y}$ direction on 2D Cartesian axes, with emission angle $\theta$ beginning from $\hat{x}$, the incident angle on the circle surface is then

$$\beta = \arcsin\left(r\cos\theta\right) \tag{4.4}$$

where $r$ is the fractional radial displacement. From which $F_T$ is obtained by normalised integration of the range of $\theta$ values for which $|\beta| < \beta_c$

$$F_T = \frac{2}{\pi}\arccos\left(\frac{\sin\beta_c}{r}\right). \tag{4.5}$$

This is the expression plotted in Fig.4.5, where for YAG surrounded by air $\beta_c = 33.1°$.

The reason for the delayed onset of increased trapping fraction in Fig.4.5 is that the increase in Fresnel reflection one would observe for a source just off the fibre center does not constitute an increased trapping fraction. Only once the source is displaced sufficiently for total internal reflection to occur is the light considered trapped.

The first analytical treatment of the trapping fraction in a fluorescent fibre which includes the contribution of rays positioned off the fibre axis was published as recently as 2015 [90]. Previous analyses have restricted their consideration to rays originating on the fibre axis, known as meridional rays [88, 89]. The justification here is that those rays originating off-axis, so called skew rays, must travel further and interact with the fibre surface more frequently than meridional rays, so are attenuated faster. Indeed for scintillating fibres where the length and concentration of the fibre are typically much greater than would be required in lighting it is remarked that experimentally one might observe one or two % more light due to the contribution from skew rays [88, 89]. We now turn our attention to the results of the full analytical treatment in [90].

The total trapping fraction ($F_{\text{total}}$) can be separated into the contributions arising from rays which are emitted at angles that are guided for all values of radial displacement, so called 'narrow-angle' rays ($F_{\text{narrow}}$), and those additional 'wide-angle' rays ($F_{\text{wide}}$) which are trapped for emission off axis [90].

$$F_{\text{narrow}} = 1 - \frac{n_1}{n_0}, \tag{4.6}$$

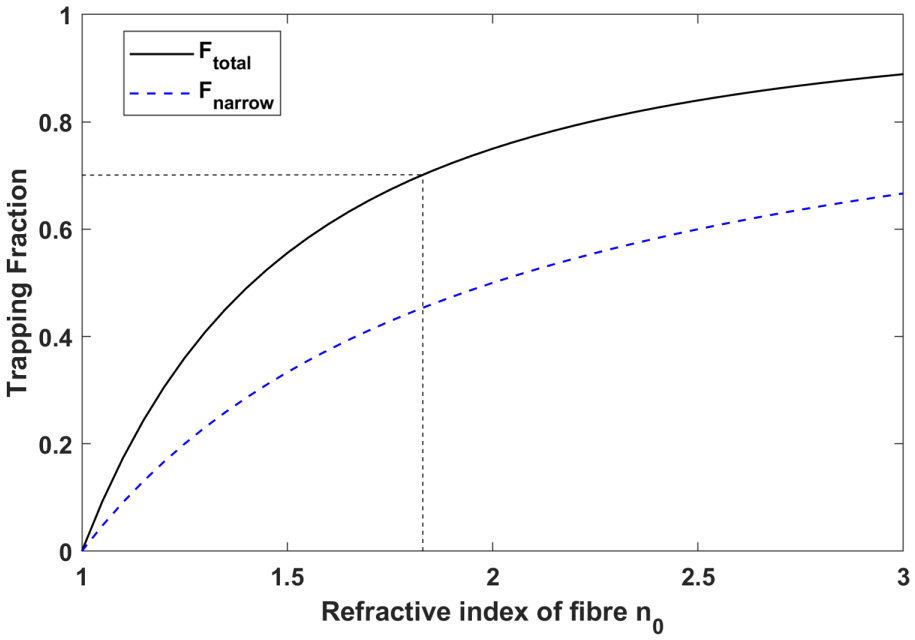


**Figure 4.6:** Variation in trapping fraction for isotropic emission within an air clad ($n_1 = 1$) fibre with fibre refractive index $n_0$. The contribution arising from narrow-angle rays (eq.4.6) shown here highlights the cost of neglecting wide-angle (eq.4.7) rays. The total trapping fraction for the proposed fibres, of index 1.83 is 70% as marked by the dashed black lines.

which is the result derived in eq.4.3.

$$F_{\text{wide}} = \frac{n_1}{n_0}\left(1 - \frac{n_1}{n_0}\right), \tag{4.7}$$

and hence

$$F_{\text{total}} = 1 - \left(\frac{n_1}{n_0}\right)^2. \tag{4.8}$$

Figure 4.6 compares the total trapping fraction with the contribution from narrow-angle rays, for an air-clad fibre of varying refractive index. The impact of including only meridional rays, given by $F_{\text{narrow}}$ is shown to be substantial and of decreasing consequence with increasing fibre index. For an index of 1.83 as in the case of YAG:$Ce^{3+}$ we see a significant increase in trapping efficiency from 45% to 70% by including the contribution of the skew rays.

For comparison we calculate the mean trapping fraction from the radially dependent simulation data of Fig.4.3, we numerically integrate the trapping fraction $F(r)$ over the circular cross-section we have sampled and normalise by that area. In the continuum this

is

$$\bar{F} = \frac{\int_0^{2\pi}\int_0^R F(r) r \mathrm{d}r\mathrm{d}\theta}{\int_0^{2\pi}\int_0^R r \mathrm{d}r\mathrm{d}\theta} \tag{4.9}$$

where $R = r_{\max}$, the radius of the circular cross-section of the fibre sampled. For our discrete data this becomes

$$\bar{F} = \frac{2\pi\Delta r}{\pi R^2}\sum_i F(r_i) r_i \tag{4.10}$$

where $\Delta r$ is the spacing between samples of the function $F(r_i)$. Noting that the data from Fig.4.3 may have some small associated uncertainty due to limitations of the simulation, details Section 5.1, we calculate the mean trapping fraction for uniformly distributed emitters to be 70.3%. This value is in excellent agreement with that obtained using eq.4.8, of 70.1% as shown in Fig.4.3.

### 4.2.2 Square Cross-section Trapping Fraction

A circular cross-section is the expected geometry for fluorescent guides manufactured using conventional fibre drawing techniques. However, where phosphor light guides are machined from a crystal boule alternative cross-sections become readily available and can even be preferable from a manufacturing point of view. With this in mind it is beneficial to understand the trapping properties of the natural cross-section when manufacturing

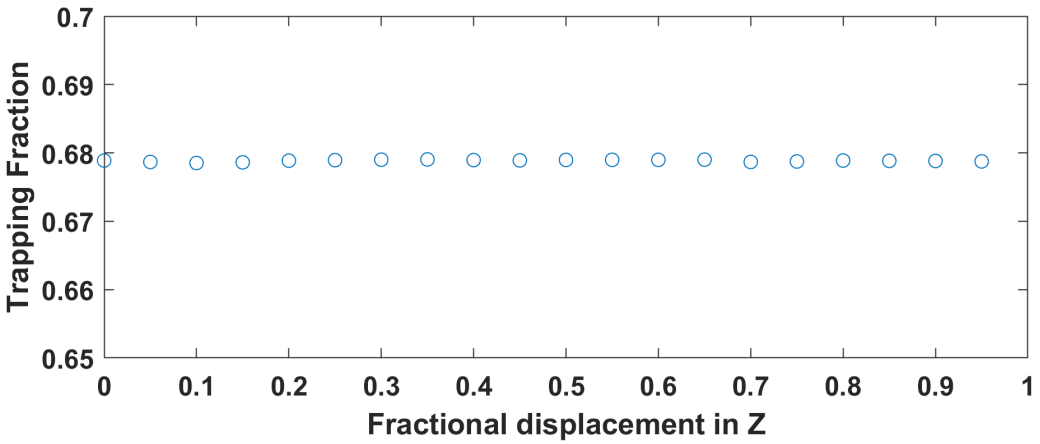


**Figure 4.7:** Confirmation of the emission position independent trapping fraction within a fluorescent cuboidal YAG guide. Results of a 3D COMSOL simulation. Where 0 is the center of the cuboid cross-section, a fractional Z displacement of 1 represents a source displaced in the direction of an edge surface-normal all the way to that edge. Simulation details in Section 5.2.

guides this way, the square cross-section.

For a square cross-section guide the trapping fraction does not vary with pump position. This is quantified by a full 3D simulation in Fig.4.7. The amount of light trapped within the fibre remains constant at 68% as expected for a loss of 8% per side, given by $\frac{1}{2}(1-\cos\theta_c)$ [91], where $\theta_c = 33.1°$. The simulation (details in Section 5.2) also confirms the expected 8% escaping per fibre end (not presented here).

Neglecting the contribution of skew rays in the cylindrical geometry could lead one to wrongly conclude that only a square cross-section can provide an acceptable trapping efficiency (68% vs 45%). As the mean trapping fraction when illuminated across the entire guide cross-section is comparable for both geometries, the preferable geometry is likely to be dictated by manufacturing feasibility.

### 4.2.3 Discussion

The results obtained in Section 4.2.1 through simulation and analytical treatment tell us that for a fully illuminated (across its cross-section), perfect, air-clad cylindrical YAG fibre the trapping fraction is impressively high at 70%. If a value close to this could be achieved experimentally then it would demonstrate that the trapping fraction is not nearly as limiting for device efficiency as one may have concluded neglecting the contribution of skew rays (eq.4.3). In a real sample the presence of surface roughness, bulk scattering / absorption and re-absorption of fluorescence by the phosphor will reduce device efficiency and may mask the predicted emission positional dependence. The dependence of the trapping fraction on the dominant loss mechanisms is explored through simulation in Section 5.4. Our experimental objective was to determine if the predicted radial dependence was still measurable in the presence of these loss mechanisms and if that led to acceptable device efficiency.

The radially dependent trapping fraction could also be actively exploited. If the surface scattering of a fibre was sufficiently low, a higher concentration of dopants towards the edges of the fibre would yield a mean trapping fraction for a uniformly illuminated fibre cross-section that exceeds 70%. The maximum attainable trapping fraction would then be $\simeq 80\%$ in the limit that all of the dopants are at the fibre edge. It is noteworthy therefore that a strong radial dispersion of dopant concentration occurs naturally in fibres manufactured by the micro-pulling-down technique [88], this can be seen in Fig.4.8. An alternative approach to achieving the same goal would be to consider non-uniform distribution of the pump light over the the fibre cross-section. A pump beam that propagates as a ring near

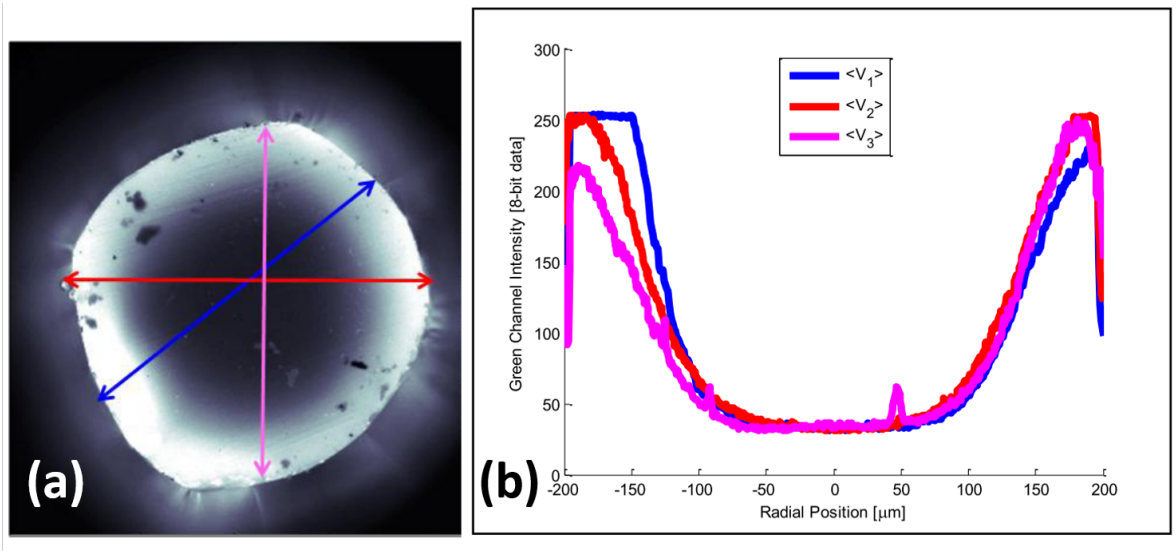


**Figure 4.8:** Figures from [88], showing a potentially exploitable fibre radial dopant dispersion. Measurement of the relative dopant concentration of a 400 $\mu$m diameter YAG:$Ce^{3+}$ fibre, exhibiting the strong radial dispersion characteristic of the micro-pulling-down fabrication method. (a) Green channel of an image of a sawn 100 $\mu$m thick slice of fibre under UV illumination. The dark spots are attributed to surface debris from the sawing process. (b) Line plots of the green channel intensity, as indicated by the coloured lines in (a). Some detector saturation is present.

the fibre edge would have a similar effect on the trapping fraction.

## 4.3 Outcoupling and Directing Fibre Emission

The high trapping fraction achievable in YAG fibres due to the high index contrast between the fibre and surrounding has the consequence by definition that many of the rays propagating down the fibre do so with large transverse wavevectors. That is, they are travelling in paths at large angles from the fibre axis. A significant consequence of this is that the emission from the fibre ends is not directional as one would observe in a fibre of low refractive index contrast. The other result is that some of the rays propagate at such high angles that they will be reflected from the end facet of the fibre (incident angle on end facet is $\pi/2$ - the angle from the fibre surface). Both of these issues may be addressed by modifying the geometry of the ends of the fibre.

To allow high-angle rays to escape the end facet of the fibre, the condition for total internal reflection eq.4.2 must be broken. For scintillating fibres, where one needs only transport the generated fluorescence to a photodetector, this can be achieved by refractive index

matching [88]. Placing a material of refractive index closely matching that of the YAG fibre between it and the detector is a simple but effective method of allowing up to 100% of the initially trapped rays to reach the detector. Whilst this is an excellent solution for scintillating fibres, it is a requirement for general lighting applications that the emission overcomes the index contrast between the fibre and free space. As such we turn our attention to the geometry of the facet, to change the incident angle $\beta$ upon it and allow these rays to outcouple.

### 4.3.1 Surface roughness

Indefinite confinement of the fluorescence generated within the fibre is a scenario that can exist only within simulation or simplified analytical treatments. For any real fibre the light will eventually be able to escape via some surface or bulk defect, or else be lost to absorption. With this in mind, one could accelerate and localise the outcoupling by introducing intentional surface roughness to one end of the fibre. Perturbations to the shape of the fibre facet change the incident angle of light on the surface locally, provided the roughness is not sub-wavelength. If this does not cause a given ray to out-couple upon first contact with the roughened end, its perturbed angle of reflection will see to it that the location of

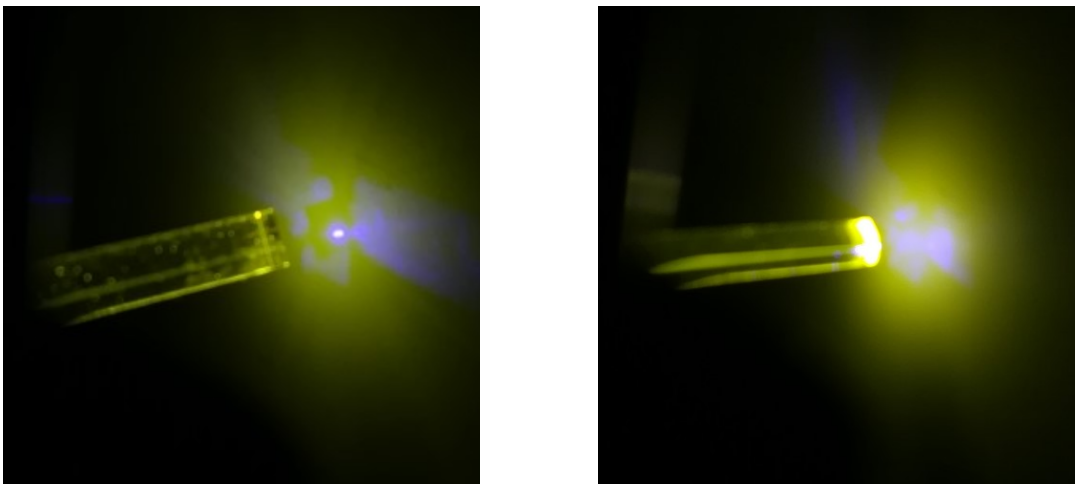

**(a)** Unmodified, flat guide end **(b)** Roughened guide end

**Figure 4.9:** Images of two cuboidal YAG:$Ce^{3+}$ guides illuminated by a thin 'pencil' of 450 nm laser light, taken with a camera phone. The images are taken from behind the end facet, such that only light which is reflected from the end facet or strikes the white screen may be seen. For the unmodified flat end (a), no observable reflected light out-couples from the guide and reaches the camera. For the roughened guide (b) [ground with a diamond disc], a significant amount of light is outcoupled in the direction of the camera.

contact for subsequent passes is essentially randomised. The appeal of this approach lies in the ease of fabrication (e.g. sandpaper of suitable hardness), but due to the perturbation of reflected ray trajectories, outcoupling may occur in directions other than out of the roughened end facet. This is detrimental for device efficiency and can be seen in Fig.4.9

### 4.3.2 Hemispherical Lenses

Maximal transmission through the end facet of the fibre is achieved when all rays are incident normal to the surface. This is the scenario depicted in the fibre cross-section simulation of Fig.4.4a. This may be achieved with relative ease by attaching a large index-matched hemispherical lens to the fibre, as in Fig.4.10. If the lens diameter is significantly larger than the fibre diameter then all rays will strike the surface at normal incidence (Fig.4.10a). However, for smaller lenses where one is not in the regime where the fibre can be considered a point source, this is no longer the case. (Fig.4.10b) illustrates the limiting case where the lens diameter equals that of the fibre and outcoupling efficiency is severely reduced. Fig.4.11 shows how the out-coupled fraction increases with the diameter of the hemispherical lens on the end of the fibre, based on a 2D simulation, for both a bare index matched

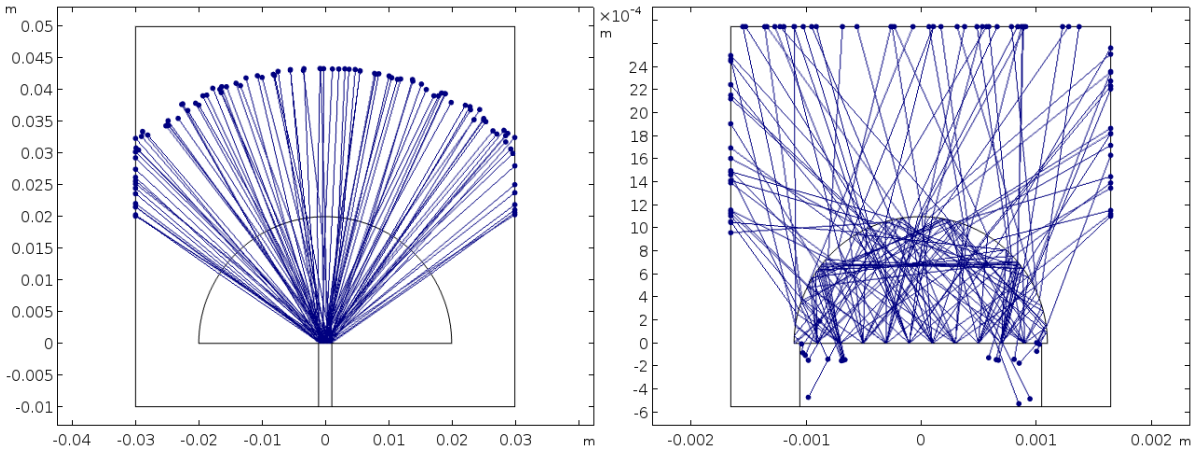


**(a)** Lens diameter >> fibre diameter

**(b)** Lens diameter ~ fibre diameter

**Figure 4.10:** Ray trajectories for emission escaping from the end of a YAG fibre by means of a hemispherical lens. (a) large lens for which the fibre is approximately a point source of emission. (b) small lens, where the spatial extent of the source limits outcoupling. Results of 2D COMSOL simulation, see Section 5.6. The ray sources are 10 uniformly spaced emission cones along a 2 mm line where the fibre meets the lens. The cone emission angles are constant and determined by the angles that are guided by the YAG fibre ($\pm$ 57°). The lens diameters and trajectory snapshot times of (a) and (b) are 20 mm and 1.1 mm and 0.2 ns and 0.02 ns respectively.

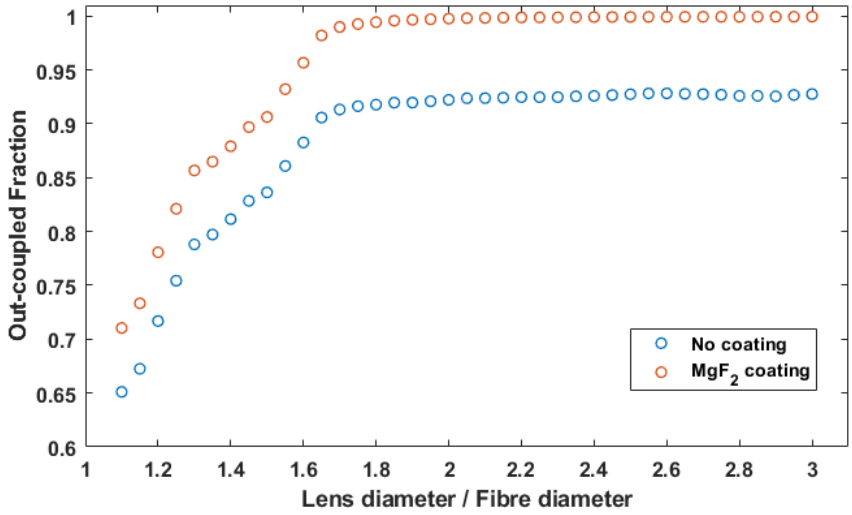


**Figure 4.11:** Fraction of rays out-coupled from the end of a YAG fibre by a hemispherical index matched lens, as a function of its diameter, as depicted in Fig.4.10. Results of 2D COMSOL simulation, see Section 5.6 for full details. The angular spread of the rays entering the lens was restricted to those that would be guided in the YAG fibre. Results are shown both with and without the presence of a $MgF_2$ anti-reflective coating on the lens surface.

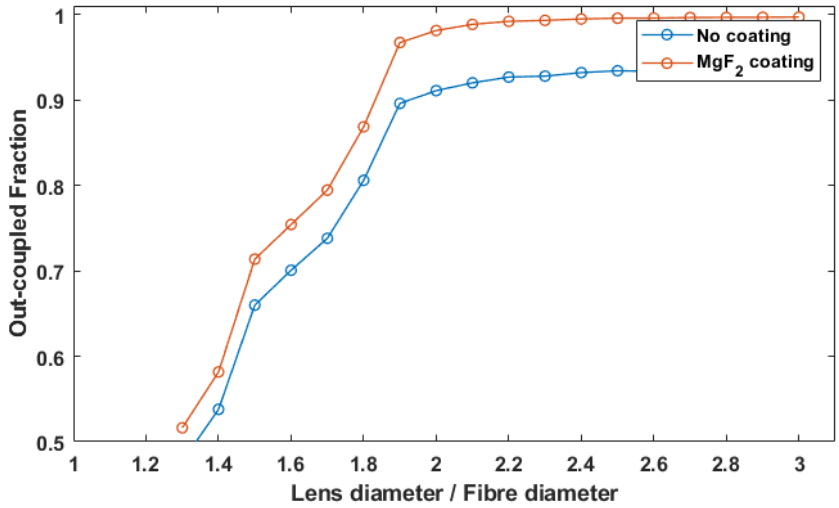


**Figure 4.12:** As in Fig.4.11, but from a 3D simulation. The effect of including skew rays is to delay the onset of efficient outcoupling with increasing lens diameter. Fraction of rays out-coupled from the end of a YAG fibre by a hemispherical index matched lens, as a function of its diameter, as depicted in Fig.4.10. Results of 3D COMSOL simulation, see Section 5.7 for full details. The lines connecting data points are visual aids only.

lens and one with an $MgF_2$ anti-reflective coating.

Whilst the 2D simulation based results of Fig.4.11 seem satisfactory given the apparent rotational symmetry of the problem, we have already seen the need for a 3D treatment in Section 4.2.1 and how the contributions from skew rays can be substantial. The result of a full 3D simulation are presented in Fig.4.12. It can be seen that the outcoupling efficiency saturates at higher values of lens diameter (around 2.2 vs 1.7 × the fibre diameter), causing the outcoupling efficiency for lenses not much larger than the fibre to be reduced considerably from that expected from the 2D simulation.

A hemispherical lens is a simple but effective means of outcoupling the fluorescence from a YAG fibre, but it offers only a modest reduction in the emission angular spread. We now turn our attention to addressing the emission angular distribution.

### 4.3.3 Ball Lenses

When considering the problem of outcoupling and directing emission from high index fluorescent fibres, the natural association of ball lenses with fibre collimation caused some interest in their suitability. The following two sections explore how they might be adapted for fluorescent fibres.

Ball lenses are as the name suggests spherical and are typically produced on the scale

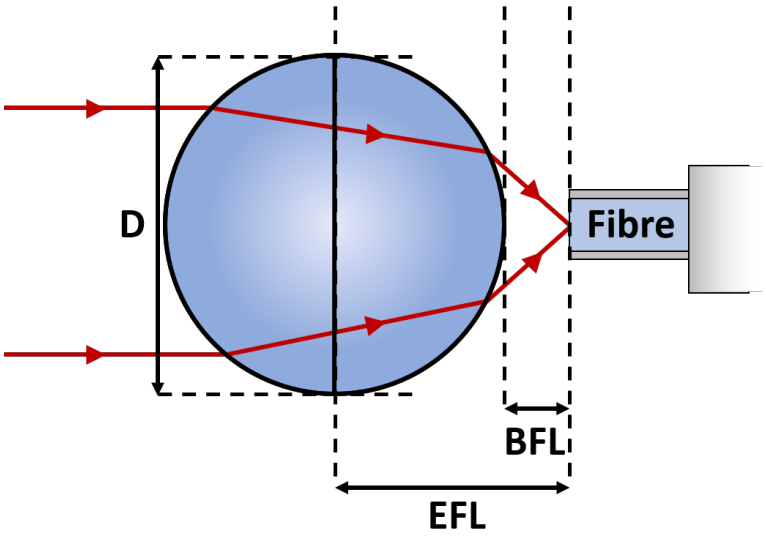


**Figure 4.13:** Schematic of key parameters of a ball lens, here shown coupling a collimated laser beam (left) into an optical fibre (right). $D$ is the diameter of the spherical lens and $EFL$ and $BFL$ are the effective and back focal lengths respectively.

of millimeters or less. Their short focal lengths lend themselves to applications such as source collimation, miniature optics and fibre coupling. Fig.4.13 illustrates a ball lens coupling light into an optical fibre, with the key parameters of a ball lens labelled. The effective focal length ($EFL$), the distance from the centre of the lens to the focal plane, is [92]

$$EFL = \frac{nD}{4(n-1)}, \tag{4.11}$$

where $n$ is the refractive index of the lens and $D$ its diameter. It is assumed the surrounding medium is air. From the $EFL$ one can easily find the back focal length ($BFL$) as follows:

$$BFL = EFL - \frac{D}{2}. \tag{4.12}$$

It is the association with fibre coupling / collimation that makes the ball lens an interesting candidate for controlling the emission from fluorescent fibres. In Fig.4.13 one can see that if the direction of light propagation is reversed, then it is serving to collimate emission from the end of the fibre. In order for a single optical element to achieve the desired emission lensing and outcoupling, it would need to be in contact with the fibre end. One can see from 4.11 that if the index of the lens $n = 2$, that the $EFL$ becomes $D/2$, or equivalently the $BFL$ becomes 0. This means the focal plane now coincides with the point at which the lens would couple to the fibre. For a large ball lens coupled to the end of a fibre, this index condition should allow for the best lensing performance.

Fig.4.14 is a snapshot of ray trajectories for an $n = 2$ ball lens collimating a point conical light source located on its circumference. The emission cone is now restricted by Snell's law 4.1 to 50° half-angle due to the slight index contrast between the fibre ($n = 1.83$) and the lens ($n = 2$). The most striking observation is that a significant number of rays undergo total internal reflection and continue to propagate around the lens never escaping as time progresses.

The second is that emission from the lens is not perfectly collimated. The latter property is of no surprise and is common to any spherical optic, to achieve perfect collimation (in simulations at least, perfection is absent in the real world) the shape needs to be aspheric. The emission cone angle for the proposed YAG fibres is unusually high causing incident angles beyond the critical angle and rays to interact with the lens far from where the paraxial (small-angle) approximation holds, where deviation from the ideal aspheric lens shape is significant. This spherical aberration is therefore particularly pronounced.

Facing the presented challenges, it seemed a natural next step to explore how simple modifications to the spherical shape might offer to mitigate these effects. This is the topic

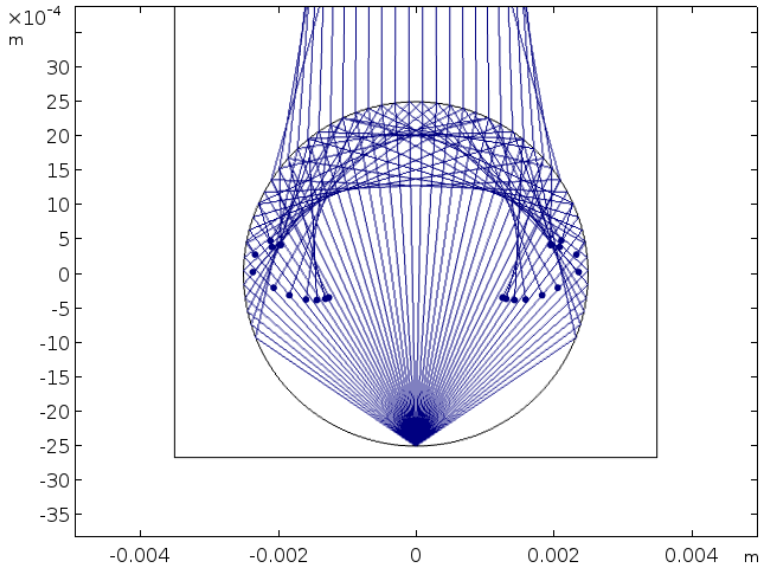


**Figure 4.14:** Snapshot of ray trajectories for a 50° half-angle emission cone into a refractive index of 2 ball lens. Blue circles mark the heads of propagating rays. Due to the high emission angles into the lens a significant portion of the light is trapped within the lens via total internal reflection.

of the following section (Section 4.3.4).

### 4.3.4 Modified Ball Lenses

Rays that propagate down the fibre at large angles from its axis strike the ball lens in such a way that they are totally internally reflected and continue to bounce around the perimeter of the ball at high incident angles from the surface normal, never escaping (as can be seen in Fig.4.14). To disrupt this scenario one could imagine grinding the lens sides flat (or conical in 3D) from the point at which the fibre is coupled to the edge of the lens at some carefully chosen angle such that those large-angle rays first strike this surface, and are then reflected onto the front (top in Fig.4.14) face at an incident angle close to the surface normal. This scenario has been simulated and is shown in Fig.4.15.

It turns out that for a point conical source of light one need not take great care choosing the taper angle and that the outcoupling and lensing performance continues to improve as the taper angle tends to 0°, see Fig.4.15b. Reflection off the tapered edges of the lens effectively folds the large-angle rays back into the lens cone, resulting in excellent collimation for small taper angles.

The remarkably effective performance of the 'modified ball lens' seen in Fig.4.15 is of course limited in the real world by the finite size of the fibre which is coupled to the lens.

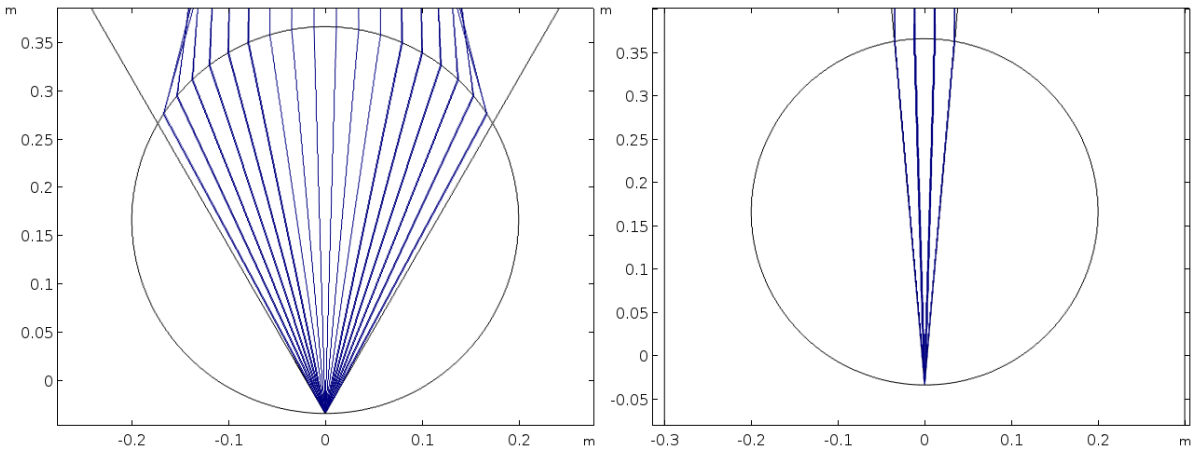


**(a)** Lens taper angle = 30°

**(b)** Lens taper angle = 5°

**Figure 4.15:** Ray trajectories for a 'modified ball lens' in air, with a conical point source of rays just above the lowest part of its surface. The intersection of the cone and the circle defines the region where $n = 2$, the modified ball lens. Results of 2D COMSOL simulation. The tapered edges reflect the large-angle rays back into the lens cone (multiple times for a taper angle of 5°). As the taper tends to just below a point source of light, this causes the reflected rays to be almost superimposed on top of each other. Full simulation details in Section 5.8

For a spatially extended source such as the end of a fluorescent fibre, the lens taper cannot converge to a single point and rays will exit the fibre end laterally displaced from the focal point of the lens. Fig.4.16 shows how the final angular distribution and outcoupling efficiency of a 1 mm line conical light source are effected by the lens taper angle and the diameter of the original ball lens.

The seemingly unusual choice of $d$=15.5 mm in Fig4.16(b) was made to match lateral size constraints for a comparison of the presented lenses with alternative approaches, not included here. One can see from these results that for a spatially extended source the optimum taper angle seems to lie around 15 / 20 °, whereas in the point source case (Fig.4.15) the taper angle could be made arbitrarily small.

If one was to apply such modified ball lenses to a fluorescent fibre based white light source, one would need to balance the requirements for good collimation and outcoupling efficiency with the increased diameter of the optical element. Whilst the presented values of original lens diameter $d$ are large with respect to the fibre diameter, one must remember that due to the tapering the diameter of the modified ball lens is generally much less than its original value. In addition, one of the key benefits offered by a fibre based solution is it should be possible to reduce the fibre diameter by increasing its length and modifying the

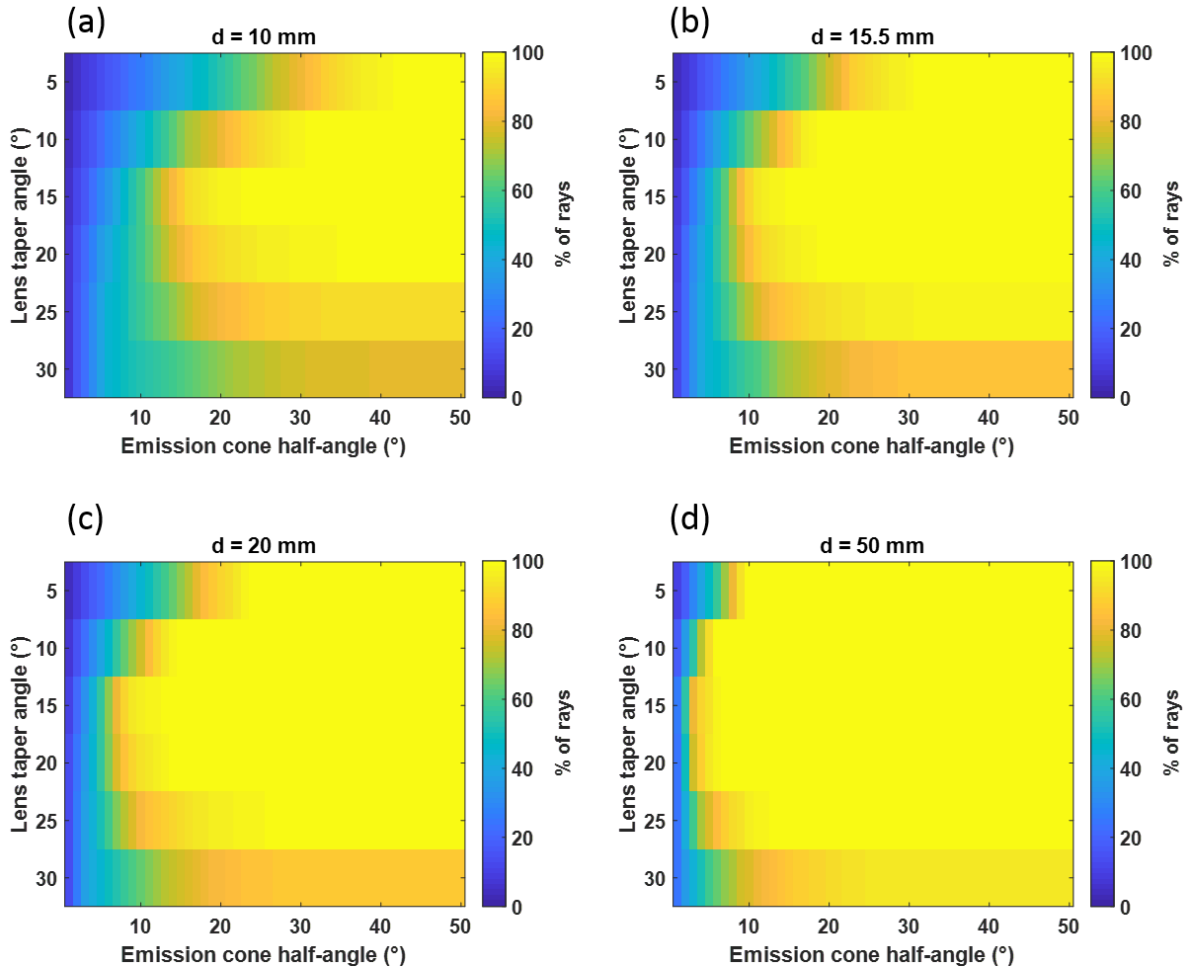


**Figure 4.16:** Effect of lens taper angle and original ball lens diameter on the distribution and amount of light exiting the modified ball lens ($n = 2$) for a spatially extended source. Results of 2D COMSOL simulation, full details in Section 5.8. The light source here is a 1 mm line of conical sources ($\pm$ 50°) on the lens surface. The parameter '% of rays', indicated by colour shows the percentage of total emitted rays which are within the specified emission cone half-angle.

dopant concentration accordingly (subject to manufacturing limitations).

The results of Fig.4.16 are for two-dimensional simulations. As we have seen for the a hemispherical outcoupling optic (Section 4.3.2) a three-dimensional simulation is necessary to make accurate quantitative predictions for device performance. It is predicted based on the changes between the 2D and 3D hemispherical lens simulations that when extending to 3D a modest increase in necessary ball lens diameter would be required to couple an equivalent amount of light into the target as was shown for the 2D simulations. If the modified ball lens were to be of particular interest commercially / academically, we would recommend using a dedicated ray tracing software package to analyse and optimise the structure in 3D.

## 4.4 Simulating Losses in Fluorescent Guides

It is expected that for most 'scatter-free' phosphor fibres the dominant loss mechanism will be scattering or absorption at the fibre surface. Indeed under normal circumstances

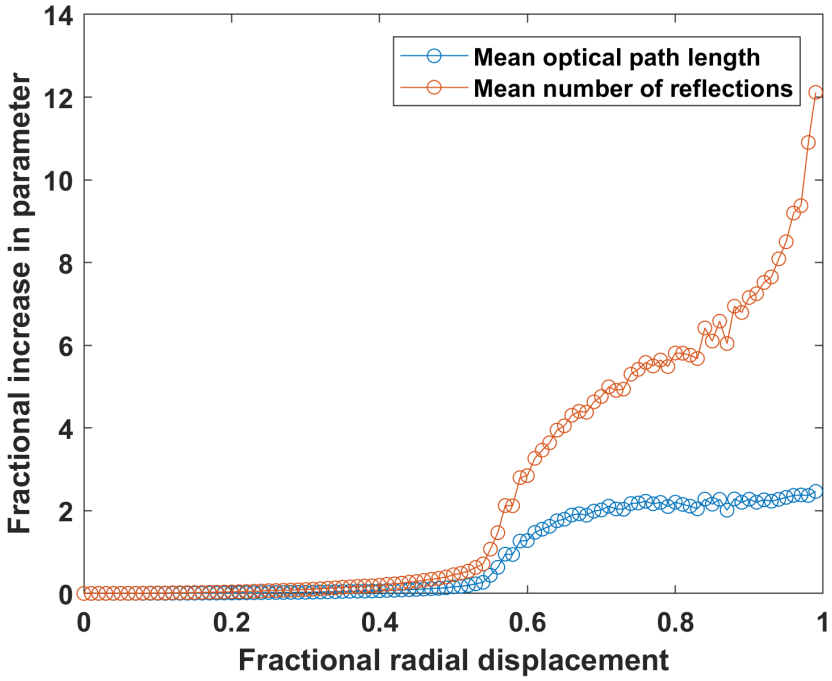


**Figure 4.17:** Fractional increase in the mean optical path length $\bar{L}$ and mean number of reflections $\bar{N}_{\text{refl}}$ of emission radially displaced from a YAG fibre's central axis (e.g. $\left(\frac{\bar{L}(r)}{\bar{L}(r=0)} - 1\right)$, where $r$ is the radial displacement). 3D COMSOL simulation, details in Section 5.3.

a cladding layer would be present to mitigate this and a well made monocrystalline YAG rod is more likely to have imperfections in its surface finish or dust present than scattering defects in the bulk.

We have used numerical simulations to support our expectation that the presence of surface losses, vs those originating in the bulk, are the most likely to mask the trapping radial variation we sought to measure experimentally. Fig.4.17 compares the fractional increase in the mean optical path length $\bar{L}$ and the mean number of reflections $\bar{N}_{\text{refl}}$ for rays

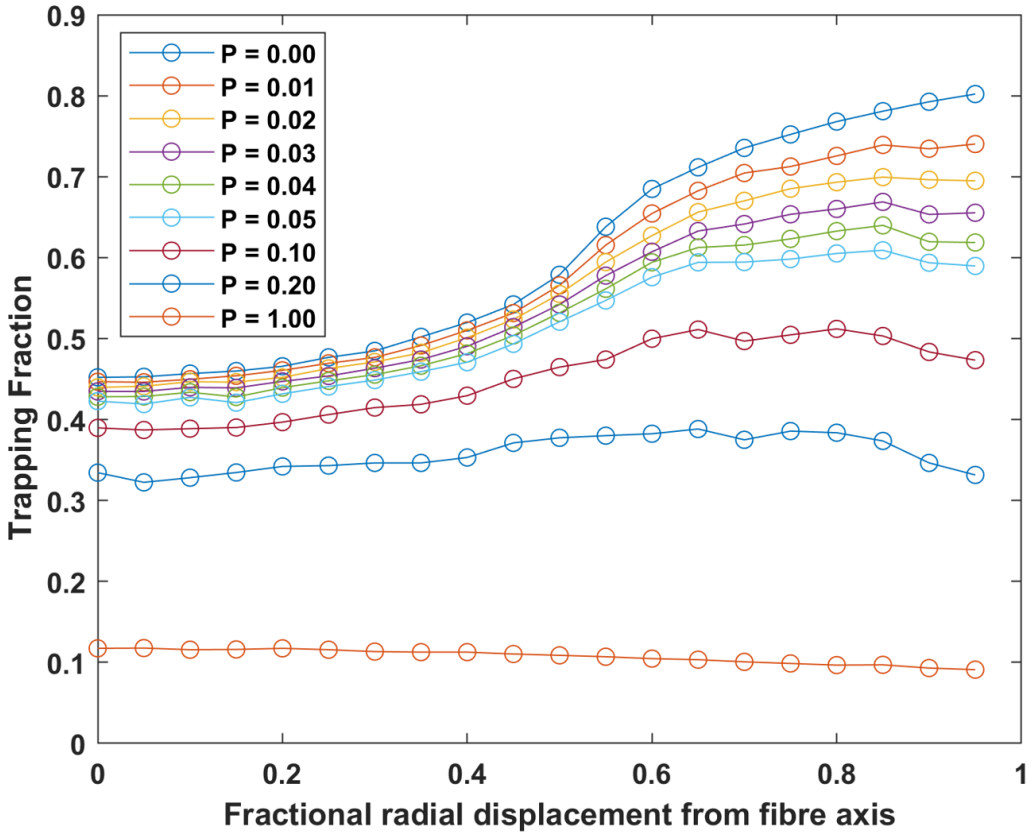


**Figure 4.18:** Variation in the excitation position dependent trapping fraction for a simulated increasingly lossy fibre surface. The legend values correspond to the probability a propagating ray disappears as it passes through a non-physical barrier just below the cylinders curved outer surface, simulating surface losses. Results of a 3D COMSOL simulation, for a 20 mm long, 6 mm radius YAG cylinder - see Section 5.4 for details. The additional rays trapped at large values of radial source displacement are attenuated much faster than those generated on axis. In the limiting case that the surface becomes a perfect absorber, emitters positioned close the the fibre axis yield (marginally) higher trapping fraction values, as one would expect.

emitted at increasing radial displacements from the YAG fibre axis. The amount of scattering / attenuation from the bulk of the fibre increases with the mean optical path length and the surface losses increase with the mean number of times a ray interacts with the surface. Fig.4.17 shows us that $\bar{L}$ and $\bar{N}_{\text{refl}}$ share qualitatively similar dependence on radial displacement, up until fractional displacements beyond ~ 0.8. However, it is clear that the predicted surface losses increase far more rapidly than those originating in the bulk.

For a given realisation of fluorescent fibre the dominant loss mechanism of course still depends on the relative strengths of the bulk and surface scattering coefficients. However, given that we expect the surface quality to be lower than that of the bulk and the radial trapping dependence to be more sensitive to surface losses (Fig.4.17), we have focused our attention on simulating surface scattering.

Fig.4.18 shows how the emission position dependent trapping fraction of Fig.4.3 varies as the probability of rays being lost at the fibre surface increases. As one would expect the effect of surface scattering (modelled as a probability of ray disappearance) is to flatten the radial dependence, with a net reduction in trapping fraction. The probability of ray disappearance at the surface is independent of the ray's incident angle in this simulation, full details in Section 5.4. Rays generated off-axis (skew rays) are attenuated faster than those generated on-axis. This is because skew rays interact more frequently with the fibre surface as they propagate down it (see Fig.4.17), increasing the probability of absorption. Experimental measurement of the trapping fraction variation in the presence of losses is the topic of Section 4.7.

## 4.5 Wavelength-scale Fibres

The work in the following section on wavelength-scale fibres was performed with assistance from Tom Vettenburg, who wrote the underlying Born series based electromagnetic wave solver. His guidance on applying his solver to the scenarios of interest was invaluable.

### 4.5.1 Motivation

With similar motives to our investigation into wavelength-scale parabolic reflectors (Section 2), we wanted to assess whether fibres with diameters on the wavelength-scale might offer an advantage over their macroscale counterparts. As the dimension under scrutiny is transverse to the direction of light propagation down the fibre, the issue of pump absorption for nano-scale layers discussed earlier (Section 3.3) is not limiting here. As a nano-

scale fibre will only support a small number of guided modes the question of how well a point dipole source will couple to them is of interest. The confined fields in the fibre might offer an enhanced Partial Local Density of States (PLDOS) for certain dipole orientations affecting the trapping fraction and even the fibre output angular distribution [74].

### 4.5.2 Method - Modified Born Series

Given the considerable size of the 3D fibre geometries we were interested in simulating and the availability of expertise within our academic group in this alternative approach, we chose not to use commercial finite element method (FEM) based software (COMSOL). Instead we've used a modified Born series based solver developed by a member of our academic group at the time, Tom Vettenburg. It offers a much improved scaling of the high-speed memory and time requirements with the size of the system vs traditional FEM and Finite-difference time-domain techniques and more transparency and physical intuition to its inner workings than a commercial package.

A detailed description of the method employed by the 'modified Born series' approach and proofs of its convergence and robustness to errors can be found in [93]. For the simplest case of a non-magnetic medium with isotropic (scalar) permittivity $\epsilon(\boldsymbol{x})$, where $\boldsymbol{x}$ is the spatial coordinate, the electric field (at fixed frequency $\omega$) $\boldsymbol{E(x)}$ is given by the vector Helmholtz equation

$$\nabla \times \nabla \times \boldsymbol{E} - k_0^2 \epsilon(\boldsymbol{x}) = \boldsymbol{S(x)}, \tag{4.13}$$

where $k_0$ is the wavenumber in free space. $\boldsymbol{S(x)}$ represents the source of radiation. To calculate the electric field $\boldsymbol{E(x)}$ for a given source term $\boldsymbol{S(x)}$ and permittivity distribution $\epsilon(\boldsymbol{x})$ the Helmholtz operator $\mathcal{O} = -k_0^2\epsilon(\boldsymbol{x}) + \nabla \times \nabla \times$ must be inverted. The direct approach, numerical inversion of $\mathcal{O}$ written as a matrix in a basis of eigenfunctions becomes too time and memory heavy for even remotely large systems. An iterative approach is therefore commonly used, such as in finite element methods [93, 94].

The iterative approach of interest is that of the Born Series, where the Helmholtz operator $\mathcal{O}$ is split into two parts, one associated with propagation through a homogeneous

medium $\mathcal{O}_h$ and another $\mathcal{O}_i$ due to the medium inhomogeneities [93, 95],

$$\mathcal{O}^{-1} = \left(\mathcal{O}_h + \mathcal{O}_i\right)^{-1} = \left(\mathbb{1}_3 + \mathcal{O}_h^{-1}\mathcal{O}_i\right)^{-1}\mathcal{O}_h^{-1}, \tag{4.14}$$

$$= \left[\mathbb{1}_3 - \left(\mathcal{O}_h^{-1}\mathcal{O}_i\right) + \left(\mathcal{O}_h^{-1}\mathcal{O}_i\right)^2 + \ldots\right]\mathcal{O}_h^{-1}, \tag{4.15}$$

$$= \left[\sum_{p=0}^{\infty}\left(-\mathcal{O}_h^{-1}\mathcal{O}_i\right)^p\right]\mathcal{O}_h^{-1}, \tag{4.16}$$

where $\mathbb{1}_3$ is the identity operator in three dimensional space. The inverse of the operator $\mathcal{O}_h$ is the simple solution for a homogeneous medium available analytically and the series can be understood intuitively as follows. The $p = 0$ term ($\mathcal{O}_h^{-1}\boldsymbol{S}$), or initial guess, can be understood as the field generated by the source term in the absence of any inhomogeneities. Each additional term then adds the fields due to scattering events of higher and higher order. The issue with this expansion is that it converges only for weak scattering systems and hence this method has been of limited use in the past. As discussed in [93], this has physical origin in the existence of bound states [95].

A recently proposed modified Born series [96], overcame this convergence issue by means of a pre-conditioning step and intelligent choice of background permittivity. The generalization to vector waves was demonstrated in [97] and the extension to magnetic and anisotropic materials by Vettenburg in [93]. The vector version of the modified Born series is as follows

$$\boldsymbol{E} = \left[\sum_{p=0}^{\infty}\boldsymbol{M}^p\right]\Gamma\overleftrightarrow{\boldsymbol{G}}\boldsymbol{S}, \qquad \text{where} \qquad \boldsymbol{M} \equiv k_0^2\Gamma\overleftrightarrow{\boldsymbol{G}}\chi - \Gamma + \mathbb{1}_3. \tag{4.17}$$

The Dyadic Green's function $\overleftrightarrow{\boldsymbol{G}}$ is defined as ($\mathcal{O}_h^{-1} \equiv \overleftrightarrow{\boldsymbol{G}}$), the solution to

$$\nabla\times\nabla\times\overleftrightarrow{\boldsymbol{G}}(\boldsymbol{x},\boldsymbol{x}') - k_0^2\alpha\overleftrightarrow{\boldsymbol{G}}(\boldsymbol{x},\boldsymbol{x}') = \mathbb{1}_3\delta^{(3)}(\boldsymbol{x}-\boldsymbol{x}'), \tag{4.18}$$

where $\alpha$ is the constant complex background permittivity of the system and $\chi$ the isotropic susceptibility with respect to $\alpha$. $\mathcal{O}_i = -k_0^2(\epsilon(\boldsymbol{x}) - \alpha) \equiv -k_0^2\chi(\boldsymbol{x})$. Convergence is shown ([96, 97]) to occur when the preconditioning operator $\Gamma = \frac{i}{\alpha_i}\chi$, where $\alpha_i$ is the imaginary component of the background permittivity, which must also satisfy $\alpha_i > \max_x(|\epsilon(\boldsymbol{x}) - \alpha_r|)$. It should be noted that the convergence speed is approximately inversely proportional to the range of susceptibilities in the simulation domain and is therefore not well suited to handling metals [93]. Finally, the system must exhibit no gain and the permittivity contain no singularities [96].

The so-called 'Macroscopic Maxwell' solver code developed by T. Vettenburg is available

to download with instructions and examples at [98], both in the Python and MATLAB languages. The MATLAB version of this code was used to assess the trapping fraction for fluorescent YAG fibres supporting a small number of guided modes. The details of how the simulation domain was defined and resulting electromagnetic fields processed are reserved for Section 4.5.4.

### 4.5.3 Results and Discussion

For an oscillating electric dipole of a given position and orientation within a YAG (n=1.83) fibre, the electromagnetic fields were calculated using the aforementioned solver to deter-

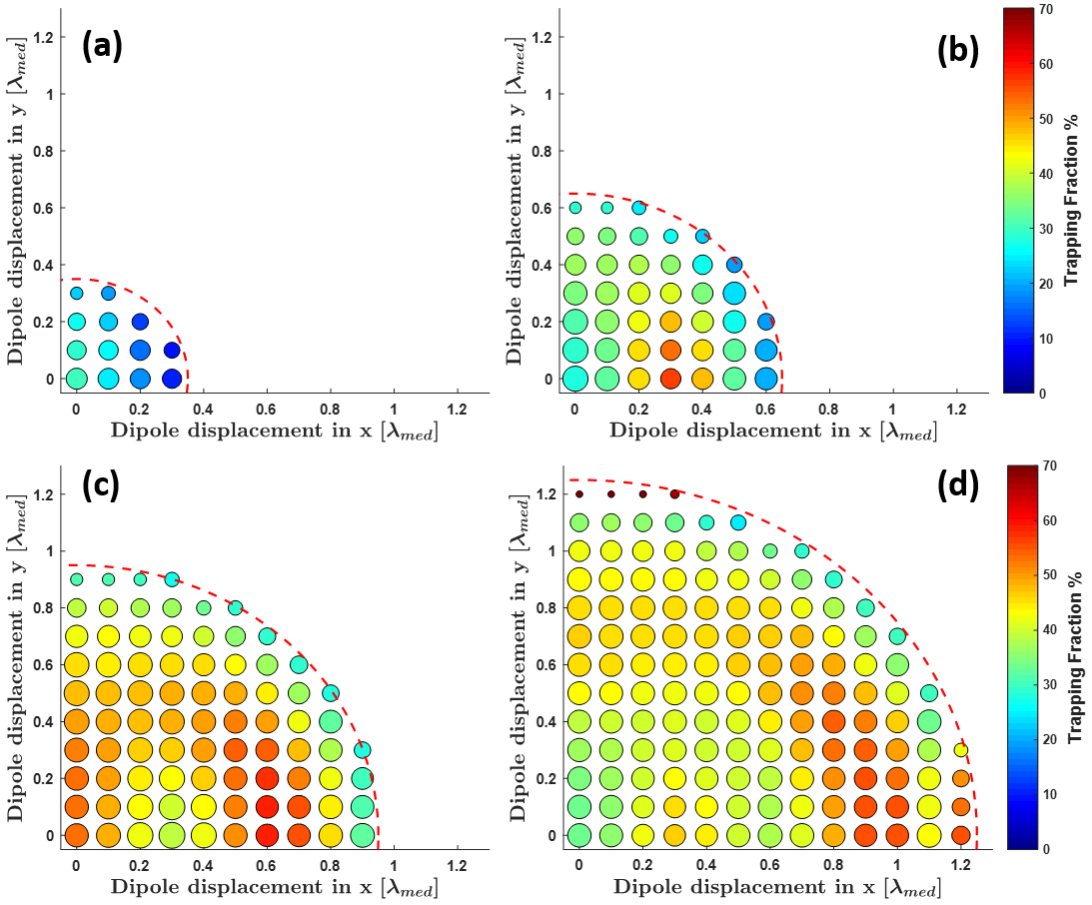


**Figure 4.19:** Spatially varying trapping fraction and LDOS for dipole moments oscillating in the $\hat{y}$ direction, for 4 different YAG fibre diameters. Results of simulations. The LDOS is proportional to the area of the marker and the trapping fraction is given by its colour. The dashed red line marks the edge of the fibre. The fibre diameters, in units of the emission wavelength in the fibre medium $\lambda_{med}$ are 0.7, 1.3, 1.9 and 2.5 for (a) - (d) respectively.

mine the trapping fraction and total emitted power (indicating changes in the Local Density Of States (LDOS)). Much like for the wavelength scale reflectors of Section 2, for a dipole moment orientated parallel to the fibre axis only a sweep of radial position is required and for a dipole moment perpendicular to the axis one quadrant of the fibre cross section must be swept. This process was repeated for regularly spaced dipole positions in fibres of diameter ranging from 0.3 to 2.5 (in increments of 0.2) times the emission wavelength in the medium. The sampled dipole orientations and positions were then combined to determine the mean trapping fraction (and density of states) for a fluorescent cylindrical fibre of the specified diameter, assuming uniform distribution of emitters. The specifics of the simulation settings and post-processing can be found in Section 4.5.4.

Fig.4.19 shows 4 examples of these quadrant source position sweeps, for a dipole moment in $\hat{y}$ (Cartesian coordinates with the fibre axis along $\hat{z}$). The trapping fraction is given by the colour scale and the relative PLDOS by the marker size (area). Fig.4.19(a) is the smallest size at which light is substantially guided by the fibre, (d) is the largest diameter sampled and (b) and (c) are representative examples that lie between these sizes. The spatially vary-

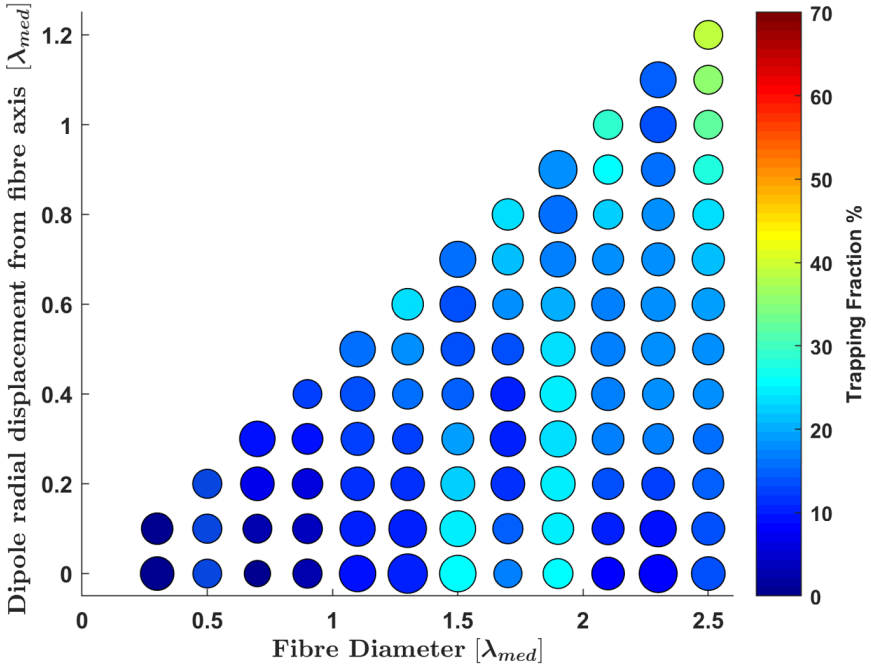


**Figure 4.20:** Radially varying trapping fraction and LDOS for dipole moments oscillating in the $\hat{z}$ direction, for a range of YAG fibre diameters. Results of simulations. The LDOS (or total emitter power) is proportional to the area of the marker and the trapping fraction is given by its colour. $\lambda_{med}$ is the emission wavelength in the fibre medium.

ing efficiency with which the emitters couple into the guided modes of the fibre can be seen, along with some significant modifications to the PLDOS when the emitter is near the fibre edge. The Trapping fraction values are generally quite high, as a perpendicular dipole has its maximum emission direction aligned with the fibre axis.

Fig.4.20 combines the one dimensional sweeps of radial displacement for a dipole moment oscillating in $\hat{z}$, for the full range of fibre diameters into a single plot. The trapping fraction is unsurprisingly significantly lower for a dipole radiating primarily radially outward from the fibre centre.

The mean trapping fraction for emitters uniformly occupying the fibre cross section, for each dipole moment orientation and their mean can be seen in Fig.4.21. At this scale the calculated mean trapping fractions are lower than in their geometrical-optics counterparts (Fig.4.3), with a general increase with fibre diameter. Fibre dimensions larger than

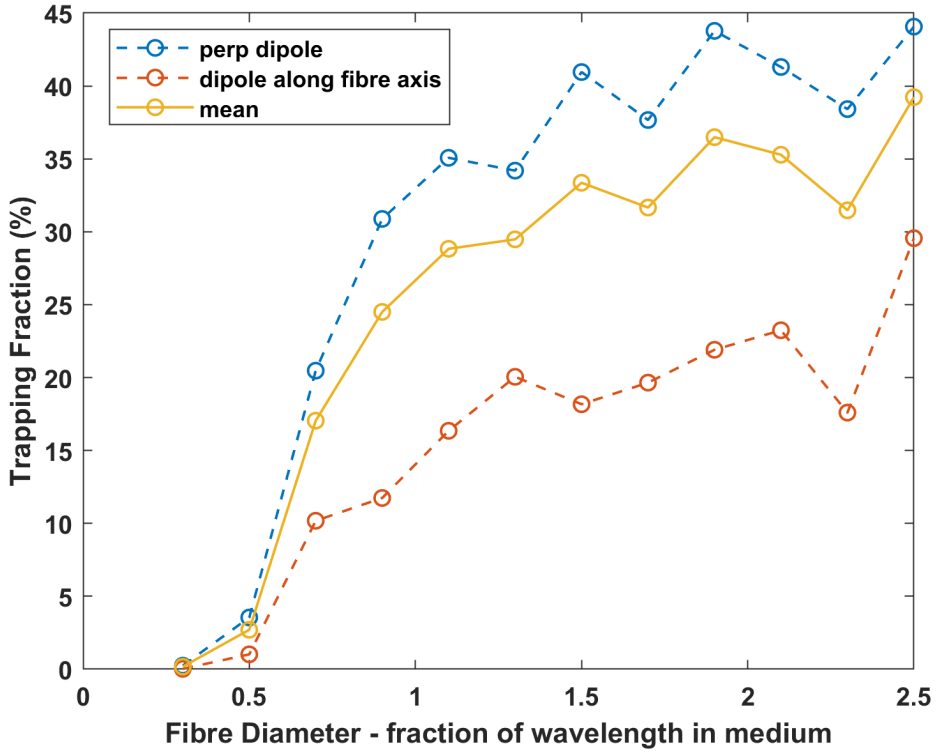


**Figure 4.21:** Spatially averaged trapping fractions for a uniformly doped fibre of varying diameter. Dependence for oscillating dipole moments parallel and perpendicular to the fibre axis as indicated. The mean solution for a collection of randomly orientated emitters is $\frac{2}{3}$ perpendicular orientation plus $\frac{1}{3}$ parallel. The lines connecting data point are visual aids only.

$2.5\lambda_{med}$ have been sampled, but due to the requirement to keep the voxels sufficiently sub-wavelength these become computationally intensive even for this solver. At a diameter of $5\lambda_{med}$ we see a variation of the orientation averaged trapping fraction from 40% to 66% between the fibre center and edge. This follows the expected trend towards the geometrical optics case of 45% to 80%.

The total power variations with fibre diameter are shown in Fig.4.22 in the same format as in Fig.4.21. The difference in PLDOS for the parallel and perpendicular dipole orientations is less than a factor of two once the diameter is large enough to provide an acceptable trapping fraction (see Fig.4.19 and Fig.4.20). This difference tends to zero as the fibre diameter increases. A reference point for the relative total emitted power in bulk YAG was calculated from that of an emitter in free space (× 1.83), with a simulated free space power of $1.03\times10^{-15}$. However upon testing the scaling of this value with changes in the refractive index of the entire simulation domain the results deviated from the expected linear dependence [99]. A coefficient of proportionality of 1.07 over an index range of 1 to 4 was calculated using a least squares linear fit.

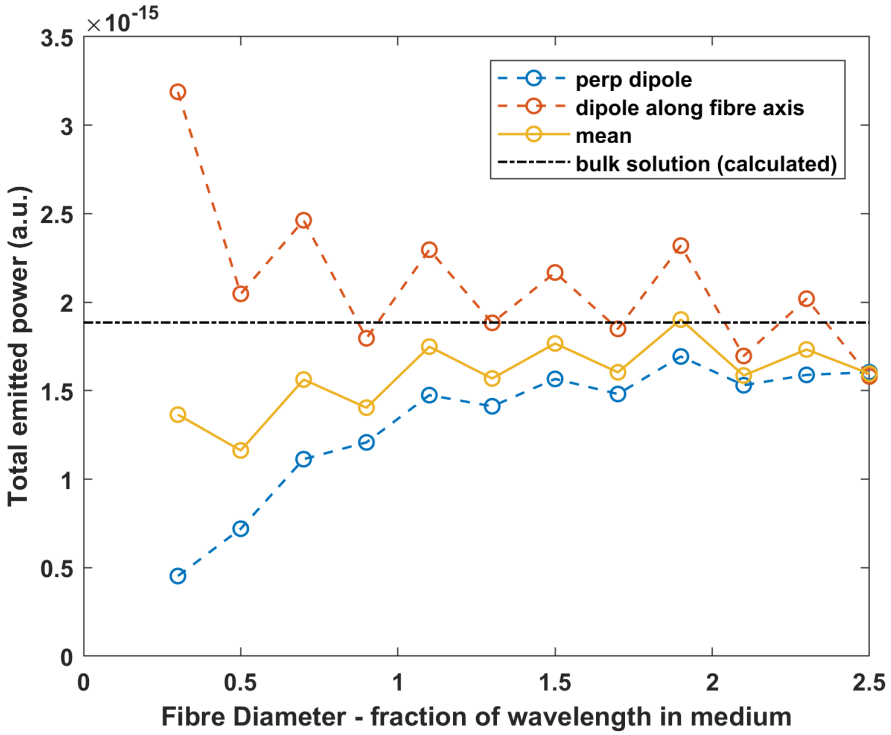


**Figure 4.22:** As in Fig.4.21 but showing variations in the spatially averaged total power radiated for a given fibre diameter. The power values have no absolute units and the calculated bulk emitted power (1.83 × that of a simulated free space dipole) is subject to some error as discussed in the text.

The results presented were deemed sufficient evidence that nanoscale fibres would offer no clear advantage over their larger (geometrical optics regime) counterparts. The modest trapping fractions calculated and similar PLDOS for perpendicular dipole moment orientations appear to offer no benefits to compensate for the added complexity. We therefore chose to focus our efforts on the large scale fibres following the geometrical optics approximation.

### 4.5.4 Simulation Details

The simulation domain size was fixed for all values of fibre diameter in the parametric sweep, ranging from 0.3 to 2.5 (0.2 increments) times the emission wavelength within the fibre. The free space wavelength was set to 500 nm, so for a diameter factor of 2 the physical size of the fibre was $\frac{2\times\ 500\ nm}{1.83}$, where 1.83 is the refractive index of YAG. The simulation domain was a 64 × 64 × 256 array, where the longer direction is along the length of the fibre, assigned the $\hat{z}$ direction. Each voxel had a width of $\frac{1}{10}$th the emission wavelength in the fibre medium.

The permittivity of voxels within the fibre diameter (which is centred in x and y and extends to the simulation boundaries in z) is set to $1.83^2$ and elsewhere is $1^2$ to simulate air cladding. At the simulation boundaries a form of anti-reflective / absorbing boundary condition was applied. Within a width of 0.6 × 500 nm of the simulation edges the imaginary component of the permittivity was linearly increased from inner to outer edge. This was done with a weighting of $0.8i$ × distance in boundary / boundary thickness. $i$ is the imaginary unit.

The source term which was supplied to the solver is a 4D matrix of the simulation size plus an additional leading dimension capturing the vectorial nature of the source. The orientation of the oscillating dipole is determined by the value of this additional dimension at the source location, with 1, 2 and 3 indicating the x, y and z directions respectively. A single non-zero term at the simulation centre defined a current density oscillating at the simulation angular frequency $\frac{2\pi c}{\lambda}$. The source amplitude was specified as $i\mu_0\omega$, where $\mu_0$ is the permeability of free space.

The arguments supplied to the function 'SolveMacroscopicMaxwell' are as follows:

- ranges : a 64 × 64 × 256 cell array of ranges to calculate the solution at.
- $k_0$ : the vacuum wavenumber, = $\frac{2\pi}{\lambda}$
- $\epsilon$ : a 64 × 64 × 256 array defining the spatially varying permittivity of the simulation domain.

- $\xi$ and $\zeta$: two terms reserved for bi-(an)isotropy and were both supplied as a single zero element as none was required.
- $\mu$: the equivalent term to $\epsilon$ but for spatially varying permeability, here supplied as 1 (uniform).
- source : a $3 \times 64 \times 64 \times 256$ cell array defining the vector source as discussed in the text.
- a handle to a progress function that tracks convergence of the calculated fields

Once the progress function determines that the estimate of the relative error in the solution is below 0.001 a 4D array containing the vector electric fields of the solution is returned. Much like for the source, the leading dimension can have value 1, 2, or 3 specifying the x, y or z component of the electric field.

The magnetic field $\boldsymbol{H}$ is determined from the electric field using the included 'calcCurl' function with prefactors as follows $\boldsymbol{H} = \frac{-i}{c\mu_0 k_0} \nabla \times \boldsymbol{E}$, [Here $\times$ is the vector (or cross) product]. From which the time averaged Poynting vector $\langle \boldsymbol{S} \rangle = Re(\boldsymbol{E} \times \boldsymbol{H}^*)$ was determined, where $Re$ denotes the real component and $^*$ denotes the complex conjugate.

The total power exiting the simulation domain was determined by integrating the normal component of the time averaged Poynting vector on the closed surface of a cuboid of the simulation domain size minus the width of the absorbing boundary regions. This provides the total power emitted for a given configuration. The total power propagating within the fibre was calculated by integrating the z component (parallel to fibre-axis) of the Poynting vector at the boundaries as above, but only within the cross-section of the fibre. The trapping fraction is simply the ratio of trapped to total emitted power.

## 4.6 Prototype YAG:$Ce^{3+}$ Light Guides

### 4.6.1 Evolution of Prototypes, Sample Descriptions

Our initial efforts to produce 'fibre-like' structures for demonstration purposes used a commercial YAG:$Ce^{3+}$ powder suspended in an optical adhesive. The fibre shape was provided by hollow translucent plastic tubing, into which a mixture of the power and adhesive was encouraged and cured with a long-wave UV light source. The result is a fibre or rod with some light guiding ability, but with high scattering losses due to the powder inclusions and poor index contrast between the core (n=1.64) and scattering cladding (assumed n~1.5). We also produced samples where the plastic tube structure was replaced by a smooth rolled tube of aluminium foil. Examples of both types of sample can be seen in Fig.4.23.

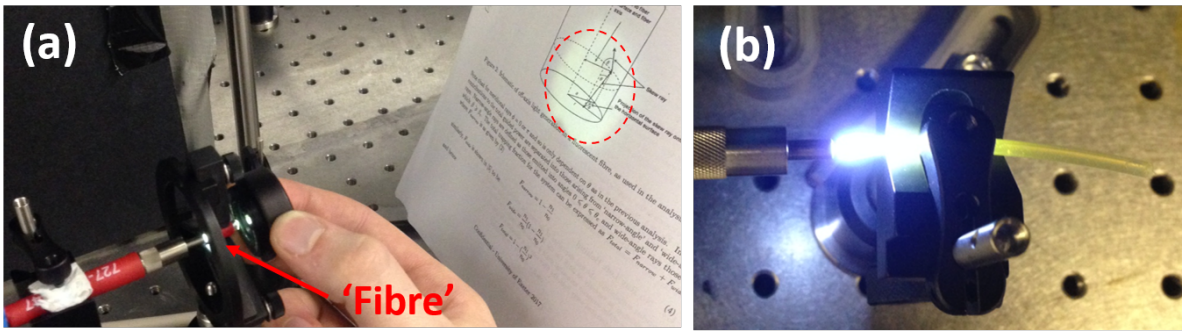


**Figure 4.23:** Powder phosphor fibre prototypes, under 450 nm blue LED illumination via an illuminating fibre. (a) Demonstration of directional emission (red dashed circle), achieved with external fibre collimation lens, for a phosphor filled aluminium tube. (b) High levels of scattering / leakage from a powder in plastic tube fibre.

The commercial phosphor powder used was Intematix NYAG4454-L, a cerium doped YAG phosphor with mean particle size of approximately 13 $\mu$m [100]. The optical adhesive was NOA164 from Norland, of refractive index 1.64 and curing soft, rubbery and tacky in air [101]. The powder was mixed with the adhesive in the ratio 2.4 g adhesive to 0.03 g powder.

With these crude prototypes we were able to demonstrate enhanced control over the angular distribution of emission (due to the small cross-sectional area (CSA)) and colour tuning by varying the length of the fibres.

To produce a fibre prototype of acceptable optical efficiency, a scatter-free phosphor would be required. At this stage we had two commercially available phosphor components at our disposal, both of which were highly scattering. The first, a phosphor powder in host material sample, see Fig.4.24(a) and the accompanying description in the text, which scatters visible light due to the index contrast between the host material and the powder and the powder particle size. The other sample, of the emerging technologies that remove the need for host material to suspend the phosphor, is a monocrystalline YAG:$Ce^{3+}$ disc Fig.4.24(b). In this case the material is scatter-free, but has had surface roughness added to achieve better colour uniformity in the output light. Efforts were made to remove the surface roughness with diamond based polish and cut thin strips of disc to produce fibre shaped scatter free phosphor. These were ultimately unsuccessful and soon superseded by the next generation of prototype.

A set of 10 bespoke moncrystalline YAG:$Ce^{3+}$ light guides were ordered from the company who manufactured the monocrystalline disc Fig.4.24(b), Crytur [80]. To demonstrate the properties of the fibre geometry at a scale which would be convenient for optical excita-

tion and yield white output the concentration of cerium dopant in the phosphor needed to be significantly less than that used in typical phosphor solutions. See Fig.3.1 in Section 3 for the pump absorption achieved in a sample of typical concentration, with absorption coefficient $\alpha \simeq 75$ cm$^{-1}$. The bespoke monocrystalline samples were machined from a boule of very low concentration YAG:Ce$^{3+}$, with absorption coefficient quoted as 'less than 1 cm$^{-1}$'. This moves the distance for equivalent absorption to the ~200 $\mu$m samples to several centimeters. As the samples are machined from a boule, as opposed to drawn like fibre, we were advised that a cuboidal guide with square cross-section would be significantly easier to produce. The guide dimensions were thus 2 × 2 × 100 mm, with all sides highly polished to preserve light guiding ability.

Prototypes made from the square cross-section samples can be seen in Fig.4.24(c),(e) and (g) and the full manufacturer specifications in Table 4.1. Fig.4.24(c), an unmodified 'rod' as supplied by Crytur. Due to the very low dopant concentration, it appears uncoloured and transparent from the side, but fluorescence from ambient light is guided to the ends producing a strong yellow colour. As discussed in Section 4.2.3 and quantified in Section 4.3.3, a simple hemispherical lens can drastically improve device efficiency by allowing almost all of the generated fluorescence to escape the guide end. A 6 mm diameter sapphire (n = 1.77) hemispherical lens has been added (high index glue) to an unmodified guide to form the sample in Fig.4.24(e). In Fig.4.24(g) the length of the guide was first reduced in an effort to tune the amount of residual blue pump and achieve white output. The colour metrics and efficiency of all the prototypes in Fig.4.24 are discussed later in Section 4.8.

Due to the square cross section of these samples, they are unable to demonstrate the radially dependent trapping fraction we expect from simulations with circular cross-section guides (Fig.4.3 vs Fig.4.7). After some experimentation with fluorescent acrylic rods (in the next Section 4.7) we were able to obtain monocrystalline YAG:Ce$^{3+}$ rods with a circular cross-section from Crytur. These can be seen in Fig.4.24(d),(f) and (h). The modifications follow the same logic as for the square cross-section guides, but as the cylinders have a larger diameter of 5 mm, a larger hemispherical lens of differing material (due to availability) was used. A 10 mm diameter half ball lens made of material S-LAH79 (n=2) was used.

Interlinked with the question of "How will the overall efficiency of the fluorescent guide type phosphors compare with the predictions for loss-free structures?" is"Will the position dependent trapping fraction be observable in these cylindrical prototypes?". We have seen in Fig.4.18 that if the surface quality of the 'fibre' is not high enough, scattering losses will dampen the radial dependence to the point where it may no longer be detectable. Demonstrating this radial dependence experimentally is the subject of Section 4.7.

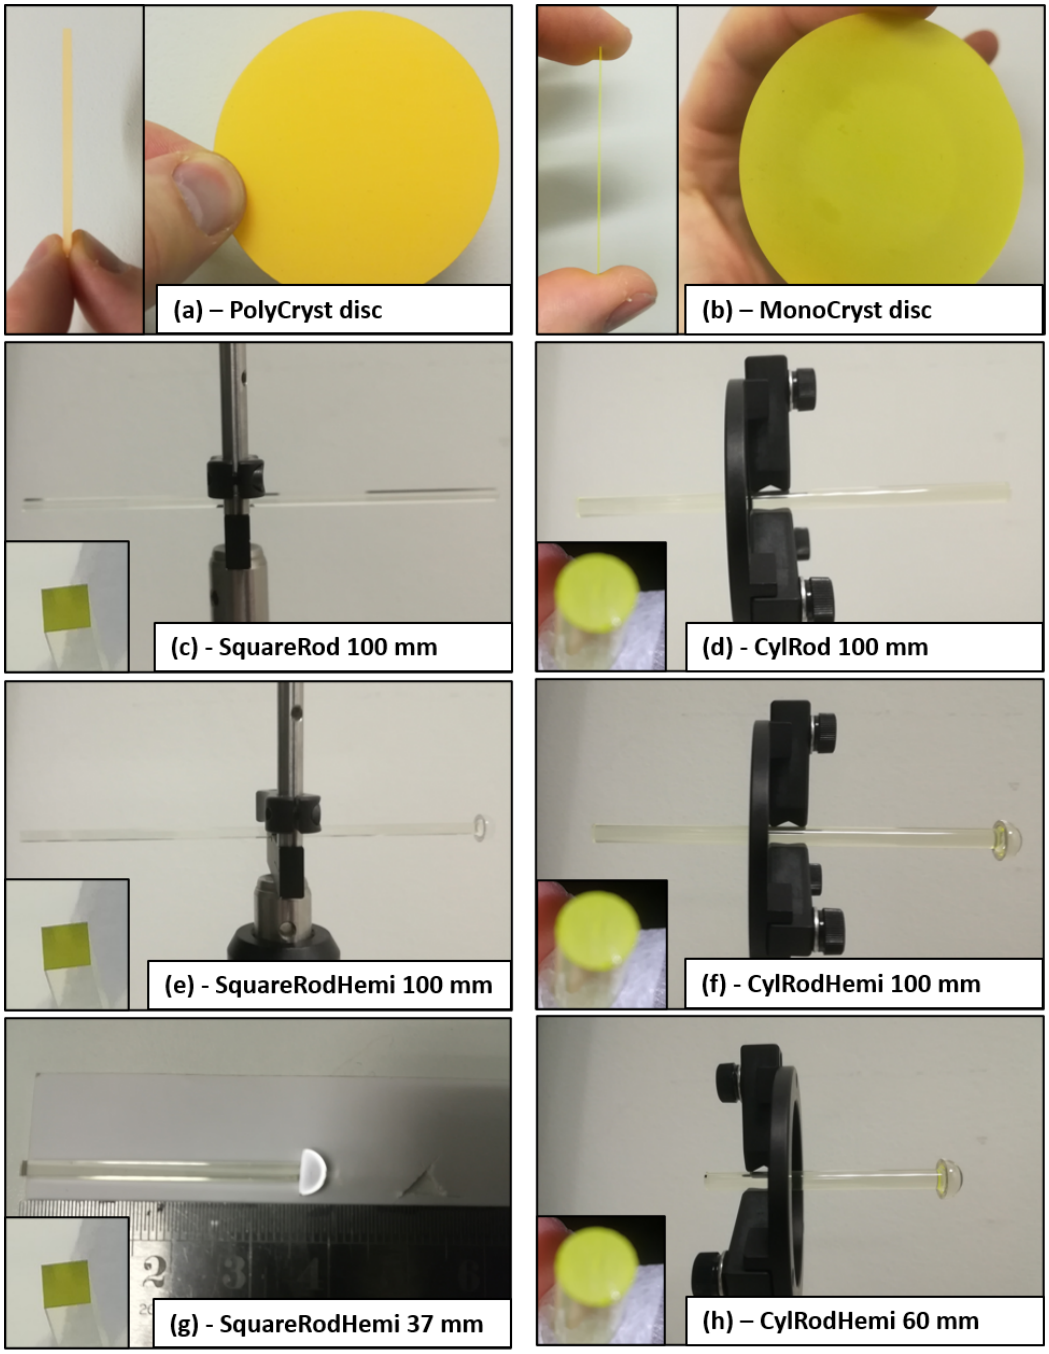


**Figure 4.24:** Images of phosphor samples of significance, establishing shorthand designation for each, full sample details in Table 4.1. (a) commercial polycrystalline phosphor from Intematix with high CRI. (b) monocrystalline YAG:Ce3+ disc, 200 $\mu$m thick, roughened surfaces available from Crytur [80]. The remaining samples are bespoke low concentration phosphor light guides from Crytur. They are either 2 mm x 2 mm square cross-section (CS) or 5 mm diameter circular CS, as indicated by the inset image in each case. Hemispherical lenses of diameter 6 and 10 mm have been added to the square CS and cylindrical rods where indicated.

| Sample | Image location in Fig.4.24 | Description<br>Manufacturer specifications + modifications |
|---|---|---|
| PolyCryst disc | (a) | ChromaLit Round 90 CRI, 2700K, 61.5 mm diameter. A powder (Polycrystalline) 'phosphor composite' on a polymer substrate from Intematix [102]. |
| MonoCryst disc | (b) | YAG:Ce single crystal phosphor plate, 200 $\mu$m thick, 54 mm diameter with roughened surfaces. Crystal growth by the Czochralski method [103]. Obtained from Crytur [80]. Intended output CCT of 5700K (blue pump, wavelength unknown). Manufacturer code: LWY\|CIE(0.434,0.543)\|D55\|C0\|200\|G/GN/N\|\|5700\| |
| SquareRod (Bare) 100 mm | (c) | Material: YAG:Ce, Low absorption coefficient < 1cm-1 (blue pump), Surface: All sides polished, Dimensions: (cuboidal) 2.0 (+/-0.1) x 2.0 (+/-0.1) x 100 (+/-0.2)mm, Code: LCY\|CIE(0.418,0.554)\|2x2x100\|\|\|A/N\|\|\|C, from Crytur [80] |
| CylRod (Bare) 100 mm | (d) | Material: YAG:Ce, Low absorption coefficient < 1cm-1 (blue pump), Surface: All sides polished, Dimensions: (cylindrical) 5.0 (+/-0.1) x 100 (+/-0.2)mm, Code: LCY\|CIE(0.418,0.554)\|dia.5x100\|\|\|A/\|\|\|C, from Crytur [80] |
| SquareRodHemi 100 mm | (e) | As in 'SquareRod (Bare) 100 mm', with a 6.0mm Diameter, Sapphire Half-Ball Lens (n=1.77) attached using Norland optical adhesive NOA 164 (n=1.64). |
| CylRodHemi 100 mm | (f) | As in 'CylRod (Bare) 100 mm', with a 10.0mm Diameter, S-LAH79 Half-Ball Lens (n=2.0) also attached using NOA 164. |
| SquareRodHemi 37 mm | (g) | As in 'SquareRodHemi 100 mm', but first the 100 mm rod was cut to 37 mm with a diamond coated cutting wheel on a rotary tool, both available from high-street hardware stores. As a consequence of the soft and rubbery cured state of the adhesive the lens had become misaligned prior to photographing but after all measurements had been performed. |
| CylRodHemi 60 mm | (h) | As in 'CylRodHemi 100 mm', but first the 100 mm rod was cut to 60 mm with a diamond coated cutting wheel on a rotary tool. |

**Table 4.1:** Detailed descriptions of the samples photographed in Fig.4.24

## 4.7 Demonstrating a Radially Dependent Trapping Fraction Experimentally

Fig.4.3 shows that a cylindrical fluorescent fibre has the potential to be considerably more efficient than one would assume if neglecting the contribution of skew rays (those originating off axis). The potential discrepancy could be as large as that between a 45% and 80% trapping fraction if illuminating with a beam or ring much smaller than the fibre diameter, or between 45% and 70% (Fig.4.6) if the entire fibre cross-section is illuminated. We have seen in Fig.4.18 how losses due to surface scattering affect rays originating far from the fibre axis more than those near its center, reducing the magnitude of the improvement one expects towards the fibres edge. The question remains, can one exploit, or even detect this effect in a 'real-world' sample with unquantified losses present?

As the square cross-section rods are not predicted to exhibit any pump position dependent behaviour and it was unclear if we would be able to obtain cylindrical guides from Crytur, we looked for more easily accessible systems that might exhibit the same behaviour. Fortunately, there exists a range of fluorescent acrylic rods of decent optical quality, generally under the description 'light gathering rods'. An assortment of light gathering rods can

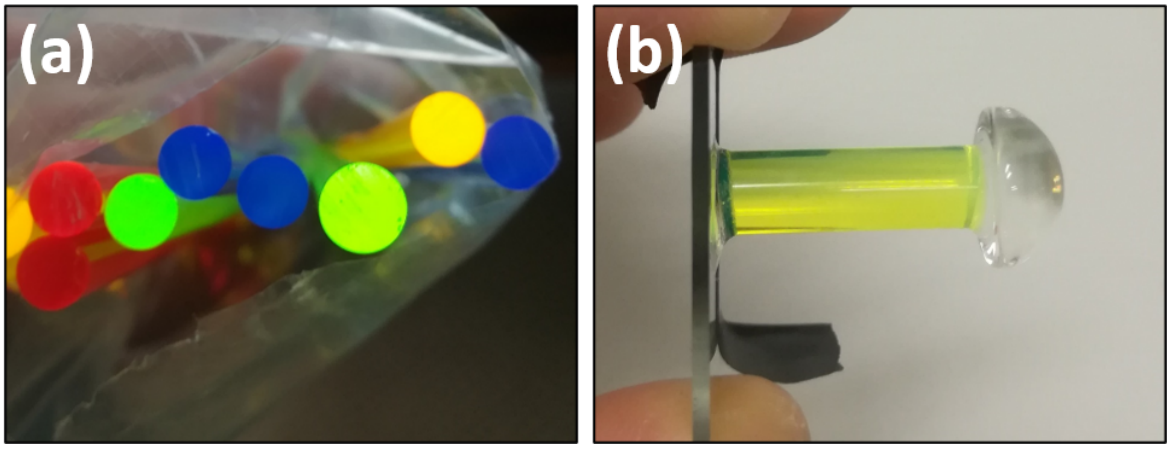


**Figure 4.25:** Fluorescent 'light gathering rods' used for investigating pump position dependent trapping fractions. (a) An assortment of so-called light gathering rods, transparent (outside of spectral absorption bands) fluorescent acrylic rods. 6 mm in diameter, obtained from 'https://www.homecrafts.co.uk/', now out of business but widely available at online craft suppliers. (b) A modified rod from (a), cut to 20 mm length with a glass slide and a 12 mm hemispherical glass lens (ebay) attached via index matched optical adhesive (NOA68) to either end. The black electrical tape visible is for ease of handling and does not interact optically under the measurements performed.

be seen in Fig.4.25(a), purchased for approximately £40 (details within caption).

To probe the varying trapping fraction across our samples we constructed an illumination system that yielded a movable thin beam of light that is roughly collimated over the length of the sample. The experimental setup is as shown in Fig.4.26. The lenses L1 and L2 are Plano-convex with $f$ = 300 mm. The purpose of L1 is to collimate the large diverging beam from the 450 nm diode 'FlexPoint' laser (FP-D-450-5P-C-C). L2 then focuses the collimated beam to a point within the sample. As the focal length of L2 is reasonably large, the Rayleigh length of the beam is long enough to maintain a thin beam width along the full length of the samples.

The 1D linear motorised stage and controller were obtained from Thorlabs, with item numbers PT1/M-Z8 and KDC101 respectively. The beam was scanned horizontally across the input face, passing through the center of the circular cross-section. The motorised stage was controlled via connection to a PC in LabVIEW, which was also used to collect power readings from the photodiode power sensor, a Thorlabs S120C connected to a Thorlabs PM100USB meter interface. The sensor had a filter F1 attached where indicated in the results and is butt-coupled to the integrating sphere measurement port. The two filters used for F1 are a 500 nm longpass (Thorlabs FEL0500) and a 450 nm passband filter of 10 nm FWHM (Thorlabs FB450-10). The integrating sphere used was a Newport 819C-IS-5.3 with the 2.5" port reduced to 1" with a 819M-PFR-1.0-2.5 reducer.

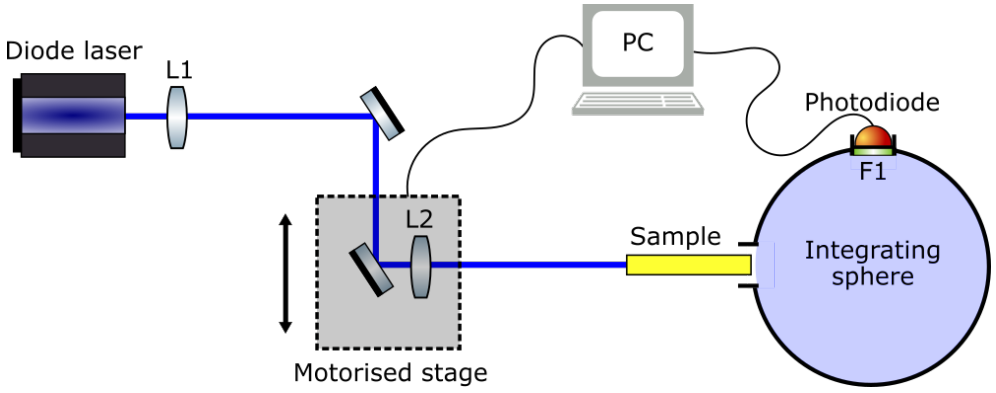


**Figure 4.26:** Schematic of experimental setup used to study dependence of sample optical output on the illumination position. A 450 nm diode laser source is focused to a thin approximately collimated beam within the sample. The incident position of the pump beam on the sample face is controllable via a 1D linear motorised stage. Light exiting the sample at the integrating sphere port is collected and evenly distributed across its inner surface. A photodiode records the amount of light detected from the sphere measurement port. A changeable filter F1 enables distinguishing of the pump light from that generated within the sample. Full details in the text.

The samples were positioned in the reduced port and measurements take from the 90° port, with the baffle between these ports. Samples were positioned at a consistent depth into the sphere port such that the end of the fluorescent rod was in line with the inner edge of the anodised black ring that frames the 1" port reduced opening. It is noted that as light from the sample is both diverging from the rod end and has a ballistic (blue) component which strikes the opposite side of the sphere, the measurement configuration (in terms of baffle use) is only optimised for the fluorescent component.

Fig.4.27 shows the variation in detected power with pump beam translation for a 15 mm long section of the green fluorescent acrylic rod shown in Fig.4.25(a). The green rods were selected for their strong absorption of the 450 nm pump light and similar fluorescence output spectrum to YAG:$Ce^{3+}$. The absorption coefficient for blue pump light is high, such that to the naked eye almost all pump light is absorbed within approximately 1 cm. The faces of the rod were finished with a facing tool in a lathe and the resulting concentric surface defects removed with fine sandpaper.

An increase in the amount of detected fluorescence can seen close to the edges of the sample. The edges are identifiable by symmetry and the drop / increase in fluorescence / pump light. It can also be seen that this behaviour is not replicated in the amount of pump light detected, so we can attribute this additional light to a reduction in the amount which leaks out of the sides of the fibre. However, the amount of additional light is small and does not appear to follow the smooth increase with distance from the rod center we expect from Fig.4.18. [It is noted that the refractive index of the rods (approx 1.5) is less than that used in the simulation, but qualitatively the same behaviour is observed in simulations for n = 1.5.]

Upon careful consideration it becomes unclear whether or not the predicted variation in trapping fraction should be expected to be measurable for this sample. This is due to the lack of outcoupling modification to the rod. Additional skew rays that are trapped as the pump moves off the rod central axis would be completely unable to escape a perfect rod, due to the total internal reflection experienced at the end facet. It is only due to the presence of defects in the fibre bulk and surfaces that these rays have the potential to escape. With the presence of losses also expected to mask the trapping fraction variation (Fig.4.18), one can now see that the results of Fig.4.27 are in fact quite encouraging.

Fig.4.28 compares the results of Fig.4.27 with a sample with improved outcoupling. It is 20 mm long with a 12 mm hemispherical glass lens on the exit face and a glass slide attached to the input face. Due to the index matching between the adhesive (NOA68) and the glass slide, scattering due to surface roughness on the input face is mitigated by its ad-

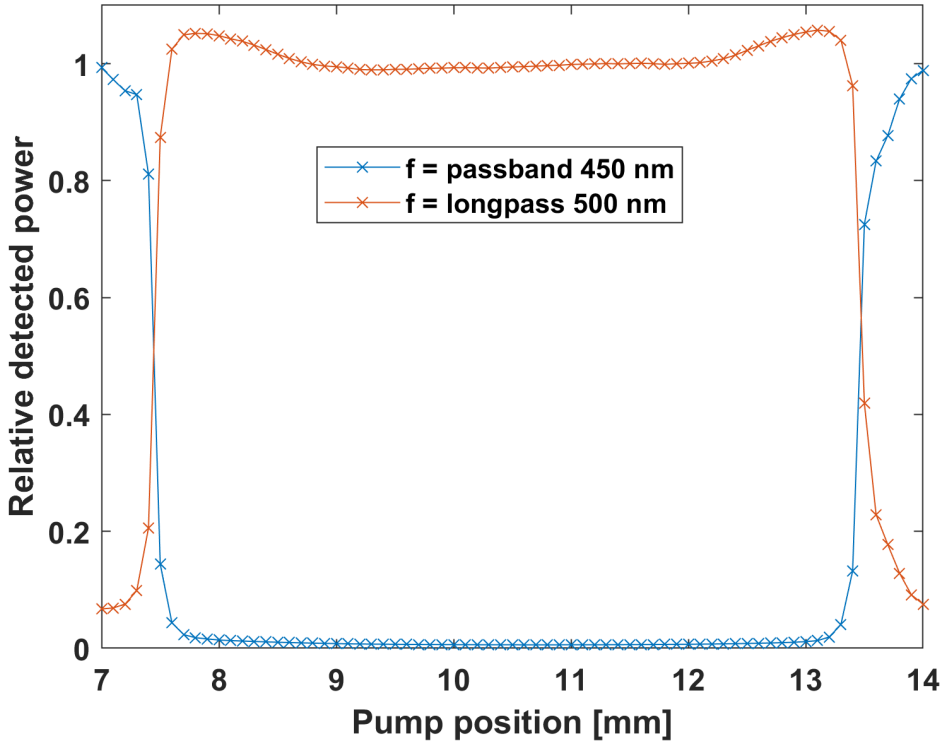


**Figure 4.27:** Variation in amount of light exiting a 15 mm long, 6 mm diameter green fluorescent acrylic rod as a 450 nm laser beam is scanned across its diameter. The two data sets distinguish light from the pump (450 nm 10 nm FWHM filter) from the fluorescence (longpass 500 nm filter). The longpass power was normalised by its pump position = 10.5 mm value and the passband power by its position = 15 mm value. The edges of the rod can be identified by sharp changes in detected power with pump position. An increase in detected fluorescence, from which we infer an increased trapping fraction can be seen close to the rod edges.

dition. The small change in length here was dictated by prepared sample availability and is not the variable under scrutiny in Fig.4.28.

A clear increase can be seen as expected and now the radial dependence follows qualitatively the trend that was predicted. The 'rounding off' of the data near the fibre edges is due to the finite width of the pump beam, causing the measurements to be the convolution of this approximately Gaussian beam shape with the trapping fraction dependence seen in Fig.4.18. The considerable amount of fluorescence observed when the pump is positioned outside of the fibre for the hemi lens sample is due to the reflections off the front face of the lens and refraction from adhesive over-spill around the rod end directing pump back into the fibre.

Fortunately, in addition to demonstrating the variable trapping fraction in an acrylic fibre we were able to purchase cylindrical monocrystalline YAG:$Ce^{3+}$ guides from Crytur, allowing us to see the magnitude of this effect in high optical quality samples and measure the efficiency of this next generation of prototype. The rods are 5 mm in diameter and 100 mm long, of the same dopant concentration as the square cross-section samples and highly polished. They can be seen in Fig.4.24(d),(f) and (h) and are described in detail in Table 4.1. The experimental setup for measuring these samples was the same as before (Fig.4.26) with one key difference, the photodiode and filter combination was replaced with a spectrometer, attached to the measurement port via 600 $\mu$m optical fibre.

The spectrometer is an ocean optics USB2000+ calibrated using a radiometrically calibrated HL-3P-CAL tungsten-halogen light source of known output power. The calibration lamp is not of the type designed for calibrating an integrating sphere measurement system,

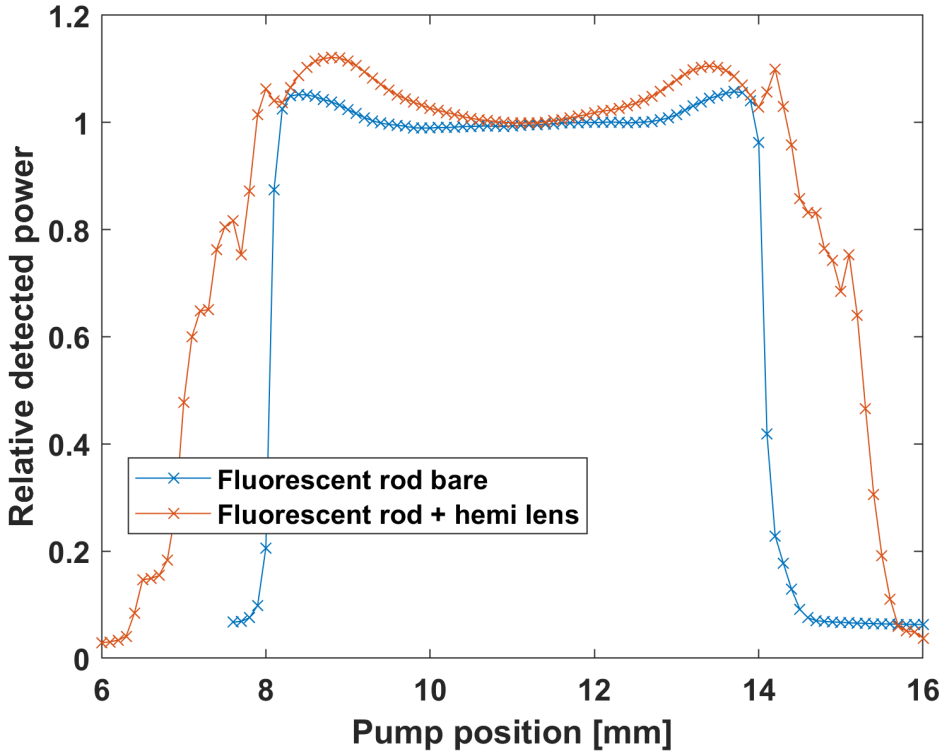


**Figure 4.28:** As in Fig.4.27, but comparing the bare 15 mm rod seen there with a 20 mm rod with a hemispherical lens on the output face and glass slide on the input face, as seen in Fig.4.25(b). The hemispherical lens allows efficient extraction of trapped light and reveals the predicted dependence on pump position (see Fig.4.18). The inclusion of the lens causes fluorescence to be generated at pump positions that are outside the limits of the rod. Both data sets have their power values normalised by the pump position = 11.5 power value.

which it is understood include and are calibrated for diffuse output from the light source. As such, the calibration was performed in the absence of the integrating sphere with the following optical path - light source - cosine corrector - 600 $\mu$m fibre - spectrometer. The calibration was performed using the OceanView spectroscopy software following the manufacturers instructions [104].

When the calibrated spectrometer is used with the integrating sphere (connected via fibre), the spectral response of the instrument response factors (such as grating and detector efficiency) will be compensated for, but the absolute power entering the sphere will remain unknown. The calibrated spectrometer was used in this way to obtain spectral power distributions where the absolute value of the 'relative power per nm' is unknown, but that subsequent spectra can be accurately and reliably compared with each other. In this way one can measure changes in output power with pump position or a lm/W$_{rad}$ figure by comparison of a pump spectrum with that from a phosphor sample.

A spectrum was obtained for pump positions at 0.1 mm intervals spanning the diameter of the samples and a little beyond either side. A constant integration time of 2000 ms was used and non-linearity correction and electric dark correction were both enabled. A dark spectrum was first acquired and then subtracted from subsequent measurements as per standard procedure [104]. Note the 2000 ms integration time causes the detector to become saturated when the pump laser enters the sphere unhindered. This means values for integrated power in the blue spectral region are unreliable outside of the fibre, however the behaviour here is only of qualitative interest and is not presented in the following results.

Fig.4.29 shows the integrated power in the range 480 - 700 nm for each spectrum obtained at varying pump position across two samples, the 100 mm square rod with hemispherical lens of Fig.4.24(e) and the 100 mm cylindrical rod with hemispherical lens of Fig.4.24(f). For the circular cross-section rod we observe a strong dependence on the pump position within the sample, of the form predicted in simulation and larger in magnitude than seen for the fluorescent acrylic rods (Fig.4.28). For the square cross-section rod, very little fluctuation is observed across the entire sample, with no indication of a radial dependence, as predicted. The spectra from each sample have not been normalised post spectral integration and therefore their relative amplitudes reflect directly the amount of light exiting the end of the sample. Indeed the square cross-section device outperforms the circular cross-section one.

It was expected that the surface quality of the cylindrical rod might be inferior to the square cross-section rod, due to the added challenges of obtaining the circular cross section from a boule of crystal. However this is unconfirmed and unquantified. This is a pos-

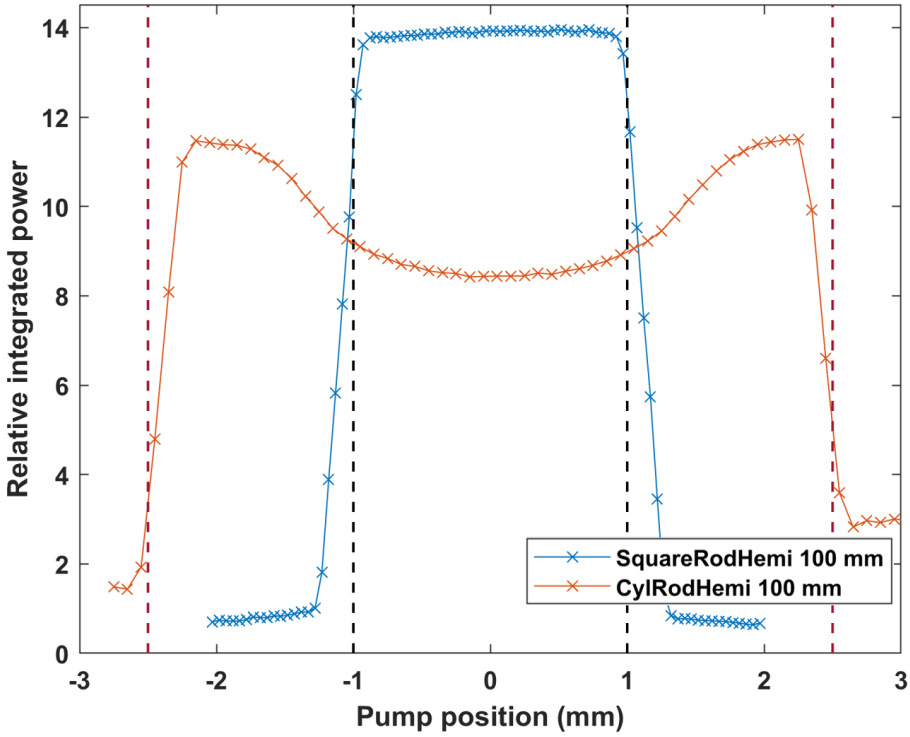


**Figure 4.29:** Variation in the amount of fluorescence exiting the end of a square vs circular cross-section monocrystalline YAG:$Ce^{3+}$ light guide, with outcoupling lens attached. Spectra obtained from the configuration depicted in Fig.4.26, but with photodiode and filter replaced with spectrometer, have been integrated in the spectral region 480 - 700 nm to determine changes in the relative power. The dashed lines mark the approximate locations of the edges of the samples - red, cylindrical sample, 5 mm diameter, black, square cross-section, 2 mm width. (In hindsight) the square sample was mounted at an angle, increasing its apparent width here.

sible explanation for the lower output from the cylindrical rod, in line with the predictions of Fig.4.18. Another likely contribution is the losses caused by the hemispherical lens. The Fresnel losses when passing from the rod, to adhesive, to lens are found to be comparably small in both cases, despite the differing refractive indices of the two lenses. The greater effect is likely to be the differing ratio of rod diameter to lens diameter in each case. Both ratios (6 mm / 2 mm for square, 10 mm / 5 mm for circular) should be large enough to reach the maximum out-coupled fraction seen in Fig.4.12. However, as the cylindrical rod has a lower diameter ratio it is more prone to experiencing additional losses due to imperfect alignment with the lens. We also note that an approximation is made by applying the results of Fig.4.12 to these two lenses, as both have refractive indices slightly shifted with

respect to the simulated index of 1.83.

Fig.4.30 compares the experimentally obtained positional dependence of the cylindrical sample in Fig.4.29 with the simulated lossless (P = 0) curve of Fig.4.18. The simulated data set has been zero padded outside of fibre limits and convolved with a Gaussian of $1/e^2$ beam width of 0.25 mm. The beam width was inferred from the experimental data. The experimental data set has been normalised to the same value as the simulation at the rod center. It can be seen in Fig.4.18 that with increasing losses the trapping fraction is reduced slowly for meridional rays at the fibre center. Choosing to normalise to the lossless trapping fraction value at the fibre center is motivated by the considerable trapping fraction variation observed experimentally, indicating the losses are in the regime where one would expect only a 2 or 3 % reduction in the central value.

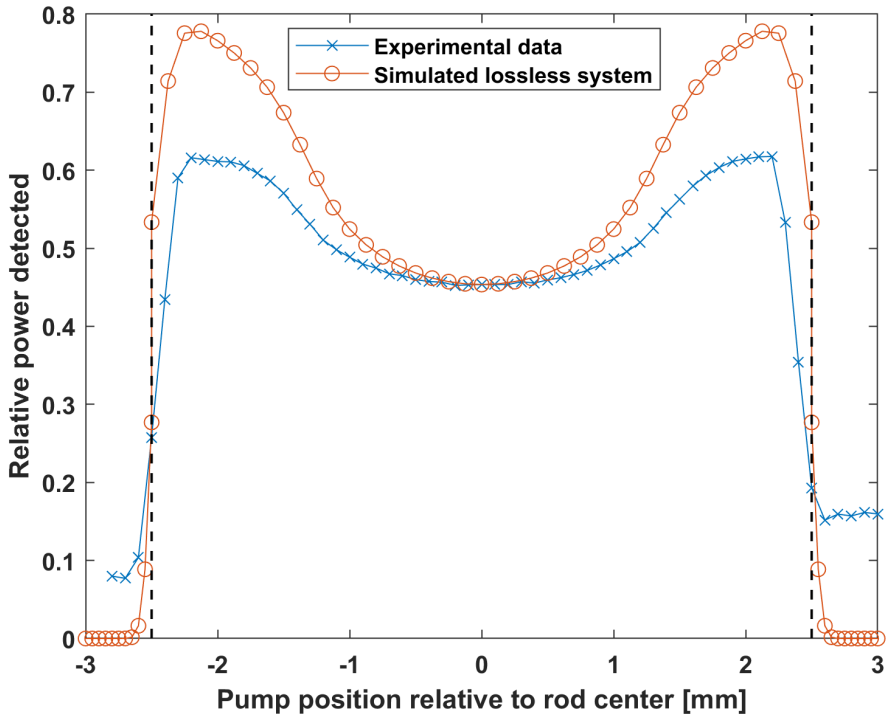


**Figure 4.30:** Comparison between the trapping fraction variation measured experimentally for the CylRodHemi 100 mm sample and the maximum variation possible, for a perfect lossless system. The simulated data comes from the previous lossless (P=0) curve of Fig.4.18, which has been convolved with a Gaussian to simulate the experimental beam width, details in text. The beam width has been inferred from the experimental data. The experimental data set has been normalised to 0.453 at the rod center, to align with the simulated data.

It is observed that the magnitude of the trapping fraction change measured experimentally is roughly half what one would expect of the perfect lossless system. It is anticipated that the magnitude of the variation measured could be increased by repeating these measurements on a shorter rod (less losses) and by reducing the beam width further.

## 4.8 Prototype Efficiency and CCT Results

Spectra obtained using the configuration depicted in Fig.4.26, but with a spectrometer in place of the photodiode, were used to assess the efficiency with which the fluorescence generated within a sample is directed out of its output facet. In addition these spectra were used to assess some of the key colour metrics of the light emitted from a given sample.

### 4.8.1 Efficiency and CCT Calculation Methods

A calibrated irradiance spectrum was acquired using OceanView spectroscopy software for the 450 nm pump beam entering the integrating sphere, in the absence of any sample. The integration time was set to 210 ms to maximise the signal without experiencing saturation. A dark spectrum was first acquired and non-linearity correction and electric dark were enabled. This is the standard procedure used unless stated otherwise. This spectrum serves to quantify the power in the pump beam and whilst this is not an absolute power value as discussed previously, it can be accurately compared with those obtained in the presence of a sample. As the signal to noise ratio is poor for a fluorescent sample at 210 ms for each sample a spectrum was acquired for an integration time of 210 ms and another at 400 ms. The 400 ms spectra have been used for all of the results in this section (Section 4.8). The purpose of the 210 ms spectra was to confirm that the software's internal calculation of the detected power per nm correctly accounts for a change in integration time and that the 210 ms pump spectrum can reliably be compared with the 400 ms sample spectra. The results of Table 4.2 were confirmed to be within 1 or 2 % of each other when calculated using the 210 or 400 ms spectra.

The samples assessed are those shown in the sample column of Table 4.2, which correspond to those depicted in Fig.4.24. Where a sample is listed as 'cent' or 'edge' this indicates that the pump was positioned either at the center of the sample cross-section or near the edge at the position where the most fluorescence was detected. Rod shaped samples were positioned in the same way as for the trapping fraction measurements - samples were positioned at a depth into the sphere port such that the end of the fluorescent rod was in line

with the inner edge of the anodised black ring that frames the 1" port reduced opening. For the disc shaped samples this was not possible due to the port opening size and available mounts. Disc shaped samples were placed in contact with the reduced port opening, such that the opening was completely covered.

The conversion from spectra of units 'power per nm' ($P(\lambda)$) to lumens ($\phi_v$) is made using a photopic luminosity curve ($\bar{y}$) and the luminous efficacy of radiation at 555 nm ($K_{cd}$ = 683 lm/W) as follows [105]:

$$\phi_v = K_{cd} \int P(\lambda)\,\bar{y}(\lambda)\mathrm{d}\lambda. \tag{4.19}$$

The unhindered pump spectrum was integrated numerically between 430 and 480 nm to determine the power (proportional, not absolute) in the pump beam. Each sample spectrum was then multiplied by the 'Judd-Vos modified CIE 2-deg photopic luminosity curve (1978)' [106]. The luminosity curve was first interpolated linearly to match the wavelength intervals of the acquired spectra (≈ 0.379 nm bins). The resulting weighted spectrum was integrated between 400 and 780 nm and multiplied by 683 to produce a lumen value. This lumen value was then divided by the calculated pump power to yield the lm/W$_{rad}$ values seen in Table 4.2. To clarify, that is lumens per watt of light incident on the sample as opposed to a wall plug efficiency.

Correlated colour temperature values were calculated for each sample with the aid of the IES TM-30-15 Advanced Calculation tool, available from [107]. The correlated colour temperature of a white light source is a useful metric for the how 'warm', or how blue or yellow the light is. Formally, it is defined as "the temperature of the Planckian radiator having the chromaticity nearest the chromaticity associated with the given spectral distribution on the diagram where the (CIE 1931 standard observer based) $u'$,2/3 $v'$ coordinates of the Planckian locus and the test stimulus are depicted ([108, 107])". CCT values are determined by the calculation tool using the method presented in [109], using a minimum search program and either a triangular or parabolic function between sample points on the Planckian locus, depending on the spectrum's distance from the locus. Spectra were first interpolated linearly to satisfy the 1 nm spacing for the tool input and 380 to 700 nm acceptable range. The resulting CCT values can be seen in Table 4.2.

To determine the percentage of pump/fluorescence photons detected the spectra were first converted from units of 'power per nm' to proportional to 'number of photons per nm'. This was done simply by multiplying the power measured at a given wavelength by the wavelength at which it was measured. Photon energy $E = hc/\lambda$, where $h$ is Planck's

constant, $c$ is the constant speed of light in vacuum and $\lambda$ is the photon wavelength. The pump photons were then counted by numerically integrating the spectra between 430 and 480 nm. Comparison with that obtained for the unhindered pump beam spectrum yields the column 'Pump photons detected' in Table 4.2. The fluorescence photons are counted in a similar way but with the integration limits of 480 to 700 nm. In addition, the normalisation is to the total number of photons in the unhindered pump beam minus the number of pump photons detected. For the polycrystalline sample the upper limit was increased to 750 nm to include the increased red emission present, see Fig.4.31. This was not done for all samples as above 700 nm the noise in the spectra become significant. The integrated value is compared with the number of photons in the unhindered pump to yield the 'fluorescence photons detected' column.

### 4.8.2 Results and Discussion

The results of the efficiency and CCT calculations are collated in Table 4.2. An illustrative selection of the sample spectra used can be see in the Figs.4.31, 4.32 and 4.33. These compare the CylRodHemi 60 mm prototype to the commercially available samples, the various efficiencies achievable with the circular cross-section rod and the square cross-section rod to the circular cross-section rod, respectively.

The sample which represents our closest prototype to the proposed fibre based phosphor component is the Cylindrical rod which has been shortened to 60 mm to achieve white output, with hemispherical lens for enhanced outcoupling. We consider the spectrum of this sample when pumped close to the sample edge 'CylRodHemi 60mm Edge' when comparing the prototype performance with some of the functional requirements in Section 4.9.1. The CCT of this sample is measured as $\approx$ 5500 K, with CIE 1931 $x$ $y$ coordinates of (0.3345,0.3998). These coordinates, along with a representative selection of the remaining samples can be seen plotted in a CIE 1931 colour space chromaticity diagram later in Fig.4.34. As a sample with CCT of $\approx$ 5500 K, its calculated efficiencies such as lm/W$_{rad}$ and % fluorescence photons are a fair representation of that expected of a phosphor component used in a typical white light.

Of the original pump photons that have undergone fluorescence (or have been lost), the measured percentage that have exited the 'CylRodHemi 60mm' sample via its hemispherical lens and entered the integrating sphere is found to be 43%. This corresponds to a yield of 148 lumens of white output per watt of incident pump light. We now consider how these obtained efficiencies compare theoretical predictions.

| Sample | Pump photons detected (%) | Fluorescence photons detected (% of pump photons) | CCT (K) | lm/W$_{rad}$ |
|---|---|---|---|---|
| (Chromalit) PolyCryst disc | 3.5 | 36 | 3563 | 84 |
| (Crytur) MonoCryst disc | 18.2 | 45 | 12133 | 136 |
| SquareRodHemi 40mm | 14.0 | 39 | 9760 | 128 |
| CylRodHemi 60mm Cent | 5.8 | 30 | 5768 | 108 |
| CylRodHemi 60mm Edge | 7.3 | 43 | 5449 | 148 |
| CylRodBare 100mm | 1.4 | 11 | 5047 | 43 |
| SquareRodBare 100mm | 1.7 | 11 | 5202 | 41 |
| SquareRodHemi 100mm | 1.2 | 39 | 4501 | 143 |
| CylRodHemi 100mm Cent | 0.9 | 23 | 4561 | 86 |
| CylRodHemi 100mm Edge | 1.4 | 32 | 4502 | 119 |

**Table 4.2:** Key performance metrics for the samples under assessment, see Fig.4.24. Measured in the configuration of Fig.4.26 (with Spectrometer). CCT = correlated colour temperature. lm/W$_{rad}$ = lumens per watt of light incident on the phosphor. Calculation details for all columns are given in the Section 4.8.1. Note that for samples with comparable input and output faces (this includes the discs), the detected light is approximately half of the total emitted. The double horizontal line half way down separates the approximately white spectra from those that are almost entirely fluorescence and lie too far from the black body locus for their CCT and lm/W$_{rad}$ values to be interpreted in the same way as those above.

We have calculated the maximum percentage of fluorescence photons we could expect to see, out of both ends of the rod, relative to the total number that do not pass through unconverted, $F_{th}$ using the following known factors:

- Fresnel reflection at the input face, assuming normal incidence, **8.6**%
- Internal quantum efficiency of phosphor, **94**% [19]
- Trapping fraction (~max), **80**%, see Fig.4.3
- Fresnel reflection between rod end, adhesive and lens, estimated of order 1 or 2 %, neglected from calculation.
- Outcoupling efficiency of hemispherical lens, $d_{rod}$ = 5, $d_{lens}$ = 10 mm, ≈ **90**%, see Fig.4.12.

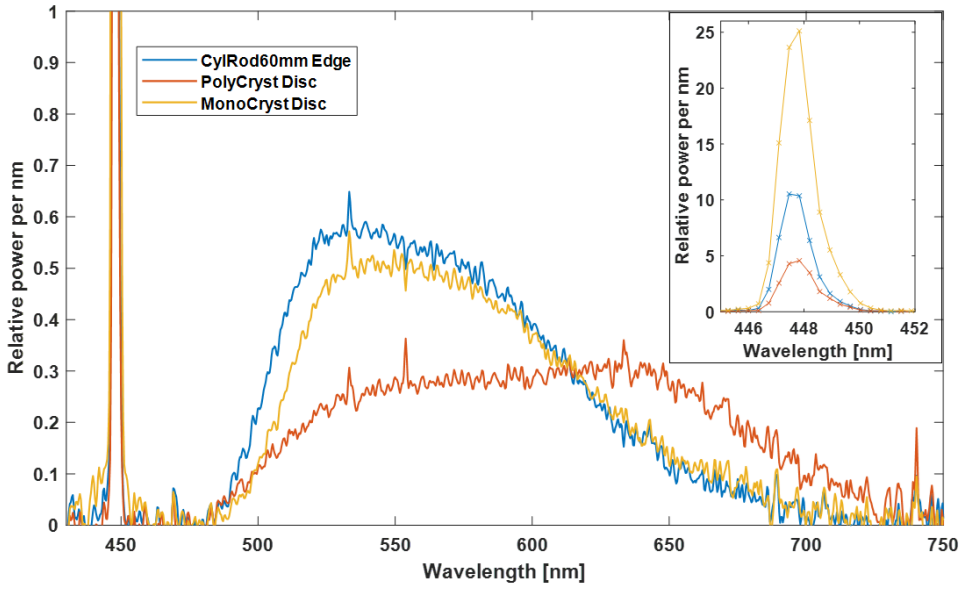


**Figure 4.31:** Relative spectral power distributions for the cylindrical prototype white light 'CylRod60mm', pumped with a 450 nm beam near its edge and two commercially available phosphor discs. Spectra have been smoothed with a Gaussian-weighted moving average filter of width 5 data points (≃ 1.9 nm). The inset image top right shows the pump peaks at larger y scale. The inset spectra are unsmoothed.

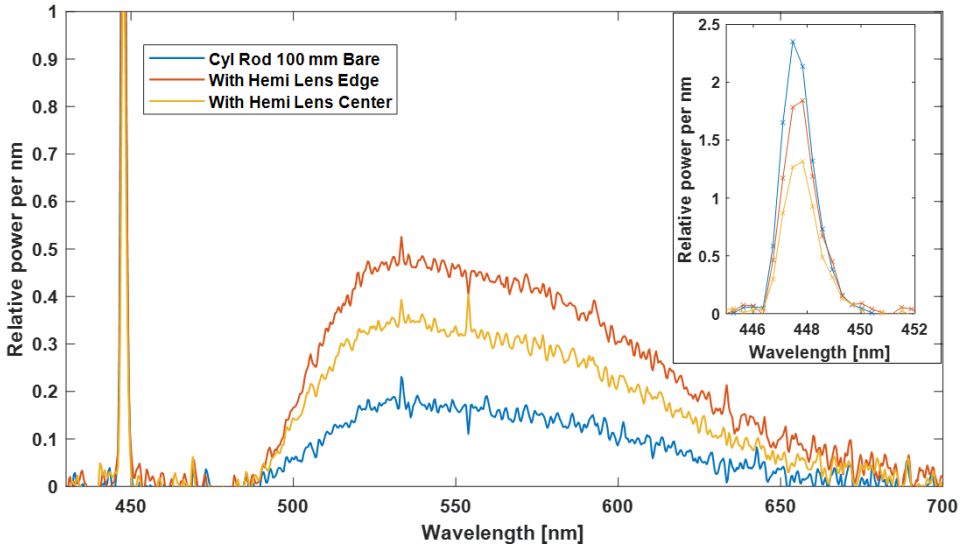


**Figure 4.32:** As in Fig.4.31 but comparing the unmodified cylindrical YAG:$Ce^{3+}$ rod with the pump position dependent spectra obtained with the addition of a hemispherical outcoupling lens.

$$F_{th} = (1 - 0.086) \times (0.94) \times (0.80) \times (0.90) = 62\% \tag{4.20}$$

The measured amount of fluorescence (43%) is $\simeq \frac{2}{3}$ of that we would expect to measure if the rod - lens system was completely free from defects and if no light escaped through the input face of the rod (62%). Note that the fraction of light exiting the unmodified rod face will be considerably less than half of the total fluorescence generated, as no outcoupling lens is attached here. Based on these values it seems reasonable to expect that a next generation prototype with the following modifications:

- Anti-reflective coatings on the input and exit (hemispherical lens) faces
- A Distributed Bragg Reflector or other dichroic element on the input face, to reflect fluorescence to the output end whilst allowing the pump light to enter the rod.
- Improved finishing of the cut rod surface and more precise lens placement.
- Pump beam width reduction - was optimised for 100 mm sample.

would achieve an efficiency of around 200 lm/W$_{rad}$. If it is found that the surface quality of the cylindrical rods is consequentially poorer than that of the square rods, there may be potential to improve this and improve the efficiency further. One is then left with managing the directionality of emission, which is likely to come at some cost to overall efficiency but as we have seen in Section 4.3.4 the directional efficiency can be kept high if one is willing to increase the cross-sectional area of the device (law of etendue [110]).

### 4.8.3 Square vs Circular Cross-section

Simulations have shown (Fig.4.3) that the trapping fraction achieved in a cylindrical rod may exceed the uniform 68% trapping fraction predicted for a square cross-section, if either the dopants are, or pump light is, distributed more towards the edges of its circular cross-section. This is the case for the simulated system, without any surface losses or defects present. If the rod is illuminated across its entire cross-section, the resulting trapping fraction is $\simeq$ 70%, roughly equivalent to that of the square cross-section rod. The monocrystalline circular cross-section rods from Crytur in conjunction with the fluorescent acrylic rods have successfully demonstrated the predicted trapping fraction variation and provided a benchmark for the magnitude of this effect for a 'real-world' lossy sample.

It is noted that the unmodified square cross-section rods perform equivalently to the unmodified circular rods in Table 4.2 and Fig.4.33. However, once a hemispherical lens is added the square rods outperform even the edge pumped circular ones. This may be due

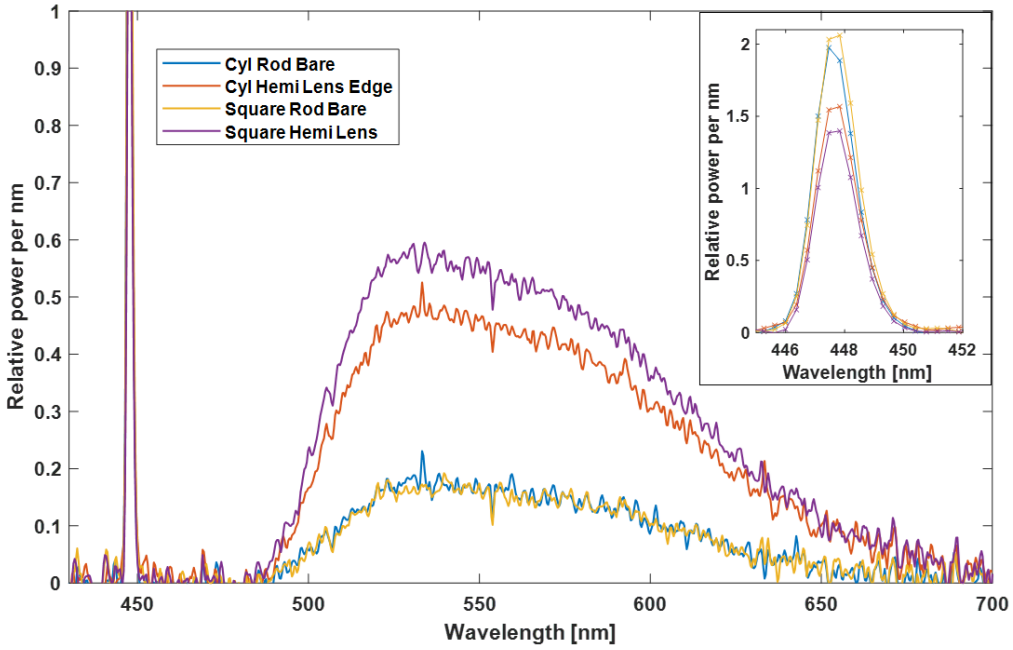


**Figure 4.33:** As in Fig.4.31 but comparing square cross-section samples with their cylindrical counterparts.

to its decreased sensitivity to lens misalignment as discussed in Section 4.7.

The question of which shape is 'better' will likely depend heavily on manufacturing feasibility and specific application. A circular cross-section has been assumed as the intended geometry as drawing the samples as fluorescent optical fibres seems an appealing manufacturing solution.

## 4.9 Conclusions

### 4.9.1 Evaluation Against Functional Requirements

With the optical properties of the proposed fibre-based solution now well understood through simulation and experiment, we begin our conclusions by looking again at the requirements outlined in Section 1.4.1.

1. **Emission Cone**

   The strength of the proposed fibre-based solution is to enable sufficient pump absorption and power output in a geometry that need not increase the area of the ele-

ment, instead distributing the material along the fibre length. A small cross-section lends itself to effective collimation with modification of the fibre ends or external optics. The prototype is presented without a modification that adds a high degree of directionality to its emission. It is expected that the optimum component / modification to provide directionality will be derived from existing optical components (e.g. the compound parabolic concentrator [111]) and will be dictated by the target application of the device it is deployed in and manufacturing constraints.
A novel modification, the modified ball lens, is explored in Section 4.3.4. The simulations there are restricted to two dimensions and the results cannot be quantitatively applied to a 3D system. The equivalent directionality seen in Fig.4.16 is anticipated to require a larger lens diameter than shown for a 2D simulation when extended to 3D, in line with the behaviour observed for the hemispherical lens (Figs.4.11 and 4.12). If such modified lenses were of further interest we would suggest the use of a dedicated ray-tracing and optical design software.

2. **Imaging Capability**

   The fibre-based solution presented was designed with imaging capability at its heart. In the same way an imaging fibre optic bundle (fibrescope) works it is proposed that a bundle of fluorescent fibre can be used to transfer spatial patterning in the pump beam into spatial patterning of the fluorescence output. The maximum resolution of the resulting fluorescence pattern is provided by the number of independent fibres, which effectively become pixels in the array.

3. **Optical Efficiency**

   The working prototpye 'ClyRodHemi 60 mm', when pumped at the edge of its cross-section, has been measured to produce ~ 150 lumens of white light from its output face per watt of incident pump. It is expected that an improved prototype would reach or exceed the 200 lm/W target, as outlined in Section 4.8.2. Replacing the out-coupling lens used in this prototype with a component achieving directionality will then come at some efficiency cost, but this can be kept minimal if one is willing to increase the cross-sectional area of the fibre + lens system sufficiently.

4. **Brightness / Power Handling**

   For the same reasons that the fibre geometry (or fibre bundle if seeking to preserve

patterning of the pump light) should be favourable for achieving a narrow emission cone it is expected to be well suited to high power applications. Thermal management aside, we have seen in Section 3.2 that the volume of phosphor needed to reach a high lumen target is very small, provided the activator ions are constantly being excited. The absorption properties of YAG:$Ce^{3+}$ were considered in Section 3.3. Achieving the necessary pump absorption to yield white output in a fibre geometry shouldn't present a problem as optical path length is available in abundance. At high pump intensities the ground state of the activator ions will begin to become depleted. That is, the number of cerium ions not already in the excited state will become small(er). This will affect the fractional pump absorption at high pump intensities. A recent article investigates this so called luminescence saturation (which is distinguished from thermally induced saturation) and confirms a "remarkably high saturation threshold" [86]. The solution to any luminescence saturation would seem to be as simple as extending the length of the fibre, with no need to increase its cross-sectional area.

5. **Colour**

The 60 mm working prototype sample has an output correlated colour temperature of roughly 5500 K. This was the CCT we were aiming for for a representative YAG:$Ce^{3+}$ only white light. The 4 circular marks in Fig.4.34 represent the 4 samples made from the exact same phosphor (from Crytur). A straight line connecting and extending beyond these points represents the accessible chromaticity coordinates for this material (not factoring in potential redshift from re-absorption for very long samples). These chromaticity coordinates (and hence varying CCT values) may be accessed by adjusting the ratio of unconverted pump light to fluorescence exiting the output end of the sample. This can be done by changing the length of the sample, as was demonstrated here, or could also be achieved by varying the pump angle or angular distribution. Both alterations vary the mean optical path the pump light travels. Colour tuning could also be achieved through intentional saturation of the emitters, but that seems unlikely to be desirable.

The colour rendering index (CRI) for the 60 mm sample was calculated to be 62 by the IES calculation tool. The expected CRI for this YAG:$Ce^{3+}$ only phosphor was around 70 [16]. This shortfall is likely due to the reduced spectral coverage of the LD pump vs the typical LED pump spectrum. To achieve the high CRIs required for quality interior lighting (around 90) it is necessary to expand the phosphor composition be-

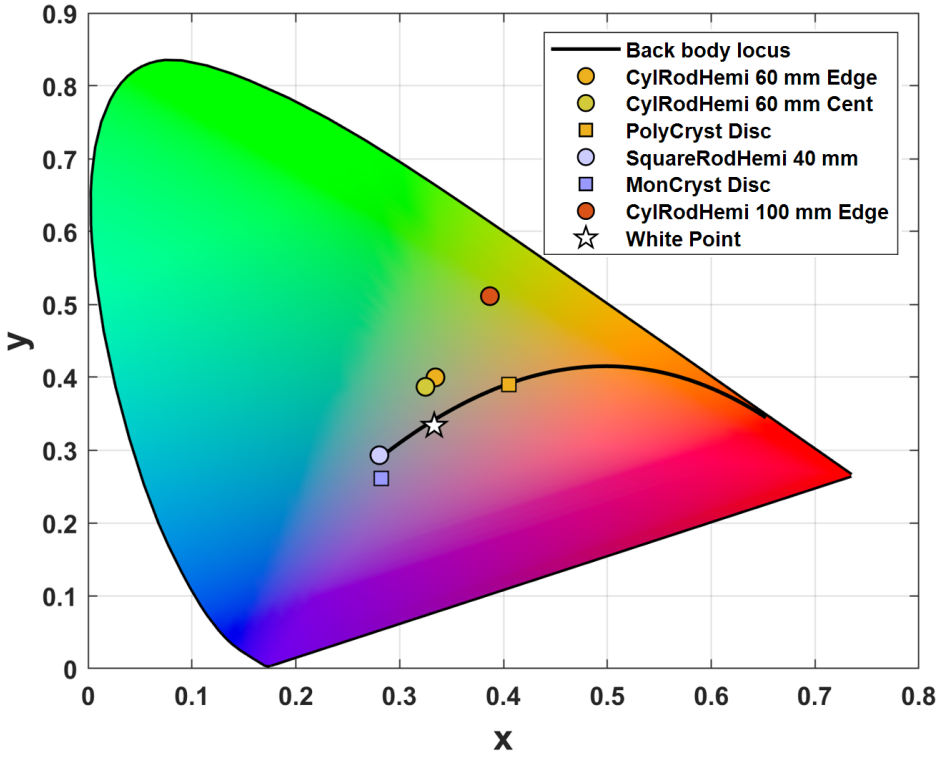


**Figure 4.34:** CIE 1931 color space chromaticity diagram with the chromaticity coordinates of a selection of the samples in Table 4.2 marked. The black body locus and theoretical white point $(\frac{1}{3}x,\ \frac{1}{3}y)$ are marked for convenience. The black body locus was constructed from a third order polynomial fit of the sample points provided in the IES TM-30-15 Advanced Calculation tool

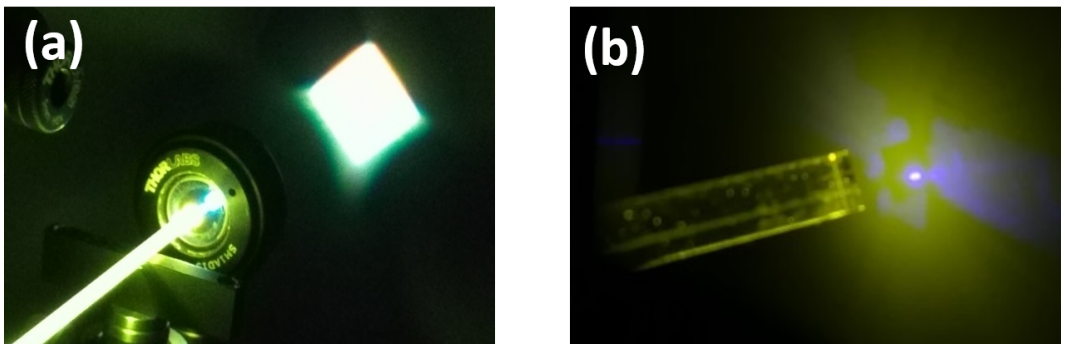


**Figure 4.35:** (a) Lensed emission from a cuboidal YAG:$Ce^{3+}$ guide illuminated by a fibre-coupled LED source. Excluding chromatic aberrations at the image edges, the imaged fibre end appears to the eye as a uniform white square. (b) emission from a cuboidal YAG:$Ce^{3+}$ guide illuminated by a thin beam from a blue diode laser. The unconverted pump light is clearly visible as a discrete beam in the output light.

yond purely cerium doped YAG (or use an alternative high refractive index, scatter-free phosphor). This may be possible whilst maintaining a single host material (e.g. YAG) with multiple dopants or could require the combination of multiple distinct phosphors as individual fibres or 'strands' in contact forming a larger fibre if index matched. [112] shows a Ce,Sm based fibre with high CRI of 83. There is also a possibility that fluorescence leaking from YAG:$Ce^{3+}$ fibres in an imaging bundle could be recycled by a red phosphor between the fibres or in their cladding if present, but controlling the colour mixing and angular emission of this extra light would likely present a significant challenge.

Another key consideration is the colour uniformity of the output light. No spectral dependence on the emission angle should be present for the fluorescence and none is observed (by eye) for the prototypes. The angular distribution of the pump light however is determined by the distribution it entered with (unless considerably altered by the outcoupling / directing optic). There is also the distribution of the pump spatially across the fibre cross-section, this should be suitably mixed if the pump light has the opportunity to diverge into the sides of the fibre and reflect off it a few times. Images of emission from a divergent (LED) pump sample and a collimated pump (diode laser) can be seen in Fig.4.35. If localising of the pump beam to high radial positions is exploited to access higher trapping fractions this would likely need to be balanced with colour uniformity of the output light.

6. **Thermal Conductivity**

   Monocrystalline YAG:$Ce^{3+}$ has a high thermal conductivity material of up to 13 $Wm^{-1}K^{-1}$ [87] and experiences minimal efficiency loss at temperatures below 300 °C [18]. The scatter-free fibre geometry offers freedom to adjust length and concentration to control the volume and surface area over which heat is generated. It is also possible that the concentration of dopants along the fibre could be varied to homogenise the heat generation over the length of the fibre, if desired.

### 4.9.2 In the Context of Other Work

Monocrystalline YAG:$Ce^{3+}$ fibres have been investigated and employed as fast scintillation detectors for detection of ionizing radiation [88, 113, 114]. A bundle of optically separated fibres coupled to a pixelated detector provides information on the location and energy of a scintillation event, enabling tracking of sources of ionising radiation [88]. In these scintilla-

tion applications the concentration of cerium dopants is typically much higher than would be suited to this application and there is no need to couple emission to free space, so the dominant issues can differ to those considered here.

YAG crystal fibres on the low micrometer scale have also seen some interest as sources of white light for applications such as optical coherence tomography [112, 115]. The most relevant work to that presented here is the use of high concentration monocrystalline YAG:$Ce^{3+}$ rods as luminescent concentrators [91, 116, 117]. In these works the YAG cuboidal guide concentrates the emission from a series of side-pumping LED sources to achieve bright fluorescence emission from the end facet. The problem of outcoupling and directing emission is addressed with the use of a compound parabolic collimator [85].

We have focused our attention on the trapping fraction of fluorescence generated within cylinders of high refractive index. The dependence of the trapping efficiency on the radial excitation position has been investigated and demonstrated experimentally, the full analytical treatment for which was published as recently as 2015 [90]. The existence of this additional trapping power in an imperfect, lossy waveguide is of significance to the overall efficiency of light sources utilising fluorescent fibres. We are not aware of any similar investigation or experimental demonstration.

### 4.9.3 Summary

The conclusions for this chapter are mostly contained within the 'Evaluation Against Functional Requirements' Section 4.9.1 and 'In the Context of Other Work' Section 4.9.2, nevertheless we will add some final remarks here.

It is clear form the preceding sections that fluorescent optical fibres show promise as a solution to the directional lighting problem central to this work. A full assessment of their suitability now hinges on factors such as manufacturing feasibility, cost, durability and their ability to meet the more specific requirements of a particular illumination system or product. We therefore conclude this investigation having found no evidence of shortcomings against any of the defined functional requirements, and recommend fluorescent fibres for further investigation in this context.

# 5 Fibre Simulation Details

The following models are all related to the work on fluorescent fibres of Section 4. As the proposed size of the fibres is much larger than the visible wavelengths in question, the following simulations all use COMSOL multiphysics' Geometrical Optics Module (Ray Tracing Study) to assess the propagation of emission, unless stated otherwise.

## 5.1 A 3D Fibre Simulation

### Geometry

The geometry used to obtain the results shown in Fig.4.3 is shown in Fig.5.1. A YAG rod ($n = 1.83$) is suspended in air ($n = 1$), with a spherical point emitter at its center. The position of the emitter is varied and the power from rays reaching different surfaces is used to determine the trapping fraction. The aspect ratio (diameter:length) of the fibre was chosen by varying the fibre length and looking for convergence to a known value of trapping fraction. This can be seen in Fig.5.2. The geometry is specified below for the radial position sweep whilst the length is fixed.

- Air cylinder - radius = 6.9 mm, length = 22 mm, Layers 1 mm on top and bottom, axis in $\hat{x}$, position (0,0,0).
- YAG cylinder - radius = 6.6 mm, length = 20 mm, axis in $\hat{x}$, position (0,0,0).
- Point (radiation source) - Variable position = (10 mm, 0 mm, DispR×6.6 mm)

### Boundary Settings

Two 'walls' are defined. One designed to catch rays that have 'leaked' from the curved sides of the structure and the other to catch everything else, namely emission from the ends of the fibre. The former is defined as the outermost curved boundary of the simulation, up until the air end caps, which are not included. With reference to Fig.5.1, this is the curved gold surface but at a slightly increased radius, such that it encloses the air layer. The latter is those simulation outer surfaces that remain. Both 'Walls' have the deposited ray power sub-node and the freeze boundary condition added to them.

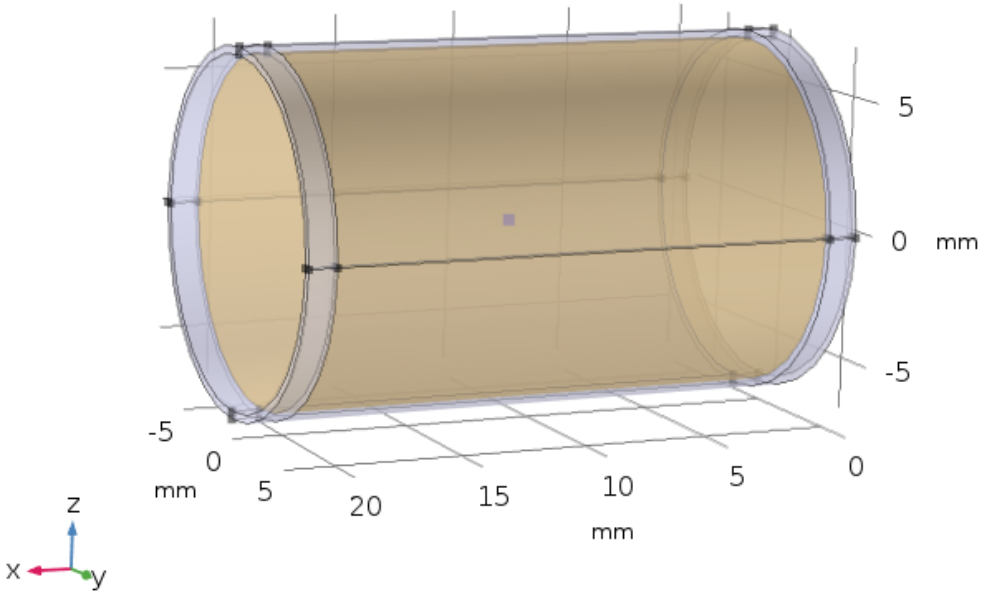


**Figure 5.1:** Schematic of the 3D COMSOL simulation geometry used to determine the dependence of the trapping fraction on radial source displacement (Fig.4.3). Here the length of the YAG (n=1.83) rod (gold colour) is 20 mm and its radius is 6.6 mm. The translucent blue outer-cylinder marks the simulation domain and is filled with air ($n = 1$). The point marked in dark blue at the center of the fibre is used as a point spherical source of rays and has its position swept radially in the $\hat{z}$ direction.

### The Trapping Fraction

Here the trapping fraction is calculated as: (total power − leaked power)/(total power), where the leaked power is that which reaches the curved boundary enclosing the fibre. The total power is 1 W. The leaked power is determined by evaluating the variable 'gop.wall1.bsrc1.Qp_int' at the last time computed [wall1 / wall2 will depend on the order in which it was defined].

### Ray Inlet

Release from Point -> Spherical -> Deterministic -> 3000 rays in wave vector space -> 1 W total source power.

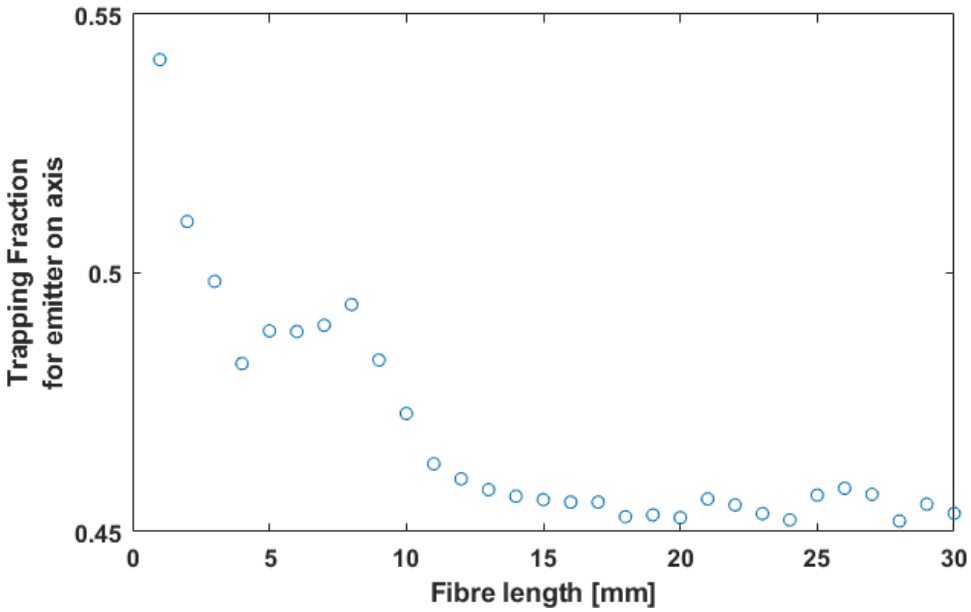


**Figure 5.2:** The trapping fraction for an emitter positioned on the fibre axis (fractional radial displacement = 0), as a function of the fibre length (fibre diameter 13.2 mm). On axis emitters are expected to have a trapping fraction of 45%, as in eq.4.3. A fibre length of 20 mm is selected based on the apparent convergence (within expected precision of simulation), to maintain result fidelity whilst minimising the required computational resources.

### Secondary Rays

Set to 3000. Required to include rays that are partially reflected upon escaping the fibre, which can partially escape again on subsequent reflections.

### Ray Tracing Settings

Time steps specified as 0 to 400 ps in 5 ps increments, with simulation terminating when no active rays remain.

### Parametric Sweep

The radial displacement of the point source was swept in the range 0 to 0.99 x the rod radius (6.6 mm), in 0.01 x the rod radius increments.

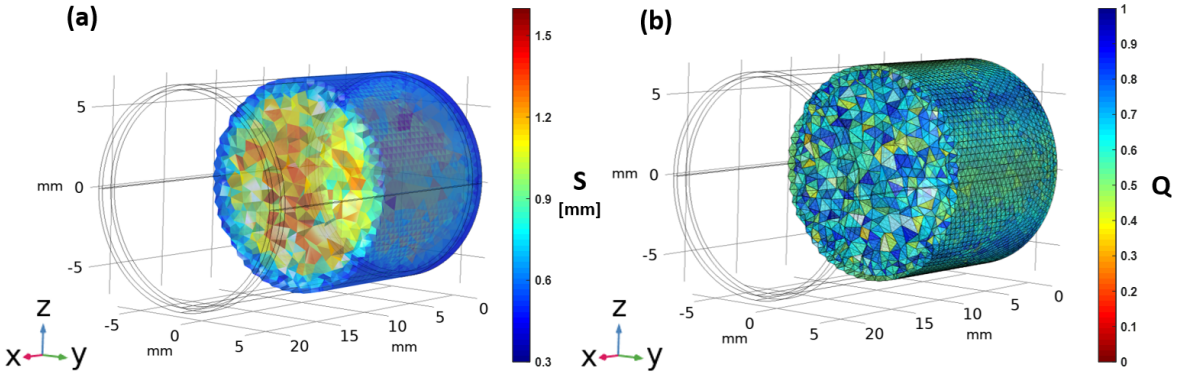


**Figure 5.3:** A visualisation of the generated physics controlled mesh cut in the $zy$ plane at the emitter position. The mesh volume elements plotted are semi-transparent and colour-coded according to their size $S$ in (a) and quality $Q$ in (b) as defined by the skewness criteria of eq.2.14.

### Mesh

Physics controlled mesh, fine, as shown in Fig.5.3.

## 5.2 3D Square Cross-section Guide Simulation

The simulation settings for this 3D square cross-section guide simulation have been adapted directly from the 3D cylindrical simulation in the previous section. Fig.5.4 shows the modified geometry, where the YAG cylinder of radius 6.6 mm has been replaced by a cuboid of width 13.2 mm, likewise for the outer cuboid. All other simulation setting remain as in Section 5.1, with the exception of the parametric sweep which was reduced in resolution to fractional Z displacements of 0 to 0.95 in 0.05 increments.

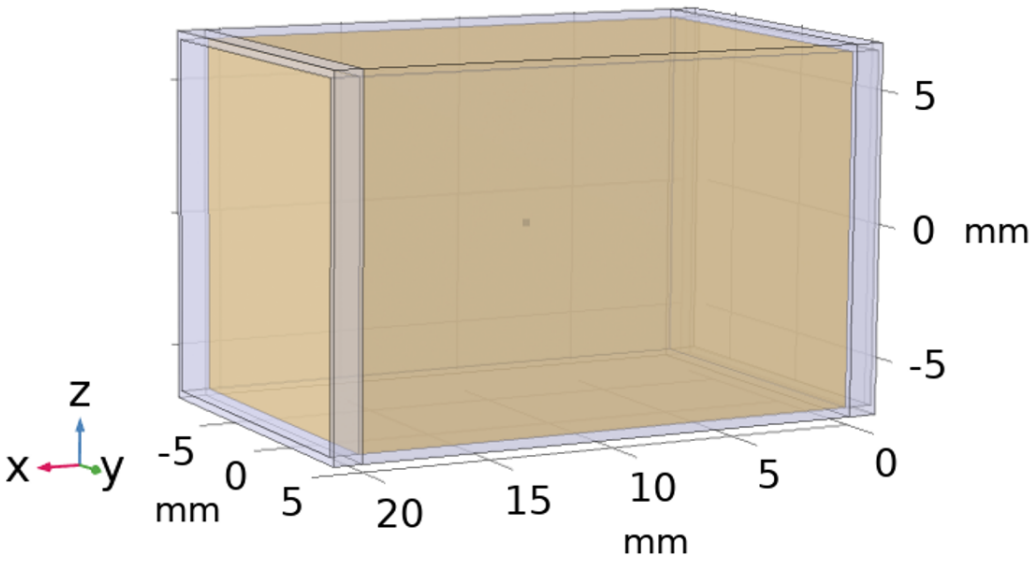


**Figure 5.4:** Schematic of the 3D COMSOL simulation geometry used to determine the lack of dependence of the trapping fraction on source displacement for a square cross-section guide (Fig.4.7). Here the length of the YAG (n=1.83) cuboid (gold colour) is 20 mm and its width is 13.2 mm. The translucent blue outer-cuboid marks the simulation domain and is filled with air ($n = 1$). The black point at the center of the guide is used as a point spherical source of rays and has its position swept in the $\hat{z}$ direction.

## 5.3 Increasing Path Lengths Simulation

### Geometry

The geometry used to obtain the results shown in Fig.4.17, which illustrate the increase in mean path length and number of reflections of rays generated off axis, unsurprisingly is based on the 3D fibre simulation in Section 5.1. The key differences are the increased length, to broaden the range of possible path lengths, and the removal of the air 'caps' on the fibre ends (as in the lossy fibre simulation of Section 5.4). The simulation geometry is shown in Fig.5.5.

- Air cylinder - radius = 6.9 mm, length = 66 mm, axis in $\hat{x}$, position (0,0,0).
- YAG cylinder - radius = 6.0 mm, length = 66 mm, axis in $\hat{x}$, position (0,0,0).
- Point (radiation source) - Variable position = (33 mm, 0 mm, DispR×6.0 mm)

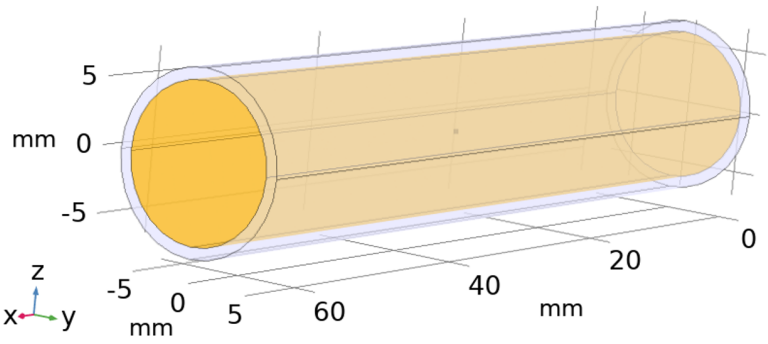


**Figure 5.5:** Schematic of the 3D COMSOL simulation geometry used to determine the increase in mean path length and number of reflections for rays emitted off axis in a cylindrical YAG guide (Fig.4.17). Here the length of the YAG (n=1.83) cuboid (yellow/gold colour) is 66 mm and its radius is 6 mm. The translucent blue outer-cylinder marks the simulation domain and is filled with air ($n = 1$). The black point at the center of the guide is used as a point spherical source of rays and has its position swept in the $\hat{z}$ direction.

### Boundary Settings

Two 'walls' are defined. The first is applied to the curved outer-cylinder of the simulation and has the disappear boundary condition applied. This removes any rays that are not trapped in the fibre, i.e. the leaked rays, from contributing to the final ray propagation statistics. The remaining surfaces, the two planar ends of the cylinder, are assigned the freeze boundary condition. By omitting the YAG - Air boundary at the end of the fibre we are assuming perfect outcoupling and preventing the indefinite reflection of some of the trapped rays form the fibre ends.

### Ray Inlet

Release from Point -> Spherical -> Deterministic -> 3000 rays in wave vector space -> 1 W total source power.

### Secondary Rays

Set to 600. In principle none should be required.

### Paths Lengths and Number of Reflections

Firstly, in the primary geometrical optics settings node, Compute optical path length and count reflections must both be checked / enabled. The variables $\bar{N}_{refl}$ and $\bar{L}$ are determined by global evaluation of the variables 'gop.gopaveop1(gop.Nrefl)' and 'gop.gopaveop1(gop.L)'.

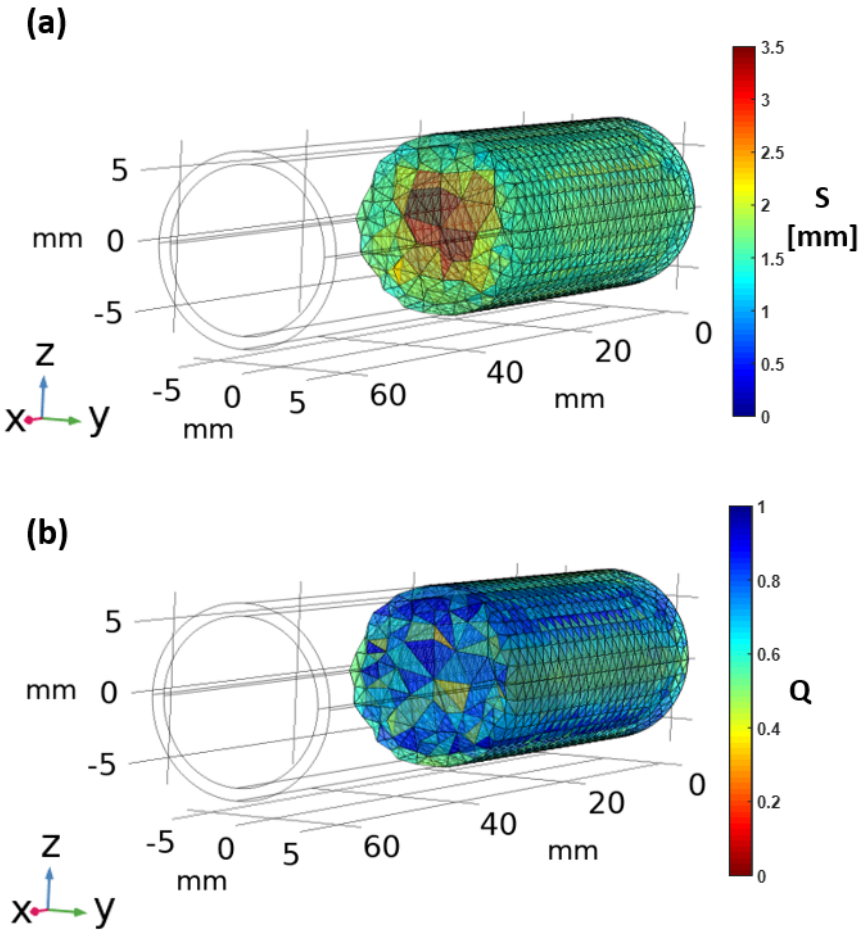


**Figure 5.6:** visualisation of the generated Fine physics controlled mesh cut in the $zy$ plane at the emitter position. The mesh volume elements plotted are semi-transparent and colour-coded according to their size $S$ in (a) and quality $Q$ in (b) as defined by the skewness criteria of eq.2.14.

### Ray Tracing Settings

Time steps specified as 0 to 30,000 ps in 1000 ps increments, with simulation terminating if no active rays remain. At 30,000 ps a small number of rays are still propagating at large angles from the fibre-axis. These are happily omitted from the propagation statistics as they would be quickly attenuated in the presence of any losses.

### Parametric Sweep

The radial displacement of the point source was swept in the range 0 to 0.99 x the rod radius (6 mm), in 0.01 × the rod radius increments.

### Mesh

Physics controlled mesh, fine, as shown in Fig.5.6.

## 5.4 A Lossy 3D Fibre Simulation

The 'bones' of the 3D fibre simulation that models the effect of a lossy fibre surface are of course similar to those of section (Section 5.1). The key differences are the inclusion of a non-physical annular layer within the YAG rod, just below the surface, via which the losses are added and the removal of the air layers on either end. Due to the present similarities simulation details are only provided where they are consequentially different from in Section 5.1.

### Geometry

- Air cylinder - length = 22 mm, radius = 6.9 mm. Cylinder axis = $\hat{x}$, position (0,0,0).
- YAG cylinder - length = 22 mm, radius = 6.0 mm. Layers on side = 0.15 mm Cylinder axis = $\hat{x}$, position (0,0,0).
- Point (radiation source) - Variable position = (10 mm, 0 mm, DispR×6.0 mm).

### Boundary Settings

Due to the presence of losses in this system, the trapping fraction cannot be calculated by subtracting the power that leaks from the sides of the cylinder from the known total

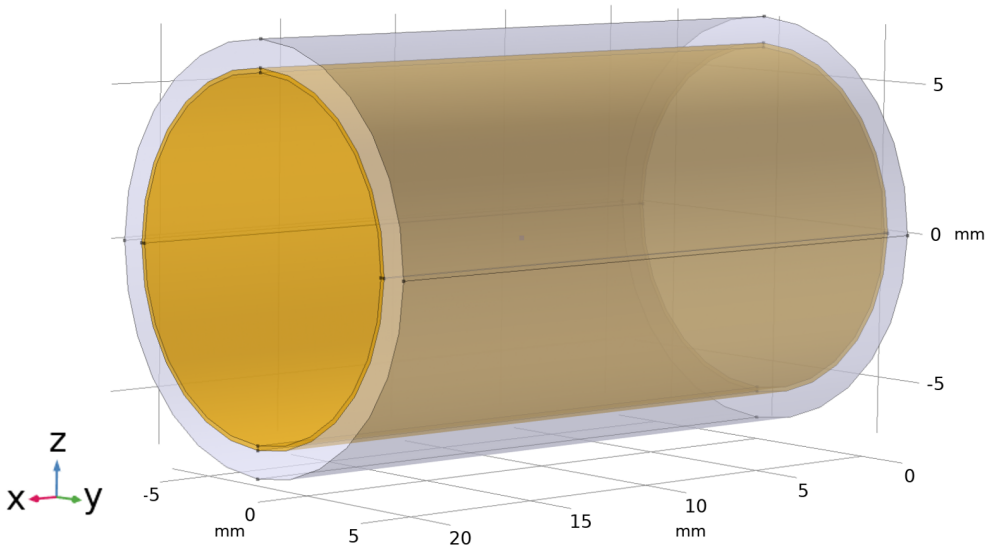


**Figure 5.7:** Schematic of the 3D COMSOL simulation geometry used determine the effect of a lossy cylinder surface on the trapping fraction (Fig.4.18). Here the length of the YAG (n=1.83) rod (gold colour) is 20 mm and its radius is 6.0 mm. The translucent blue outer-cylinder marks the simulation domain and is filled with air ($n = 1$). A layer within the YAG cylinder at 0.975× its radius serves as the lossy 'membrane'. Note that in contrast to Fig.5.1 the YAG cylinder continues to the ends of the simulation domain, without air 'end-caps'. The point marked in dark blue at the center of the fibre is used as a point spherical source of rays and has its position swept radially in the $\hat{z}$ direction.

power. Instead one wall (wall2) is defined as the entire circular end facet of the YAG rod, on both ends (taking care to include the thin ring at full 6 mm radius. This tracks all the power that propagates to the ends of the rods, without having to overcome a YAG-air index barrier as in Section 5.1 and quantifies the amount of light 'trapped' within the rod. Wall 1 is then all remaining outer surfaces of the simulation, tracking all rays that escape to the air layer. Both Walls have the deposited ray power sub-node and the freeze boundary condition added to them.

The final lossy wall (Wall3) was assigned to the inner YAG cylinder at radius 0.975×6 mm. Ideally this wall would coincide with the physical YAG-air boundary of the YAG cylinder, but attempts to do so caused the simulation to fail to converge, so this phantom layer just beneath the surface was added. The wall is set to have the disappear, or otherwise pass through condition with variable probability $P$, as shown in the results of Fig.4.18. Note that each ray that reflects off the rod surface has passed through this lossy barrier twice.

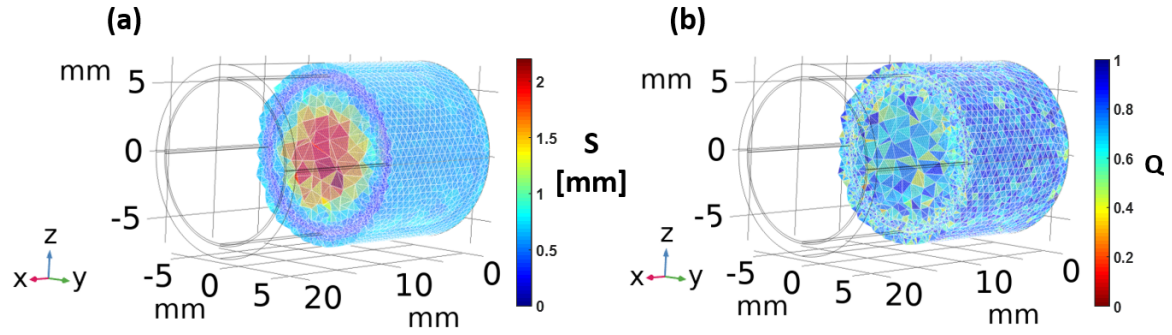


**Figure 5.8:** A visualisation of the generated physics controlled mesh cut in the $zy$ plane at the emitter position. Here DispR = 0.5. The mesh volume elements plotted are semi-transparent and colour-coded according to their size $S$ in (a) and quality $Q$ in (b) as defined by the skewness criteria of eq.2.14.

### The Trapping Fraction

The trapping fraction is determined by evaluating the power reaching the ends of the rod, via the variable 'gop.wall2.bsrc1.Qp_int' at the final time simulated, relative to the 1 W total system power.

### Secondary Rays

None required, left at 600 to identify erroneous behaviour where present.

### Mesh

Physics controlled - fine.

### Ray Tracing Settings

A significant increase in simulation run time is necessary with respect to Section 5.1. Unlike in the lossless case one must wait for trapped rays to propagate down the fibre. Time steps were 0 to 1500 ps in 5 ps increments, with the simulation terminating if no active rays remaining. At 1500 ps, for the more radially displaced sources active rays were still remaining, albeit a small fraction.

### Parametric Sweep

For each value of the loss probability $P$ (0,0.01,0.03,0.04,0.05,0.1,0.2,1), the radial displacement of the point source was swept in the range 0 to 0.95 x the rod radius (6 mm), in 0.05 × the rod radius increments. The range is restricted to less than 0.975 x the rod radius due to the offset of the lossy layer from the rod surface.

## 5.5 2D Fibre Cross-section Simulation

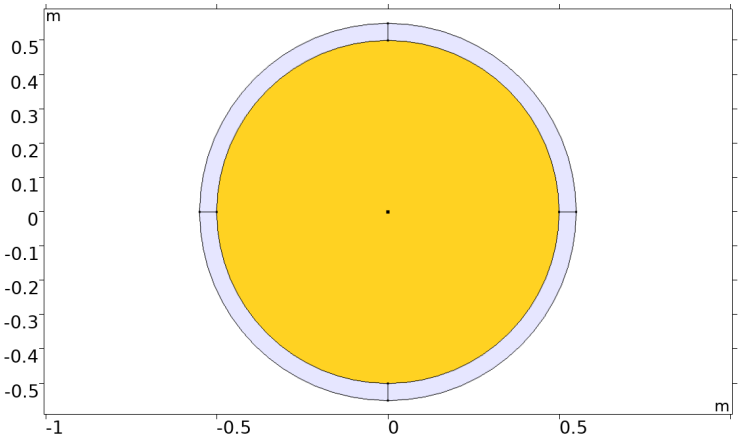


**Figure 5.9:** Schematic of the 2D COMSOL simulation geometry used to obtain the results of Fig.4.5. The gold shaded region represents the YAG (n=1.83) fibre, surrounded by air. Conventional horizontal $\hat{x}$ and vertical $\hat{y}$ axes apply

### Geometry

The simulation geometry is shown in Fig.5.9. It consists of two concentric circles, one YAG of index 1.83 and radius 0.5 m and another air of index 1 and radius 0.55 m (produced using the layer feature). Conventional horizontal $\hat{x}$ and vertical $\hat{y}$ axes apply. The point at 0,0 is swept to 0.99×0.5 m in $\hat{x}$ to radially displace the point of emission.

### Ray Inlet

Rays are released from the central black point in Fig.5.9, which has its position varied as discussed in the geometry section. The rays are released with a 'spherical' distribution with 500 rays in wave-vector space, 1 W power. The distribution sampling is deterministic.

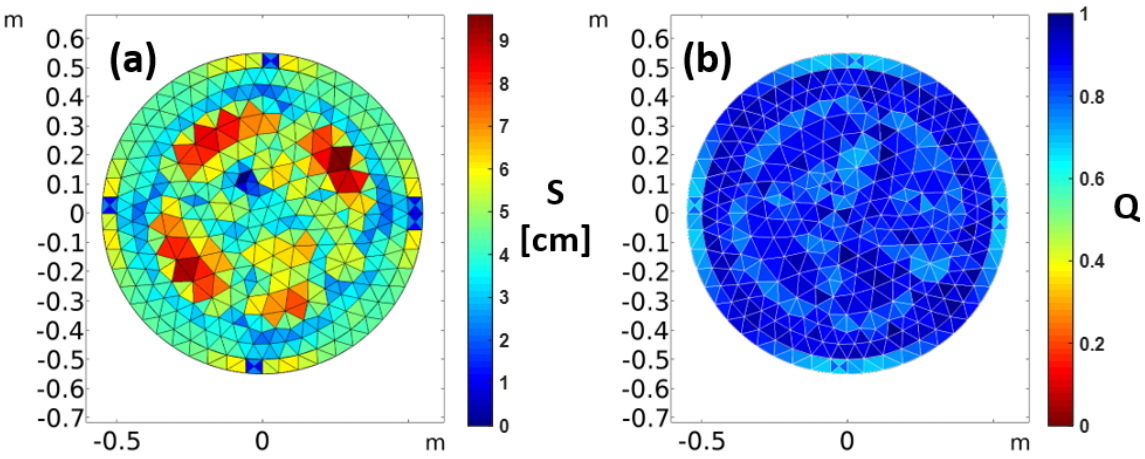


**Figure 5.10:** Normal physics controlled mesh when the emitter is displaced 0.25 m in $\hat{x}$. Mesh elements are colour-coded according to their size $S$ in (a) and quality $Q$ in (b) as defined by the skewness criteria of eq.2.14. Conventional horizontal $\hat{x}$ and vertical $\hat{y}$ axes apply.

### Boundary Settings, Deposited Ray Power

The outer-most circular boundary is set as a wall with the freeze boundary condition. A deposited ray power sub node is added to track power escaping the fibre.

### Mesh

Physics controlled mesh, Normal, as shown in Fig.5.10.

### Secondary Rays

Set to 1000.

### Trapping Fraction

The trapping fraction is obtained from the 1 W total power by subtraction of the global evaluation of variable 1-'gop.wall1.bsrc1.Qp_int' at the final time simulated.

### Parametric Sweep

Displacement of the central source point is swept in $\hat{x}$ from 0.0 to 0.99 (in 0.01 increments) $\times$ 0.5 m.

### Ray Tracing Settings

Time steps specified as 0 to 50 ns in 0.1 ns increments.

## 5.6 2D Hemispherical Lenses Simulation

### Geometry

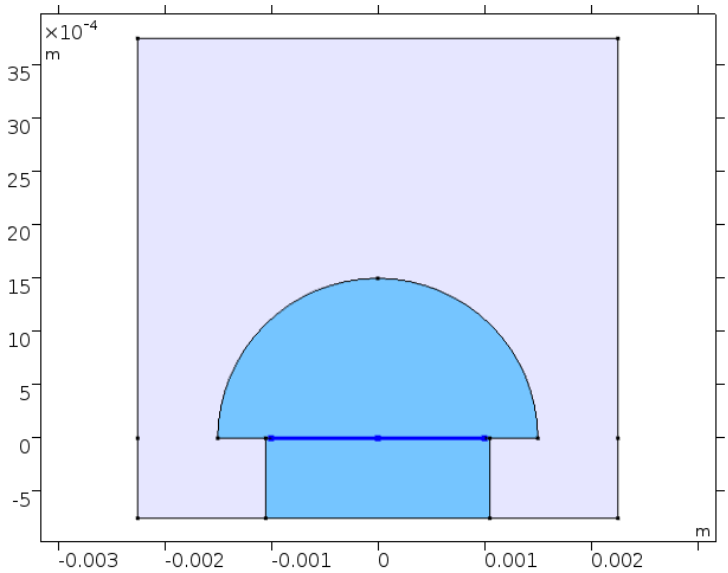


**Figure 5.11:** Schematic of the COMSOL simulation geometry for lens radius = 1.5 mm. The translucent blue rectangle marks the simulation domain and is filled with air ($n = 1$), the solid blue regions represent the YAG / lens domains, where $n = 1.83$.

The geometry for the simulation used to produce Fig.4.10 and Fig.4.11 is shown in Fig.5.11, for a fixed lens diameter. The colour-coded domains are explained in the caption. Rays are released from the dark blue line, via the 'inlet' feature. The simulation boundary is defined by a rectangle of width and height 3 x the variable lens radius. This rectangle is offset from being origin centred in positive $\hat{y}$ by the value of the lens radius. The lens is a circle of radius 'lens radius' with a sector angle set to 180 degrees, it is centred on the origin.

The rectangle that represents the end of the fibre has width 2.1 mm, such that it is slightly

larger than the 2 mm line from which rays are released. Its height is half the lens radius and its position is defined from its corner as x = -1.05 mm and y = -(lens radius)/2.

### Ray Inlet

Any deviations from the default settings are mentioned and explained below:

- Initial Position - Uniform distribution, specifies the distribution of point source along the line.
- Number of rays per release - 10, number of points sources on the line, this introduces some small error as none of the sources are at the very edge of the 2 mm width line source, where outcoupling performance should be worst. Alternatively one could perform the simulation for a series of point releases at well defined locations and calculate their mean result (as is done in Section 5.8].
- Ray direction vector - Conical, Cone axis $1\hat{y}$, specified direction of emission cone.
- Number of rays in wave vector space, 20, (reduced to 10 to produce Fig.4.10 for clarity) number of emission angles sampled.
- Cone angle - 57°, $\pi/2$ - $\beta_c$, the maximum incident angle on the end facet from a ray that can be guided in the fibre.
- Sampling distribution - Random.

### Boundary Settings, Deposited Ray Power

The purpose of the YAG rectangle representing the fibre end, which is 2.1 mm in width (to avoid overlapping of vertices with the ray inlet line) is to simulate rays travelling back down the fibre. It is assumed any ray which does so is never outcoupled. Any ray which reaches the edge of the simulation is terminated there, those than do so above the $y = 0$ line are considered outcoupled and those below it 'lost'. This is achieved by defining two 'walls' following the criteria mentioned, both with the 'freeze' wall condition applied and a 'deposited ray power' sub node defined for each. Note 'deposited ray power' is only available if compute intensity and power is selected in the intensity computation section of the geometrical optics settings main tab. The power out-coupled can then be evaluated using: Results -> Derived Values -> Global Evaluation and evaluating the variable 'gop.wall2.bsrc1.Qp_int' at the last time computed [wall1 / wall2 will depend on how it was defined].

### Anti-reflective Coating

Where indicated in the results an anti-reflective coating was applied to the curved surface of the hemispherical lens. Material Discontinuity -> Thin Dielectric Film -> Thin dielectric films on boundary = Add layers to surface. Release reflected rays = Always. Thin dielectric film1 -> Film refractive index = 1.38 and Film thickness = 600[nm]×0.25/1.38 [118].

### Geometrical Optics Settings

Number of secondary rays set to 400 (200 initial rays), this limits the number of additional rays that can be generated when partial reflections occur, if the number is too small energy will noticeably not be conserved in the system. Intensity computation set to compute intensity and power.

### Secondary Rays

Set to 400.

### Parametric Sweep

Lens radius is swept between 1.1 mm and 3 mm in increments of 0.05 mm. The size of the simulation domain (translucent blue in Fig.5.11) scales with the radius but the size of the line source does not. The width of the fibre rectangle is also fixed, but its height scales such that it always reaches the edge of the simulation geometry.

### Ray Tracing Settings

Time steps specified as 0 to 0.1 ns in 0.002 ns increments, with simulation terminating when no active rays remain. Chosen to allow multiple steps before contact with lens for the smallest lens diameter and for all rays to reach a wall in the largest case.

### Mesh

Physics controlled mesh, Normal, as shown in Fig.5.12.

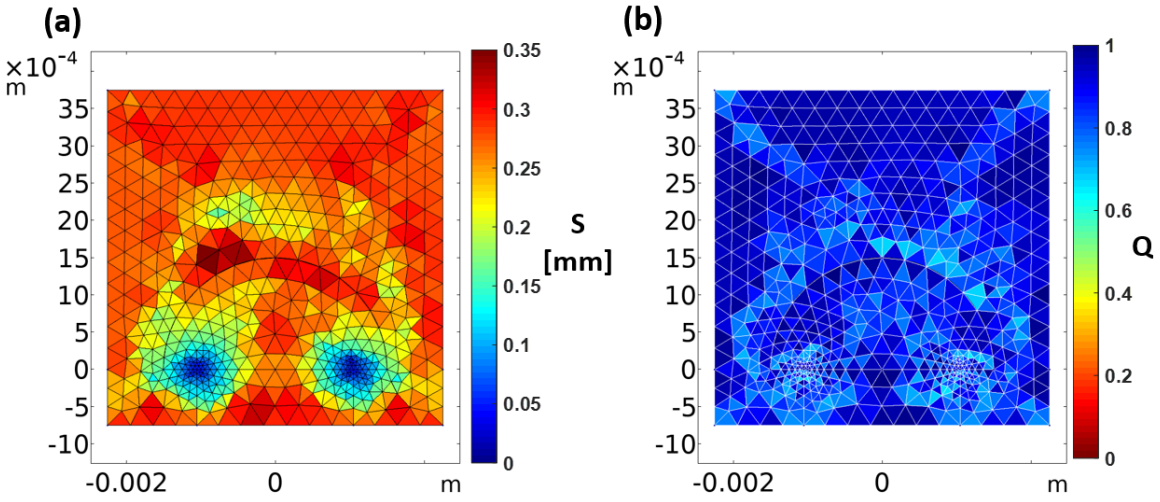


**Figure 5.12:** Normal physics controlled mesh when the lens radius is 1.5 mm. Mesh elements are colour-coded according to their size $S$ in (a) and quality $Q$ in (b) as defined by the skewness criteria of eq.2.14. Conventional horizontal $\hat{x}$ and vertical $\hat{y}$ axes apply

## 5.7 3D Hemispherical Lenses Simulation

As the 3D hemispherical lens simulation follows on directly from the 2D simulation above, we will only specify where the settings have changed with respect to the 2D model.

### Geometry

Fig.5.13 shows how the 2D simulation of Fig.5.11 has been extended into a 3D model. The boundaries of the simulation have simply been extended symmetrically in $\hat{x}$ and $\hat{y}$ to have width, depth and height 3 × the variable lens radius. The block defining the simulation boundaries has been shifted in the positive $\hat{z}$ direction by 1 × lens radius. The sphere used to define the hemispherical lens has the variable radius 'lens radius' and is centred about the origin. A partition with the x-y (z=0) workplane separates the two halves of the lens for later differing material designation. The slightly over 2 mm wide block representing the fibre in Fig.5.11 has been replaced with a cylinder of diameter 0.0021 mm, height half the lens radius and z offset of negative half the lens radius.

### Release From Grid

The rays are released from a hexapolar grid centered at the origin with cylinder axis in the positive $\hat{z}$ direction. The radius is set to 1 mm, to equate to the 2 mm line source of Fig.5.11

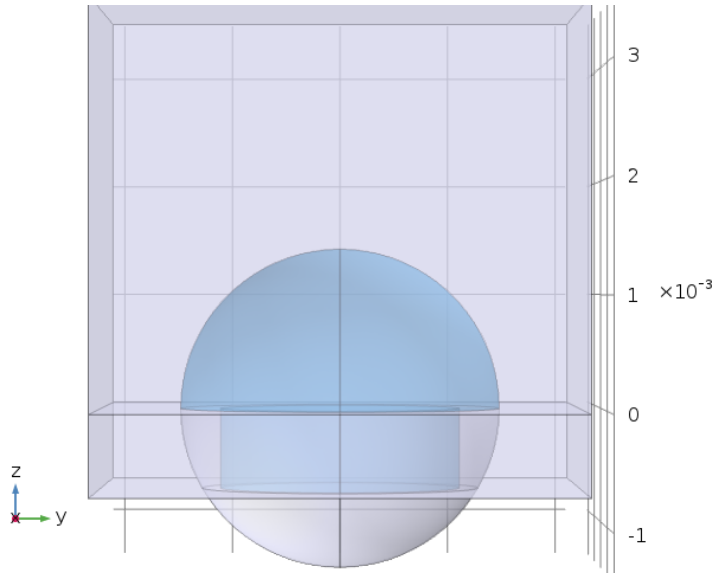


**Figure 5.13:** Geometry used in the 3D COMSOL simulations to produce the results of Fig.4.12, shown here with the variable parameter 'lens radius' set to 1.4 mm. The dark blue sections represent a refractive index of 1.83 (the fibre end and the hemispherical lens), the remaining entities in light blue have an index of 1 (air).

and the number of radial positions is set to 5, yielding 91 release points evenly spread across the fibre facet.

The Ray direction vector is set to conical, of uniform density conical distribution. 10 random wave vector rays are released from each point within the specified $\hat{z}$ axis cone of cone angle (half angle) 57°.

### Ray Tracing Settings

The ray tracing settings were set to the follow time steps in ns 0 to 0.1 in 0.001 increments.

### Boundary Settings

Two walls are defined. With reference to Fig.5.13, the wall that counts outcoupled radiation (wall2) is the cuboidal simulation boundaries above $z = 0$. The other is the outer boundaries of the lower cuboid, below $z = 0$ and ignoring the spherical section that protrudes out of it. Both walls have the freeze boundary condition and deposited ray power sub-node applied. At the last time evaluated the variable 'gop.wall2.bsrc1.Qp_int' provides the outcoupled fraction of the total 1 W power.

### Parametric sweep

The lens radius was swept over the range 1.3 to 3 mm in 0.1 mm increments.

### Mesh

Physics controlled mesh, Normal, as shown in Fig.5.14

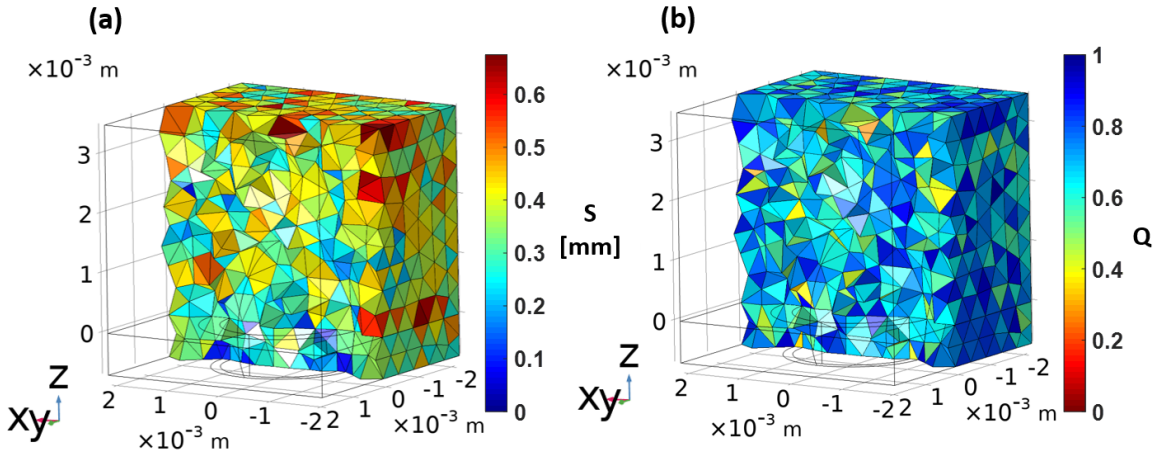


**Figure 5.14:** Normal physics controlled mesh when the lens radius is 1.4 mm. Cut in the $zy$ plane at $x = 0$. Mesh elements are colour-coded according to their size $S$ in (a) and quality $Q$ in (b) as defined by the skewness criteria of eq.2.14.

## 5.8 Modified Ball Lens Simulation

The modified (or tapered) ball lens simulations follow on from the previous Hemispherical Lens Simulation section, as such many of the details remain the same. The key differences are detailed below:

** NOTE ** the simulations were performed for a 2 mm wide fibre source, but the results have been presented for a 1 mm source by reducing all the corresponding dimensions by a factor of 2 (perks of geometrical optics). A 1 mm fibre source was desired for the lateral size constrained comparison mentioned briefly in Section 4.3.4.

The geometry for the simulation used to produce Fig.4.16 is shown in Fig.5.15, for a lens radius of 10 mm and a cone angle of 30°. Note that the rectangle at the base of the ball lens is included to ensure the line over which source positions are swept remains within

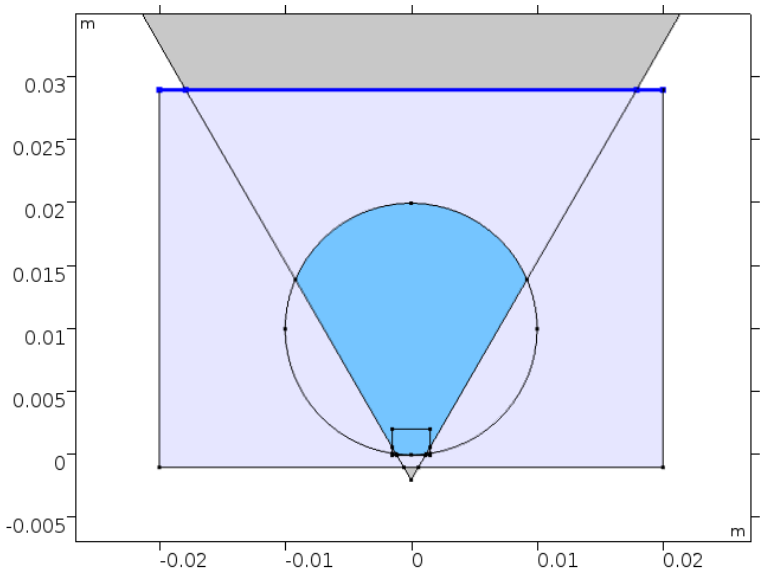


**Figure 5.15:** Schematic of the COMSOL simulation geometry for (variable) lens radius = 10 mm and cone angle of 30°. The translucent blue rectangle marks the simulation domain and is filled with air ($n = 1$), the solid blue regions represent the lens domains, where $n = 2$. Rays are released from the $y = 0$ at $x$ displacements as specified ($0 - 1$ mm lateral displacement). The dark blue line is set to the 'freeze' boundary condition and the rays that reach it have their angle of incidence analysed.

the domains defined as the lens. Some care must be taken in defining the lens domains, as is illustrated in Fig.5.16.

To obtain a better understanding of the ray trajectories in this case, and to have greater control over the source positions than one has when using a line source 'Ray Inlet' feature, we obtain results for values of lateral source displacement (0:0.1:1 mm) individually. The resulting values of intercept angle with the dark blue line in Fig.5.15 are post-processed in MATLAB to obtain the mean for an extended line source representing a fluorescent fibre.

### Cone Definition

The tapered ball lens is defined by the intersection of the circle (of varying lens diameter) and the large cone than extends beyond the simulation domain. The cone is defined by a sector of a circle with the following parameters:

- Radius: lens radius×8
- Sector angle: coneAngle×2 (here coneAngle is the lens taper angle and the sector angle is not a half-angle).

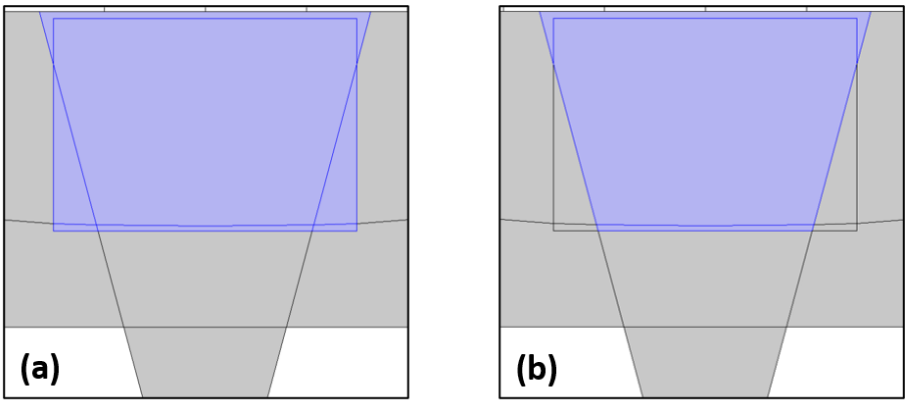


**Figure 5.16:** Close up of the simulation geometry in Fig.5.15, in the region where the ray emission line meets the circle defining the unmodified ball lens. The solid blue regions represent the YAG / lens domains, where $n = 1.83$, elsewhere (grey) $n = 1$. The line source, not shown here, is horizontal and at a height that coincides with the minimum of the circle. Thus for small circle radii the lens domain needs to be extended to enclose the full source width. (a) the crude solution applied initially, causing sub-optimal outcoupling. (b) the refined solution, allowing for more efficient redirecting of rays, yielding the results of Fig.4.16.

- Position: Base: Center, $x = 0$ and $y = -6$ mm/(coneAngle/10), this fixes the cone at a $y$ position that ensures the cone is just wider than the swept source positions at $y = 0$ for all cone angles.

### Anti-reflective Coating

The circular line section of the modified ball lens, that from which rays are intended to escape, has an anti-reflective coating applied. The coating has index 1.38 ($MgF_2$) and thickness 600[nm]×0.25/1.38, where 600 nm is the simulation wavelength [118].

### Parametric Sweep

In COMSOL a parametric sweep is performed for source lateral displacement values (dispx) ranging from 0 to 1 mm in 0.1 mm increments. In addition the cone angle and lens diameter are swept manually by repeating the parametric sweep over dispx in each case.

### Release from grid (ray source)

- Grid type: All combinations: $q_y = 0$, $q_x = dispx$ (the parametric sweep variable ranging from 0 - 1 mm).
- Ray Direction Vector: Conical, $1\hat{y}$, 50° and 100 rays in wave vector space.
- Sampling from distribution: Deterministic, not intentionally chosen over random as previously, but of little consequence with 100 rays in wave vector space.

### Boundary Settings, Freeze vs Disappear

Similar to the hemispherical lens case, we have two distinct wall definitions. The deposited ray power is no longer the parameter of interest here, we instead seek information on the angular distribution of rays exiting the lens. As such we have 'Wall 2', as shown by the dark blue line in Fig.5.15 defined as a 'freeze' boundary condition, and 'Wall 2', the remaining 3 sides of the simulation domain rectangle defined as a 'disappear' boundary. When the angle of incidence parameter is evaluated at the last time calculated of the simulation (see below), only the rays which were frozen will produce values, anything reaching any other boundary is considered lost.

### Secondary Rays

The number of secondary rays here is set to 0. This is because the results to be processed from this simulation are based on the number of rays meeting a certain criteria with respect to the total number of rays. Introducing a potentially variable number of secondary rays, of differing intensity, could add considerably to the complexity and for very little gain in accuracy. This approximation equates to assuming that our anti-reflective coating (which $MgF_2$ would do an excellent job of [118]) does a perfect job.

### Mesh

Physics controlled mesh, Normal, as shown in Fig.5.17

### Ray Evaluation 'gop.phii'

The parameter being evaluated for each value of source position and lens diameter is the distribution of incident angles on the dark blue line in Fig.5.15. This is obtained by: Derived

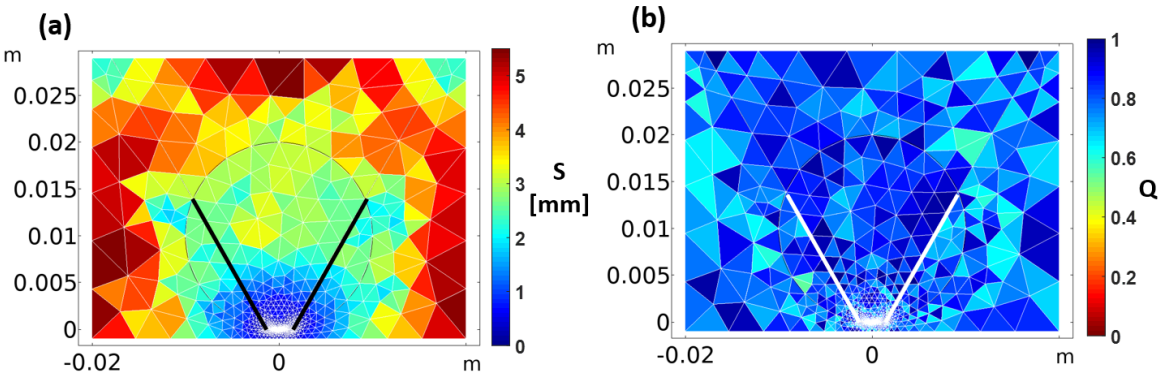


**Figure 5.17:** Normal physics controlled mesh when the lens radius is 10 mm and cone angle is 30°. Mesh elements are colour-coded according to their size $S$ in (a) and quality $Q$ in (b) as defined by the skewness criteria of eq.2.14. Bold black and white lines have been added to highlight the lens taper angle in (a) and (b). Conventional horizontal $\hat{x}$ and vertical $\hat{y}$ axes apply

Values -> Ray Evaluation and evaluating the parameter 'gop.phii'. The data set 'Ray 1 / 2' is defined in the 'Data Sets' section and should be set to the parametric sweep over source position'. Time selection should be set to last and the Parameter selection (dispx) to all. The evaluated table of 'gop.phii' values can then be exported and processed in MATLAB.

### Post-processing 'gop.phii' distributions in MATLAB

To produce the results of Fig.4.16 we first obtained a mean emission angular distribution from the exported 'Phi' values. The first two columns of the evaluated table contain the dispx value and the time value and are not required, the remaining values of incident angle are in radians. Any rays which did not reach the specified 'freeze' boundary, have a value of 'NaN' and are set to have an incident angle of 90°. As the dispx sweeps range only from 0 - 1 mm, we duplicate the rows corresponding to dispx values of 0.1 - 1 mm to represent a full 2 mm line source. This larger table containing 2100 rays in total represents the summed contributions of points on a 2 mm line source and was then processed into a 90 element vector of the number of rays within 1 degree angular bins, ranging 0 - 90° by the MATLAB function 'histogram'. For each value of lens diameter, lens taper angle and emission cone half-angle, the colourscale in Fig.4.16 is the sum of this 90 element vector between its first value (number of rays in bin 0 - 1°) and the element corresponding to the emission cone half-angle. This value is then normalised to the total number of rays, 2100.

# 6 Ghost Imaging

## 6.1 Conventional Ghost Imaging

The term ghost imaging is generally associated with an unconventional imaging technique in which an image is reconstructed through the combined information from two detectors. One detector without spatial resolution, often called a 'bucket' detector, measures the total transmission / reflection from an object that is illuminated with structured light. Simultaneously a second detector, with spatial resolution, measures an illumination pattern which shares some non-zero correlation with that reaching the object. Individually neither detector has the necessary information to reconstruct an image of the object, but through the spatial correlation in these two light beams together they can [33].

In the original demonstration of ghost imaging the spatial correlations were the result of photon pair generation by parametric down conversion [119]. Despite this quantum origin it was soon demonstrated that ghost imaging could be performed with classical light exhibiting purely classical intensity correlations [120, 121, 122, 123]. In this work we remain entirely within the classical aspect of the ghost imaging process.

Fig.6.1(a) depicts one way of performing a ghost imaging measurement. In this case the illumination patterns are randomly generated speckle patterns by coherent illumination of a rotating diffuser. For short exposure times the object can be considered illuminated with a stationary speckle pattern. The illuminating speckle's spatial intensity distribution is obtained from a CCD camera measurement made equidistant from a beam-splitter placed in the illumination path. The rotation of the diffuser then provides a fast and convenient means of generating a new independent illumination pattern. The reconstruction process is then to sum the recorded patterns weighted by their corresponding transmission values as determined by the bucket detector. Some examples exploiting speckle illumination in this way are [122, 123, 124]. Note that the spatial correlations between the illumination and reference patterns need not be perfect or even very strong at all. In [125] weak correlations between the transmitted and reflected speckle patterns of a scattering layer ([126, 127]) were exploited to enable non-invasive imaging through scattering layers with blind ghost imaging.

In 2008 the technique of computational ghost imaging was introduced [34]. The key distinction here is the use of one of a class of devices known as spatial light modulators to generate the illumination patterns computationally. This is the scenario depicted in Fig.6.1(b). Here the patterns are determined by the user instead of measured. The devices that en-

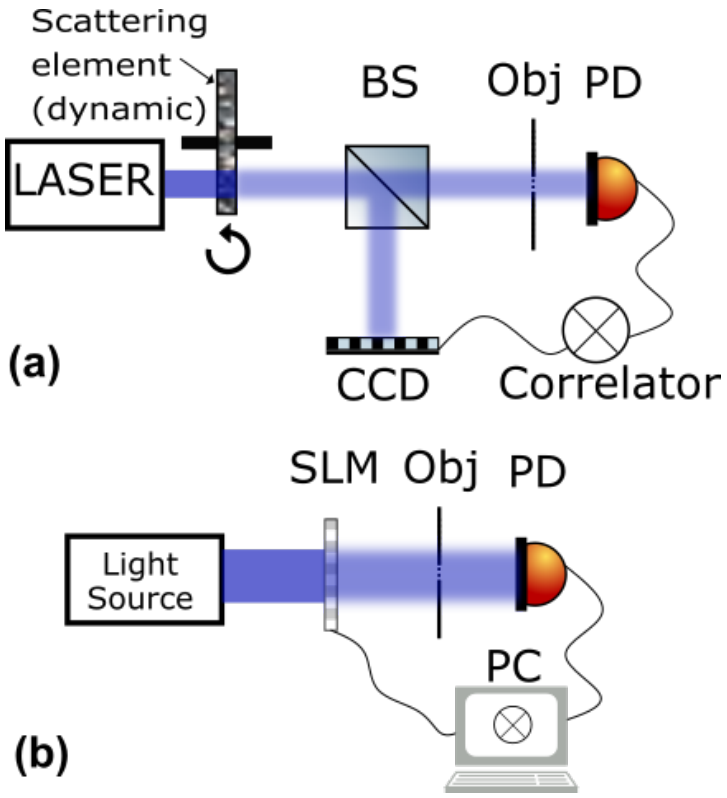


**Figure 6.1:** Classical vs computational ghost imaging. (a) typical setup for a classical ghost imaging experiment. A pseudo-thermal light source, formed here by a laser and a rotating diffuser illuminates an object with a series of speckle patterns. The spatial distribution of light hitting the object is recorded indirectly by means of a beam-splitter and a camera (CCD). Correlating the recorded patterns with their total transmission recorded by the photodiode (PD) allows image reconstruction. (b) as in (a) but the structured illumination is now user-specified by means of a spatial light modulator.

able spatial control of the illumination wavefronts come in two main types. Liquid crystal (LC-SLM) spatial light modulators manipulate the phase of the wavefront pixel by pixel by changes in orientation and thus refractive index of the liquid crystals. The other type, digital micromirror devices (DMDs) are arrays of small mirrors which can be flipped between stable tilt angles in a binary fashion to achieve wavefront amplitude control. DMDs are widely used in projection technology and can operate at KHz frequencies [128], whereas LC-SLMs are limited by the orientation timescales to around 100 Hz [129]. One can even achieve phase control with a DMD via Lee-Holography [130]. The main disadvantage of the DMD is its reduced optical efficiency.

### 6.1.1 Ghost Imaging vs Single Pixel Cameras

Fig.6.1(b) shows the principle behind a computational ghost imaging experiment, if the detector and light sources were swapped, we would then have a schematic for a single pixel camera. The distinction is merely one of where one patterns or masks their wavefront. If the object is illuminated with a patterned wavefront, it would be considered ghost imaging. If the object is uniformly illuminated but the spatial patterning is performed on the reflected or transmitted light, it is considered a single pixel camera [33]. The conventional distinction between computational ghost imaging and single pixel cameras is inconsequential for the work that follows in this chapter, the principles discussed may be applied to either technique.

### 6.1.2 Choice of Basis

The ghost imaging measurement process can be described as follows in Bra-Ket notation, with $N \times N$ pixel projection patterns or images represented as $N^2 \times 1$ element vectors for convenience. The reconstructed image $|I\rangle$ of the object $|O\rangle$ can be written

$$|I\rangle = \sum_j \langle \psi_j | O \rangle \, |\psi_j\rangle \,, \tag{6.1}$$

where $|\psi_j\rangle$ is the $j^{\text{th}}$ pattern in the basis $\boldsymbol{\psi}$ illuminating the object. The inner product $\langle \psi_j | O \rangle$, which becomes the weighting coefficient $c_j$ of $|\psi_j\rangle$ in the reconstruction, measures the spatial correlation between the projected pattern and object and is recorded with a bucket detector. Assuming a planar object of spatially varying transmissivity

$$\langle \psi_j | O \rangle = c_j = \int \psi_j(x,y)\; O(x,y)\; \mathrm{d}x\, \mathrm{d}y, \tag{6.2}$$

where $x$ and $y$ are Cartesian coordinates forming the object plane. $N^2$ orthogonal illumination patterns are required to reconstruct $|O\rangle$, unless one exploits prior knowledge of the object such as sparsity in some basis. Object sparsity is the key to compressed sensing techniques, which are covered in Section 6.2.2.

When the chosen basis of illumination patterns is the standard raster basis, the physical interpretation of equations 6.1 and 6.2 becomes particularly intuitive. For a set of raster patterns of resolution $N \times N$ the complete raster basis is described by an identity matrix of size $N^2 \times N^2$. The $j^{\text{th}}$ row of this matrix describes an $N \times N$ illumination pattern with a single non-zero element of value unity. Comparing successive patterns in the sequence

one would observe the 'on pixel' being raster-scanned across the illumination matrix. To image the object $O$ one raster-scans a single pixel of light across the object and at the $j^{\text{th}}$ (of $N^2$ total) position records the total light transmitted as $c_j$. Reconstructing the image as in eq.6.1 is simply a matter of weighting each pattern in the raster basis with its coefficient $c_j$, which, in the canonical basis is just reshaping the $N^2$ elements of $c$ into an $N \times N$ image.

The raster basis has the appealing properties of simplicity and orthonormality, however it has one major drawback in computational imaging. The fraction of the illumination beam being used for each pattern is just $\frac{1}{N^2}$, which can quickly cause problems with obtaining a sufficient signal-noise ratio (SNR) experimentally. Imaging at higher resolutions comes at the expense of significant reductions to the SNR.

One alternative to the raster basis is a series of random illumination patterns. There are a number of ways of constructing an illumination matrix which could be described as random, or pseudo-random, but a natural choice here would be binary patterns where each pixel has a 50% probability of having a value of one (light on) and otherwise has value zero (light off). This matrix would be described by a Bernoulli distribution with probability parameter $p = 0.5$ [131]. Such matrices have a clear advantage over their raster counterparts, the fraction of total light used is now $\frac{1}{2}$, maximising the signal on the bucket detector. This 50% on / off property can either be left as statistically true or guaranteed for each pattern by the choice of matrix construction method. The cost of using such random matrices is that the property of orthogonality of the patterns has been lost and thus one needs to make more than just $N^2$ measurements. It should be noted that random matrices can be orthonormalised [132, 133] and remain a popular choice for compressed sensing.

Hadamard matrices offer 'the best of both worlds'. Each pattern in a Hadamard basis (with the exception of the first) contains 50% 1s and 50% -1s. The negative values are required to provide orthogonality between each pattern (scalar product between basis rows is 0) and can be implemented experimentally with a difference measurement between two positive only project-able patterns. A Hadamard basis is complete and orthonormal with $N^2$ patterns that use the illumination beam with maximal efficiency.

A Hadamard matrix of order $2n$ can be constructed from the preceding matrix in the sequence, of order $n$, using the Sylvester method [134] as follows

$$H_{2n} = \begin{bmatrix} H_n & H_n \\ H_n & -H_n \end{bmatrix}. \tag{6.3}$$

The sequence begins with $H_1 = [1]$, the first few elements are therefore

$$H_1 = \begin{bmatrix} 1 \end{bmatrix}, \qquad H_2 = \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}, \qquad H_4 = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & -1 & 1 & -1 \\ 1 & 1 & -1 & -1 \\ 1 & -1 & -1 & 1 \end{bmatrix}. \tag{6.4}$$

For example, the third illumination pattern in $2 \times 2$ resolution Hadamard imaging sequence would be given by the $3^{\text{rd}}$ row of $H_4$ and implemented as the difference measurement as

$$|\psi_3\rangle = \begin{bmatrix} 1 & -1 \\ 1 & -1 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 1 & 0 \end{bmatrix} - \begin{bmatrix} 0 & 1 \\ 0 & 1 \end{bmatrix}. \tag{6.5}$$

The SNR advantage of using the a Hadamard basis for ghost imaging is demonstrated experimentally in Fig.6.2. In the interest of a fair comparison, a repeat measurement is made

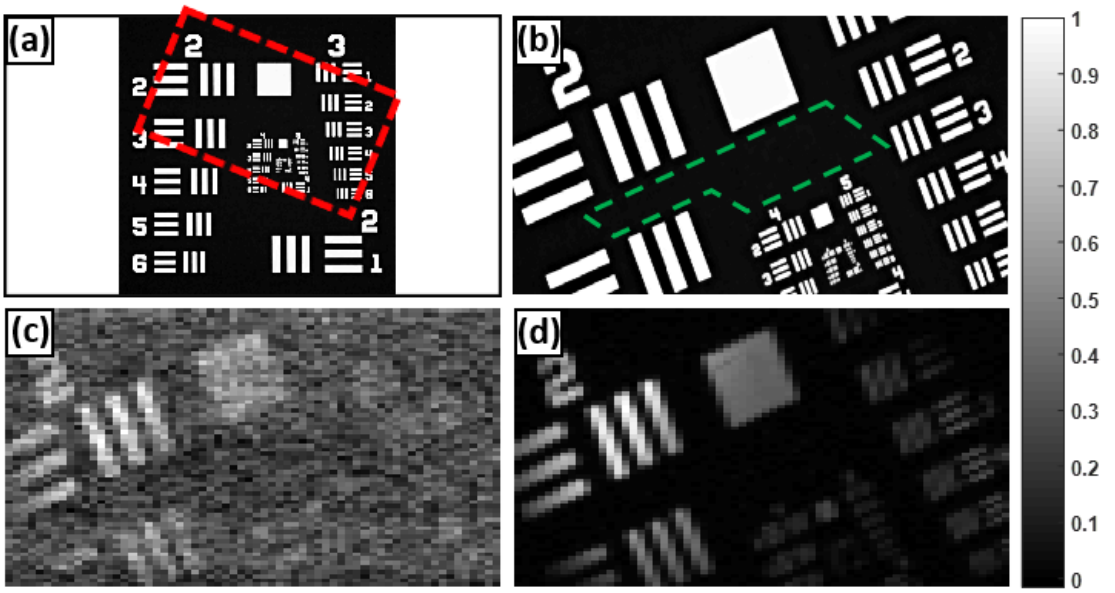


**Figure 6.2:** (a) USAF 1951 negative resolution target (manufacturers image [135]) with the region imaged in the experiments of this chapter marked by a red dashed rectangle. (b) The region highlighted in (a), with the green dashed shape indicating the region defined as the background for later SNR calculations. (c) A ghost image of (b) obtained using a raster measurement basis, showing poor SNR. (d) The same measurement obtained using a Hadamard measurement basis. The raster image in (c) has had a large background subtracted prior to normalisation. Both (c) and (d) have a resolution of 64 × 64 (rectangular) pixels.

for each pattern in the raster basis and the two values for $c_j$ are averaged. This is to match the $2N^2$ measurements needed for the difference measurements in the Hadamard case. Full details of the experimental setup are given in the following section (Section 6.1.3). The benefits of a well informed choice of illumination basis are clear when comparing Fig.6.2(c) and (d).

### 6.1.3 Experimental Method

The experimental setup used for all of the ghost imaging results of this chapter is shown in Fig.6.3. A blue LED (LLS-455 Ocean Optics) is coupled to a 600 $\mu m$ fibre optic patch cord and collimated by plano-convex lens L1 (f = 35 mm). An angled glass slide near the fibre output acts as a beam-splitter and allows fluctuations in luminous output of the source to be tracked and compensated for computationally. Both photodiodes used are Thorlabs Silicon mounted photodiodes (SM05PD1B) connected to the PC via a National Instruments data acquisition (DAQ) device (BNC-2110 connected to PCIe-6361). The DAQ input terminals are set to floating source and are read in differential mode though LabVIEW.

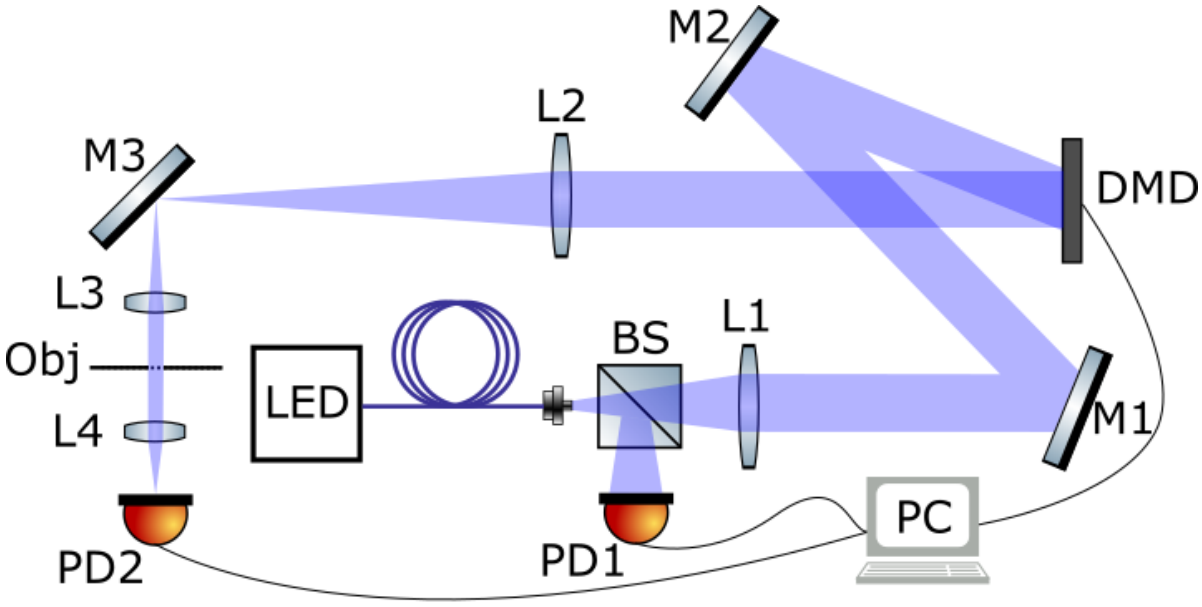


**Figure 6.3:** Schematic of experimental setup used here for computational ghost imaging, not to scale. A fibre coupled blue LED is collimated and illuminates a DMD. The DMD is imaged at reduced magnification by lenses L2 and L3 onto a planar transmissive object (Obj). The transmitted light is then collected by lens L4 and focused onto photodiode 2 (PD2). The photodiode 1 and beam-splitter combination allows fluctuations in source intensity to be compensated for.

The collimated light is angled onto the DMD using a pair of 2" mirrors (M1 and M2) such that the reflection from the DMD on-state is normal to its housing. The DMD module is a SuperSpeed V-9501 from Vialux which includes control electronics for the 1080p resolution 0.95" DMD itself [128]. The DMD is controlled on the PC using LabVIEW. The spatially modulated beam is then imaged onto the object at reduced magnification by plano-convex lens L2 (f = 200 mm) and biconvex lens L3 (f = 35 mm). The object is a Thorlabs 1" Negative 1951 USAF test target (R1DS1N), and is shown along with the illuminated area in Fig.6.2(a) and (b). Finally, the light transmitted by the object is collected by biconvex lens L4 (f = 25.4 mm) and focused onto photodiode 2.

In order to maximise the used area of the DMD the projected macropixels are rectangular with aspect ratio matching that of the DMD display (1920 × 1080). For example, for $N = 64$ we can fit 64 macropixels, 30 DMD pixels wide, perfectly into the 1920 width and for height we take the largest integer macropixel height that allows 64 repeats, which is 16 DMD pixels. The left over pixels in the vertical direction are not used and so our aspect ratio differs very slightly from that of the full DMD and depends on the value of $N$.

The photodiode signals proved to be a persistent source of difficulty. Firstly, after testing all available acquisition settings, on the DAQ block and in software, we were unable to obtain two photodiode signals which were not coupled to each other in some electrical way. This meant that instead of acquiring simultaneous power normalization readings for PD1 for each projected pattern, we had to take a single measurement after each pair of repeat / difference measurements, so that no signal would be simultaneously present in PD2.

Additionally, a non-negligible degree of UK mains electrical supply frequency (50Hz) noise was present in both signals. This led us to restrict our integration times to integer multiples of the period 20 ms, to avoid gaining an additional noise contribution from sampling a non-integer number of these oscillations. Being restricted in available integration times had the knock-on effect of us reducing the lamp intensity considerably, so that we could sample a broader range of signal-noise-ratios in a practical time frame.

The LED power variations were observed to be linear and only ~ 0.2% over a ~ 30 minute period. This high stability is attributed to the reduced LED intensity and its continuous operation over many days. However, this variation is over 4 times larger than the typical signal from 64 × 64 resolution raster illumination of the object and when combined with a large stray-light background signal warranted compensation via PD1.

## 6.2 Ghost Imaging in Alternative Bases

So far we have discussed and demonstrated only the well established basics of computational ghost imaging. We now seek to expand on the idea of choosing the right illumination basis by introducing a method by which traditional image post-processing steps can be incorporated into the illumination basis itself. The principle motivation behind doing so is mitigation of the familiar noise-amplifying effects of certain image filters. We explore this principle through theory, simulation and experiment in the following pages.

We shall begin with the method, repeating eq.6.1 that describes ghost imaging here for convenience

$$|I\rangle = \sum_j \langle \psi_j | O \rangle \, |\psi_j\rangle \,. \tag{6.6}$$

Here the illumination basis $\boldsymbol{\psi}$ is the same basis in which the image is reconstructed, but this need not be the case. A change of basis for an illumination basis $|\psi\rangle$ to a new basis $|\phi\rangle$ can be written as

$$|\phi\rangle = B\,|\psi\rangle \,, \tag{6.7}$$

where B is the matrix that performs the basis change. If one makes this substitution for the illuminating basis in eq.6.6

$$\begin{aligned} |I\rangle &= \sum_j \langle \phi_j | O \rangle \, |\psi_j\rangle \,, \\ &= \sum_j \langle \psi_j | B^T | O \rangle \, |\psi_j\rangle \,, \end{aligned} \tag{6.8}$$

one can see by comparison with eq.6.6 that the reconstructed object effectively becomes $(B^T\,|O\rangle)$. The matrix $B$ can then be chosen such that it performs any operation that can be expressed as a matrix multiplication, directly during the measurement process.

We now consider an example of a matrix $B$ that performs a useful operation upon the object $O$. Many of the ways in which one might process an image before further analysis would be by convolution with a linear filter matrix or kernel. We choose convolution with a simple edge-detection filter as an example with which to demonstrate this process.

The output value of a linear filtering operation at a given pixel location is a function of the neighbouring pixel values in the source image, in a way dictated by the filter kernel. The 2D discrete convolution between an image matrix $I(u,v)$ and filter kernel $K(i,j)$ can be written [136]

$$I'(u,v) = I * K = \sum_{(i,j)\in\Re_H} I(u+i, v+j)\cdot K(-i,-j). \tag{6.9}$$

The following kernel describes a basic edge-detection difference filter

$$K = \begin{bmatrix} 0 & -1 & 0 \\ -1 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix}. \tag{6.10}$$

As one of the simplest useful filters one could conceive, we use this filter kernel as our test case for the remainder of the chapter.

Fig.6.4 shows the result of convolution of a detailed natural image with the edge-detection kernel $K$. The key features of the original object are quite clearly recognisable even in this cluttered image, in fact, almost all of the information in the original image is contained within the regions of sharp changes in intensity extracted by the filtering operation. Indeed edges hold a great deal of significance to the human visual system [136] and edge-detection forms a key component of machine vision and image segmentation processes [137]. In Fig.6.5 we show the same filter applied to the mask describing the USAF target used in the experiments (Fig.6.2(b)), after first resizing it to a resolution of $64 \times 64$ pixels. This serves as a ground truth for subsequent processed experimental measurements at this resolution.

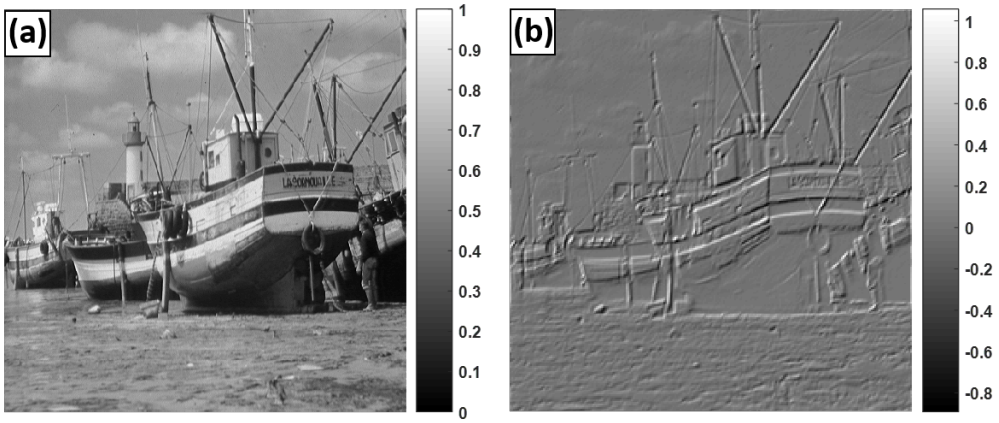


**Figure 6.4:** (a) Traditional fishing boats drying out in harbour. Also known as 'boat.png' from a widely used database of standard test images [138]. (b) The result of convolution of (a) with the edge-detection kernel $K$ given by eq.6.10. The image (a) was first reduced to $128 \times 128$ resolution.

A 2D convolution can be written as a matrix multiplication by writing one of the matrices in the form of a Toeplitz matrix, in which each element is repeated several times

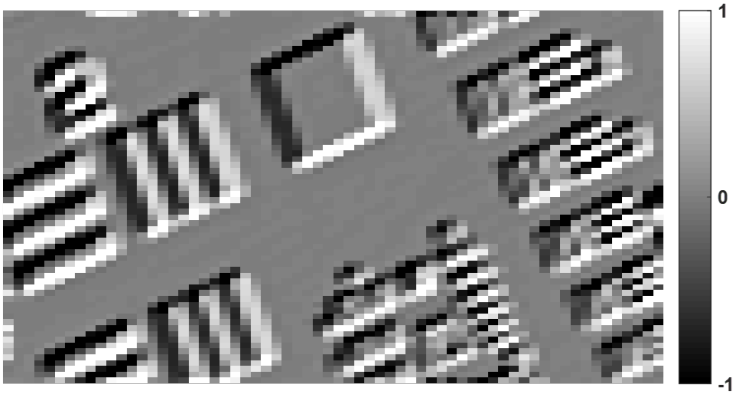


**Figure 6.5:** Ground truth for the 64 × 64 resolution experimental measurements to come. Obtained by convolution of a downsized Fig.6.2(b) with edge-detection kernel $K$ as given by eq.6.10.

across Toeplitz matrix's columns [139]. Briefly, this process involves; zero padding the filter to match the size of the convolution output, forming each row of the filter matrix into a Toeplitz matrix with the same number of columns as the input image and stacking these Toeplitz matrices into a so-called doubly blocked Toeplitz matrix [140]. A vectorised image multiplied by the doubly blocked Toeplitz matrix produces the convolution of the the image with the kernel (as a vectorised output). In this way the matrix $B$, given by the doubly blocked Toeplitz matrix for the canonical case, can be constructed from the filter kernel $K$.

When the unmodified projection basis is the raster basis with 16 × 16 illumination patterns, an example of one of the resulting projected patterns ($\langle\psi_{85}|\,B^T$) can be seen in Fig.6.6(b). The equivalent 'edge-modified' version of this pattern in the popular Hadamard basis can be seen in Fig.6.6(d). It can be seen that each row of the matrix $B$ that performs the convolution, when reshaped to a matrix, is simply the convolution of the original illumination pattern with the kernel $K$. This realisation opens the door for some much simpler methods of generating $B$ in the raster basis, from which other bases can be 'edge-modified' by multiplication with the canonical case.

Just as when post-processing an image with a filter, an issue arises at the borders of the image / projected pattern. When applying a filter to an image there can in principle be no filter result at positions where the filter matrix extends beyond the original image [136]. One must then choose whether to sacrifice some of the size of the processed image or to artificially increase the size of the unprocessed image to compensate. The latter is usually preferable, especially when multiple filters need to be applied. Some common choices for what values to pad the edges with would include zeros, the values at the unpadded image

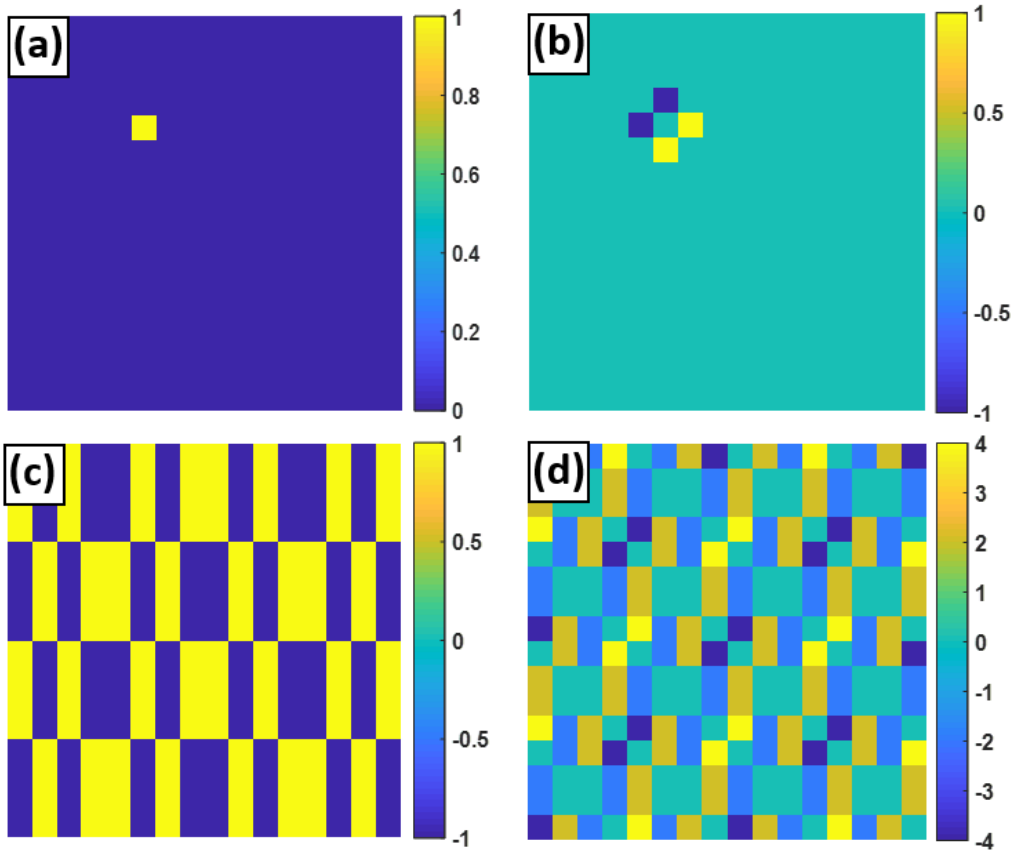


**Figure 6.6:** (a) the 85$^{th}$ projected pattern in a $16 \times 16$ resolution raster basis. (b) The equivalent pattern in the 'edge-modified' basis. Modified such that it performs the edge-detection operation of the kernel $K$ in eq.6.10. (c) and (d), the Hadamard basis counterparts to the top row ((a) and (b)).

edge and the values from the opposite edge. We have chosen the latter, generating patterns of the same size as the originals with periodicity in the horizontal and vertical directions. This periodicity results in some abnormal values at the edges of an image where the original signal was non zero, as can be seen by careful inspection of Fig.6.5.

Now we wish to compare imaging with this 'edge-modified' basis with the results of imaging with a conventional raster or Hadamard basis and then post-processing with the edge-detection Kernel $K$. We can see from the colourscales in Fig.6.6(a) and (b) that for each projected pattern in the raster edge-modified basis a difference measurement between two binary DMD patterns will be needed, whereas in the raster case a single pattern is sufficient. In the interest of a fair comparison, we use the same number of illumination patterns in each case, with each pattern held (and the photodiode signal integrated for) the same time.

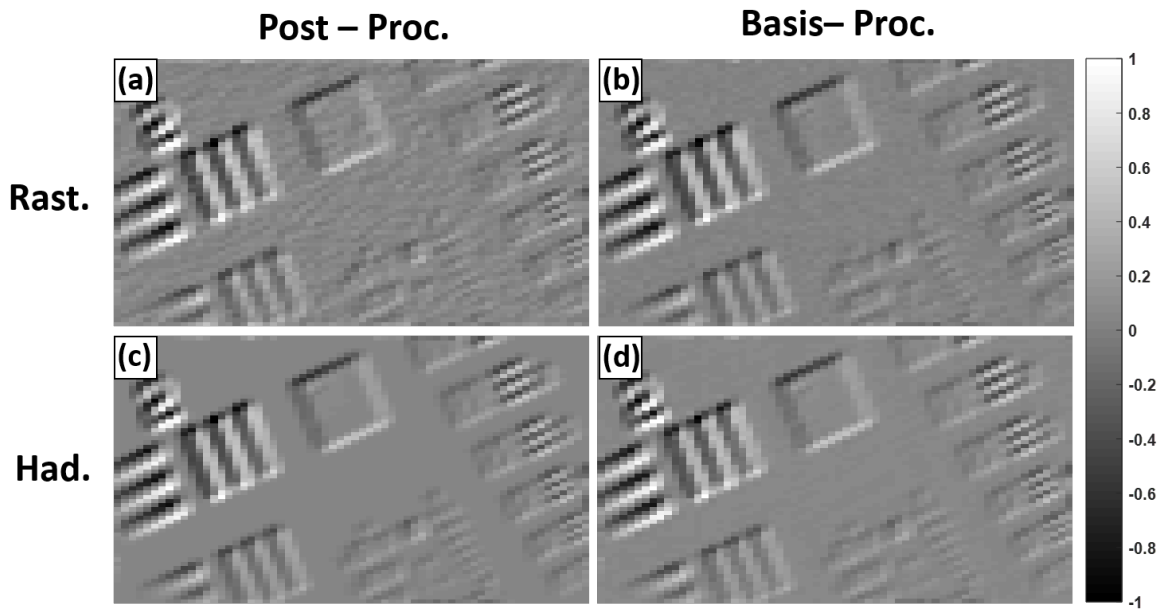


**Figure 6.7:** Comparison of traditional ghost imaging with a post-processing edge-detection step with so-called 'basis-processing' with a modified illumination basis. (a) and (b) - using the raster and edge-modified raster bases. (c) and (d) - using the Hadamard and edge-modified Hadamard bases. Results of experimental measurements with a detector integration time of 400 ms per illumination pattern throughout.

We implement this as an additional repeat measurement for each of the raster patterns, taken immediately following the original in the projection sequence. The final coefficient $c_j$ for a given raster pattern is the mean of the signal for these two measurements.

For the Hadamard illumination basis a difference measurement was already a necessity. However, we can see that after modifying the basis the difference measurement now needs to be performed between two non-binary patterns. This is achieved in experiment by utilising the grayscale capabilities of the DMD. During the specified integration time the DMD repeats the projection of a series of bit planes with reducing duration, such that the fraction of time in which a given pixel was in the on state over the cycle corresponds to its fraction of the full 8-bit depth.

The experimental results of the comparison between post-processing and 'basis-processing', for both raster and Hadamard starting bases can be seen in Fig.6.7. It is clear that incorpo-

rating the filtering step into the illumination basis does indeed perform the intended edge-detection operation, as in each case the images match well with the ground truth in Fig.6.5. The key difference seemingly lies in a lack of spatial uniformity in the illumination and or collection efficiency achieved in experiment.

The images in Fig.6.7(a) and (b) are very similar, with an apparent increase in the background noise in the post-processed raster measurement in (a). In the higher signal-to-noise ratio (SNR) results obtained using the Hadamard basis in (c) and (d) the results are almost indistinguishable, with slight background noise visible in the basis-processed case. As one might expect the two methods appear to be visually indistinguishable when the SNR is high, with some subtle differences arising at the lower SNRs in seen in Fig.6.7(a) and (b). We now compare the results for varying noise levels by adjusting the projection / integration time.

We shall focus on comparing the raster and edge-modified raster bases for the remainder of this work. The reason being that whilst the edge-modified raster patterns can be implemented as two step difference measurement between binary patterns, it can be seen in Fig.6.6 that the edge-modified Hadamard basis requires the projection of intermediate grayscale values. This adds to the overall measurement time and complicates the process of constructing a fair comparison between the two methods. In addition, one can quickly see that the edge-modified Hadamard basis will be outperformed in SNR by its unmodified counterpart, as it no longer retains its efficient 50:50 on/off pixel ratio.

Fig.6.8 compares our method of image processing with a modified projection basis ('basis-processing'), with conventional post processing of the raster basis. The experimental images were obtained at three different integration times, to offer a comparison over a range of SNRs. The basis-processed image is the higher fidelity reconstruction of Fig.6.5 in each case, with the visual improvement becoming less pronounced with increasing SNR. In addition, the spatial characteristics of the background noise appear to differ in each method. We will return to this observation in Section 6.2.1 after first quantifying and understanding the apparent improvement in SNR for the basis-processing method.

To estimate the signal-to-noise ratio in our experimental images we compare the intensity values in defined peak ($\langle I_P \rangle$) and background ($\langle I_{BG} \rangle$) regions as

$$SNR = \frac{\langle I_P \rangle - \langle I_{BG} \rangle}{\sigma_{BG}}, \tag{6.11}$$

where $\langle .. \rangle$ denotes the mean and $\sigma_B$ the standard deviation of the intensity in the background region. The background region is selected manually and is shown in green in Fig.6.2(b).

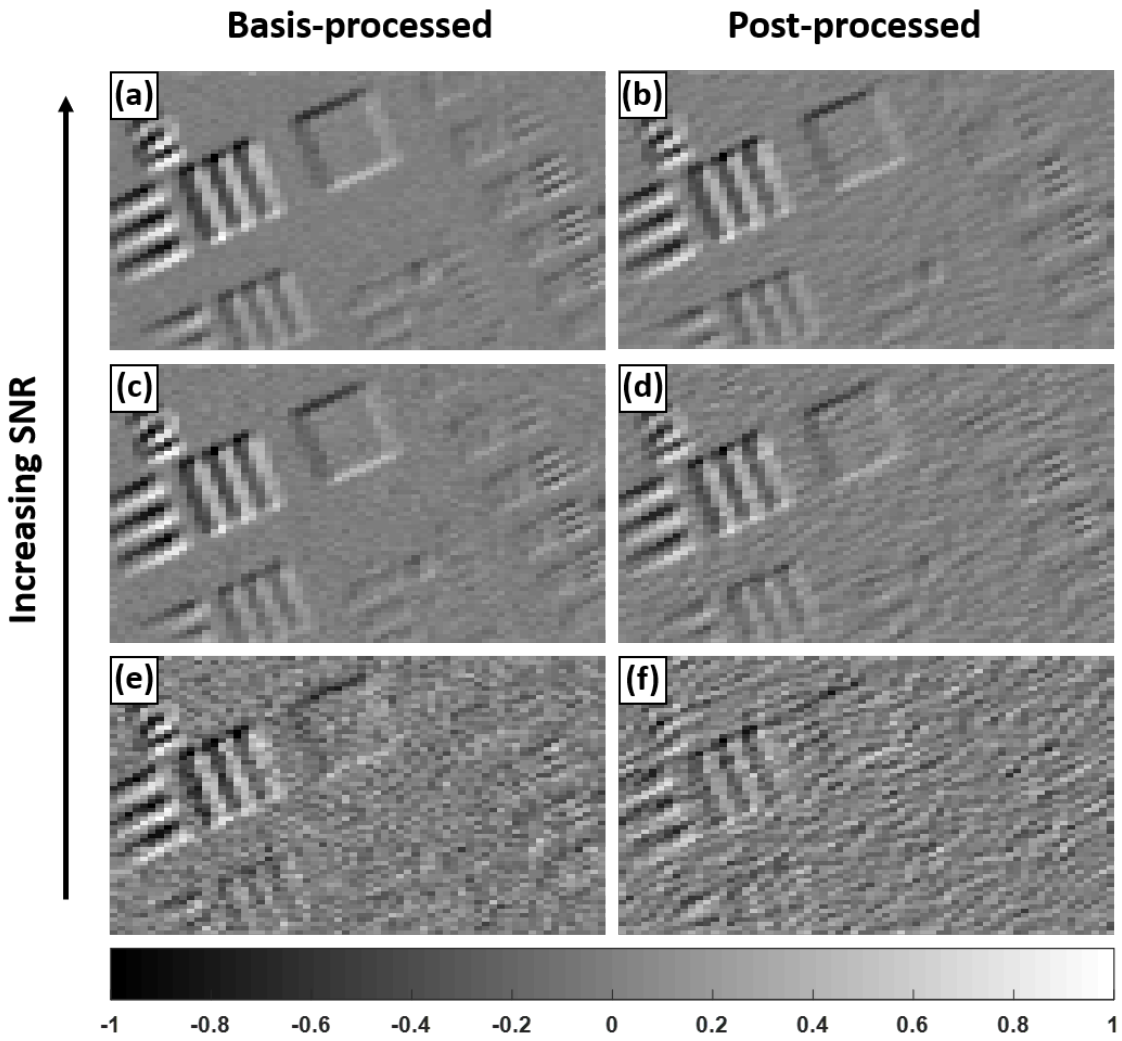


**Figure 6.8:** Comparison between the ghost images obtained using a modified projection basis ('basis-processed', left column) and those measured with a raster basis and then convolved with the edge detection kernel $K$ ('post-processed', right column). The three rows show varying detector integration times, increasing from bottom to top as 20, 100 and 220 ms. The images are 64 × 64 resolution.

As the peak intensity values in the images processed for edge detection are somewhat sprawling, we define the peak regions as the regions with the top 10% of intensity values, taken from a high SNR experimental image. The pixels at the image edges (within 2 pixels) were excluded from availability for defining the peak regions to avoid including artefacts of the choice boundary conditions. The calculated SNRs from experimental images for the two methods at a range of exposure times are shown in Fig.6.9. The visual advantage of the basis-processing method is indeed seen here quantitatively over the full range. The ratio

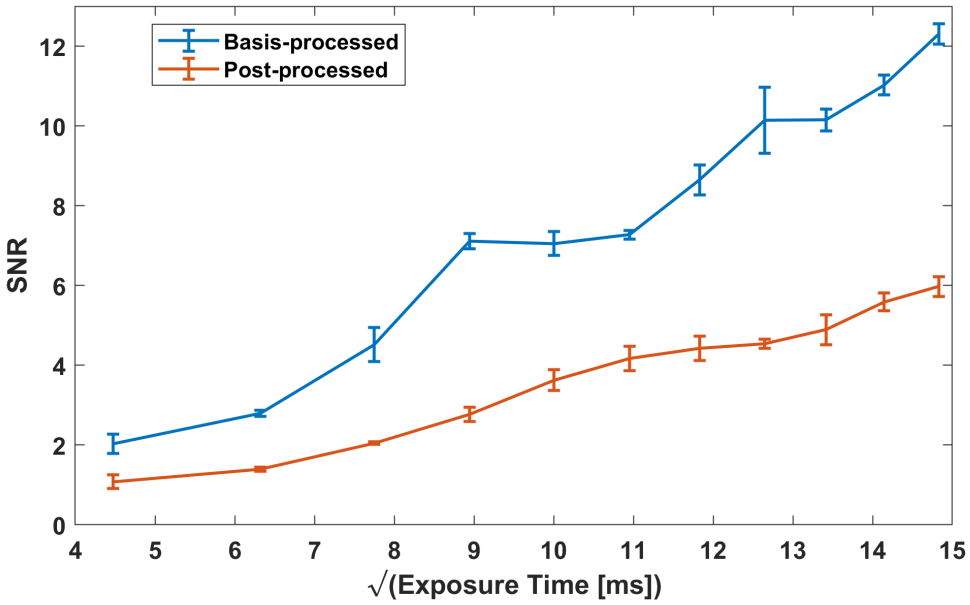


**Figure 6.9:** Comparison of calculated SNR from experimental images acquired using the basis vs post-processing methods. SNRs are presented against the square root of the image exposure (integration) time, where roughly a linear dependence is expected. The error bars are calculated from three repeat measurements. Images were 64 × 64 resolution.

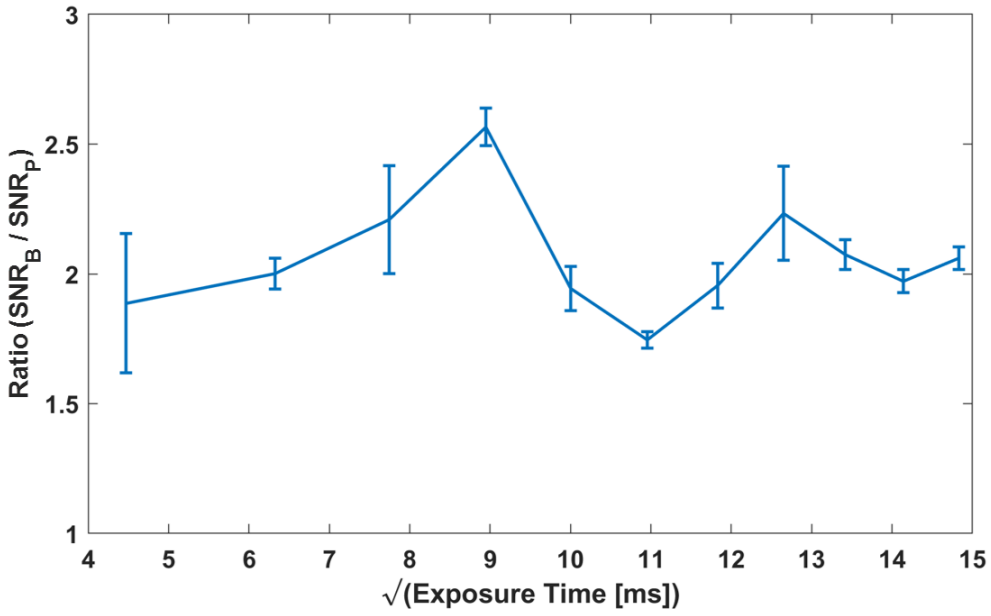


**Figure 6.10:** Ratio of the SNRs from the two methods compared in Fig.6.9, Basis-processed ($SNR_B$) / Post-processed ($SNR_P$).

between the two data sets is shown in Fig.6.10 and shows a mean SNR improvement of a factor of ~2.

We now consider a simple analytical treatment of the expected SNRs for each method. We do not include the additional power normalisation step in the measurement process. If the repeat raster measurements in the post-processing method are averaged, then the noiseless signals for each method can shown to be equivalent. The difference then lies in the image noise.

The motivation for exploring this technique of 'basis-processing' was that the noise amplification that can come with post-processing steps could be avoided by integrating this step into the illumination basis. We consider additive white Gaussian noise as a model for the detector noise in our measurement system, which we assume here to be the dominant noise source. If an image is corrupted by additive white noise, a given filter kernel $K$ will increase the standard deviation of the noise by a factor or $\sqrt{E_k}$, where $E_K$ is the "energy of the filter" defined as [141]:

$$E_K = \int_{\Re^2} |K(x,y)|^2 \mathrm{d}x\mathrm{d}y. \tag{6.12}$$

For $K$ equal to the simple edge-detection filter of eq.6.10 used in the presented experiments, $E_K = 4$ and the filter increases the standard deviation of the noise by a factor of 2.

The full signal (including noise), which we denote $S_P$ for the post-processing method can be mathematically described as

$$S_P = K * \left( \frac{(O+\sigma_1)+(O+\sigma_2)}{2} \right), \tag{6.13}$$

where $K$ is the edge detection kernel in eq.6.10 and $*$ denotes the 2D convolution as in eq.6.9. The additive Gaussian detector noise introduced with each measurement of the object $O$ is denoted $\sigma_1$ and $\sigma_2$ of equal standard deviations. Combining the two independent sets of detector noise in quadrature and denoting the edge-processed object as $O_E$ we obtain

$$S_P = O_E + K * \left( \frac{\sqrt{2}\sigma}{2} \right). \tag{6.14}$$

Finally, following eq.6.12 for the factor or 2 increase in the standard deviation of the noise following filtering

$$S_P = O_E + \sqrt{2}\sigma. \tag{6.15}$$

Similarly for the basis-processing method we describe the measurement process as

$$S_B = ((K_1 * O) + \sigma_1) - ((K_2 * O) + \sigma_2), \tag{6.16}$$

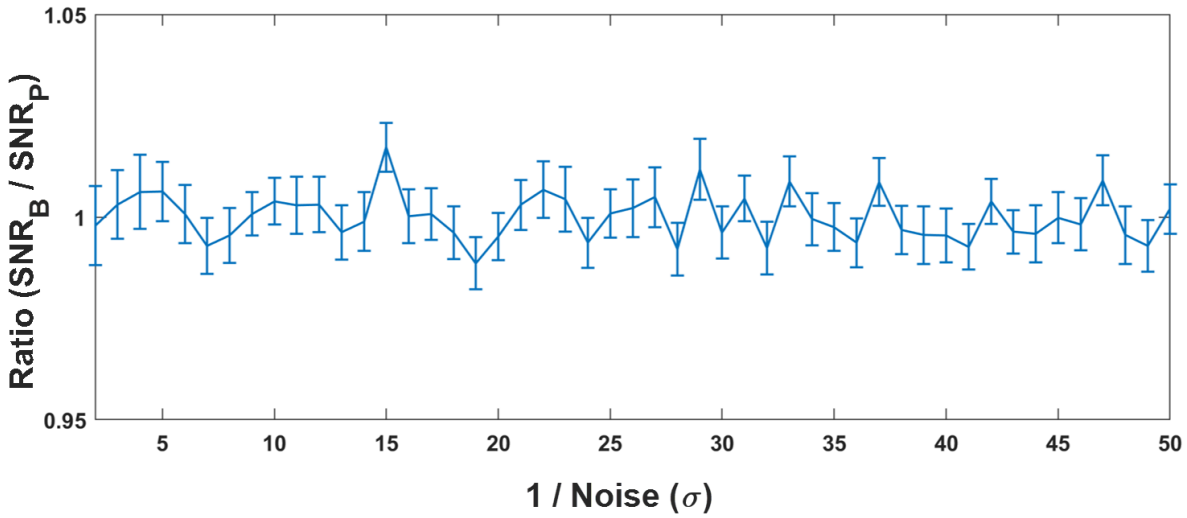


**Figure 6.11:** Simulated ratio of the SNRs from the methods two methods (Basis-processed ($SNR_B$) / Post-processed ($SNR_P$)). The reciprocal of the magnitude of the simulated detector noise takes the place of the square root of the exposure time in the experimental data. The error bars are calculated from 50 repeat simulations with new random noise sets.

where $K_1$ and $K_2$ are the projectable binary filters that form $K$ with their difference. The resulting full signal for this method is then

$$S_B = O_E + \sqrt{2}\sigma. \tag{6.17}$$

The predicted ratio of $SNR_B$ / $SNR_P$ is then clearly unity in this case. The reduction in noise through making the edge-detection process 'before the measurement noise' is perfectly counteracted by the need to make twice as many measurements.

This result is confirmed through simulated measurements in Fig.6.11. The simulations are performed on the image shown in Fig.6.2(b) with its resolution reduced to $64 \times 64$. Additive white Gaussian noise of zero mean is applied as shown in eq.6.13 and eq.6.16. At noise level $\sigma = 1$, the standard deviation of the added Gaussian noise is unity. For comparison the simulated object $O$ has been normalised to a peak intensity of unity. SNR calculations are performed in the same way as for the experimental data. Neither the analytical treatment nor the simulations (in the from described thus far) describe the trends in the experimental data, in which a clear advantage to basis-processing is observed. We find an explanation with a slight increase in the sophistication of our simulated measurements.

We have neglected the additional power normalisation measurement from the treatment given above. This step is included to counteract small fluctuations in the light source in-

tensity, but is only a requirement for the post-processing method. The difference measurement procedure used in the basis-processing method does an excellent job of compensating for these fluctuations provided they are slow relative to the measurement time. Nevertheless it has been applied to both methods in the spirit of a fair comparison.

If the noise introduced during power normalisation is added to the analysis shown in eq.6.13 and eq.6.16 as a single multiplicative term common to both measurements, as it is in experiment, it remains as such and has no bearing on the ratio of the two methods. This is confirmed in simulations.

There is however evidence to suggest that the power normalisation step is indeed the culprit behind the differing SNRs observed in experiment and the analyses presented above are in some way incomplete. If we remove the power normalisation step from the experimental data presented in Fig.6.8, the basis-processing SNRs remain favourable. In addition the post-processing SNRs become noticeably non-linear, levelling off at higher exposure times. This is all consistent with background noise introduced in the raster basis measurements by the observed source intensity fluctuations, which increase in spatial frequency with increased measurement time.

Fig.6.12 shows the calculated SNRs from experimental measurements with varying exposure, using both the basis-processing and post-processing methods as in Fig.6.8, but at $32 \times 32$ resolution and without the power normalisation step. Without the power normalisation step the experimental data show good agreement with the predicted equivalence of the SNRs, but only for quick measurements where the light source is relatively stable.

Finally, we find evidence through simulation to support that the differing SNRs between the two methods seen in experiment is a result of the power normalisation step. We have discussed previously the advantages of a difference measurement in compensating for fluctuations in the background light levels. In the experimental setup the level of stray light reaching the detector PD2 is substantial when compared to that from a $64 \times 64$ resolution raster pixel and so the details of the reconstructed raster image sit atop a large background. Fig.6.13 shows how adding a constant background to the simulated measurements can reproduce the change the gradient of the SNR vs reciprocal noise level plot for the post-processing method, when the power normalisation step is also applied. The basis-processing method remains unaffected by the presence of a constant background, leading to an SNR advantage. In our simulations the peak object transmission value is normalised to one as it has been throughout and the value of the added background $\beta$ is as shown in Fig.6.13. It is only in the presence of the multiplicative noise introduced during power normalisation that the background effects SNR. In our simulations the noise added through

power normalisation is implemented by multiplying both measurements by a factor of $(1 + \sigma_3)$, where $\sigma_3$ has the same statistics as $\sigma_1$ and $\sigma_2$.

We now have an explanation for the SNR advantages for the basis-processing method demonstrated experimentally and have seen good agreement with theory in situations where the power normalisation step is not needed (Fig.6.12, low exposure times). We've seen that if the number of measurements is kept the same and a secondary detector provides source power normalisation, both the post-processing and basis-processing methods should yield the same SNRs. This is however subject to the accurate removal of any background signal from the detector PD2, without which basis-processing will be advantageous. Furthermore if no power normalisation is applied the difference measurements that form this basis-processing example remove the background noise introduced by an unstable light source, as shown in Fig.6.12. Finally, if no secondary detector is used and power normalisation is instead obtained by replacing the repeated raster measurement with a background subtraction in which all DMD pixels are off, a factor of 2 advantage in SNR will be observed following the logic presented in eq.6.13 and eq.6.16.

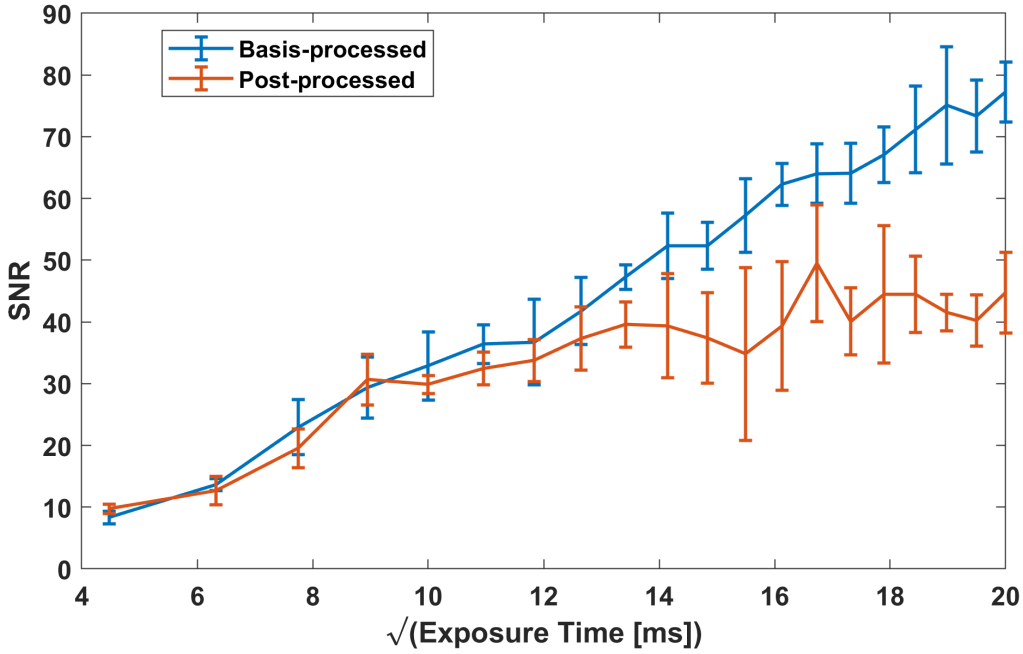


**Figure 6.12:** Comparison of calculated SNR from experimental images acquired using the basis vs post-processing methods, as in Fig.6.9 but at 32 × 32 resolution and no power normalisation step performed. The error bars are calculated from five repeat measurements.

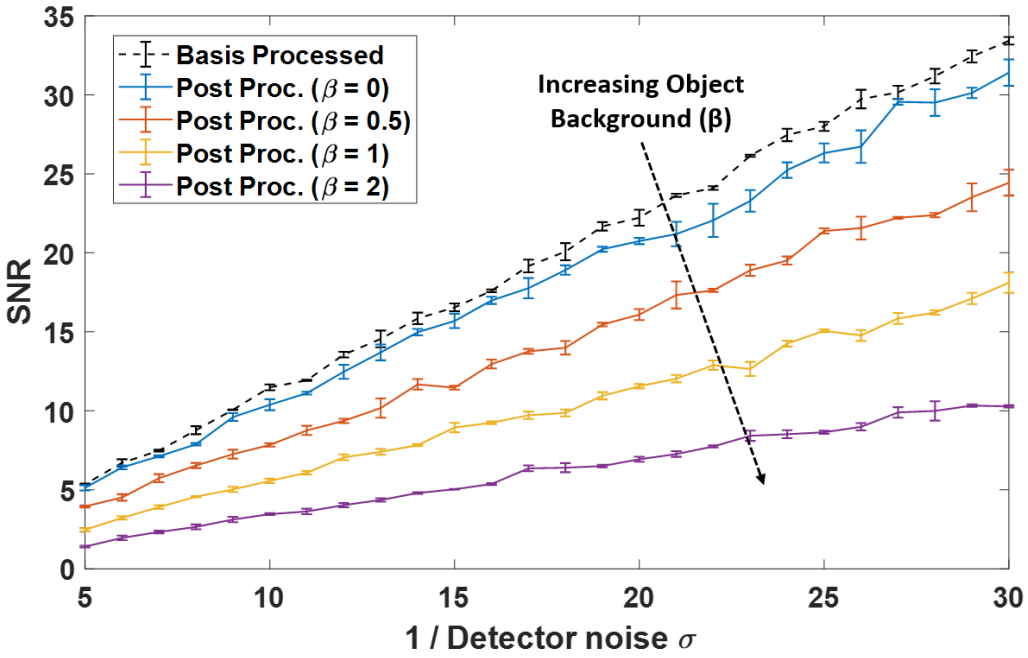


**Figure 6.13:** Effect of background signal on the SNRs the basis-processing and post-processing methods. Results of simulated measurements with varying detector noise level ($\sigma$) and constant background added to the object (as indicated by the parameter $\beta$). The small offset at $\beta = 0$ is attributed to a small initial background in the object JPEG.

### 6.2.1 Noise Spatial Correlations

In addition to the SNR differences that can be seen between the basis-processing and post-processing methods shown in Fig.6.8, one can see a difference in the spatial characteristics of the noise in each case. The explanation is fairly straightforward, in the post-processing method the noise has had its spatial statistics modified by the application of the edge-detection filter.

In Fig.6.14(a) we see a 64 × 64 image containing random Gaussian distributed noise with zero mean and standard deviation of 1. This noise is white, meaning that has a constant power spectral density (over the range accessible here). This can be seen in image In Fig.6.14(b), the magnitude of the Fourier transform of (a), as a proxy for the power spectrum. A realisation of this noise has equal probability of occupying the spatial frequencies shown in (b). In the basis-processing method the noise remains white.

In Fig.6.14(c) shows the same noise as in (a) after it has been convolved with edge-detection kernel $K$ (eq.6.10). The noise, denoted $N$ here, now exhibits some spatial correlations. By

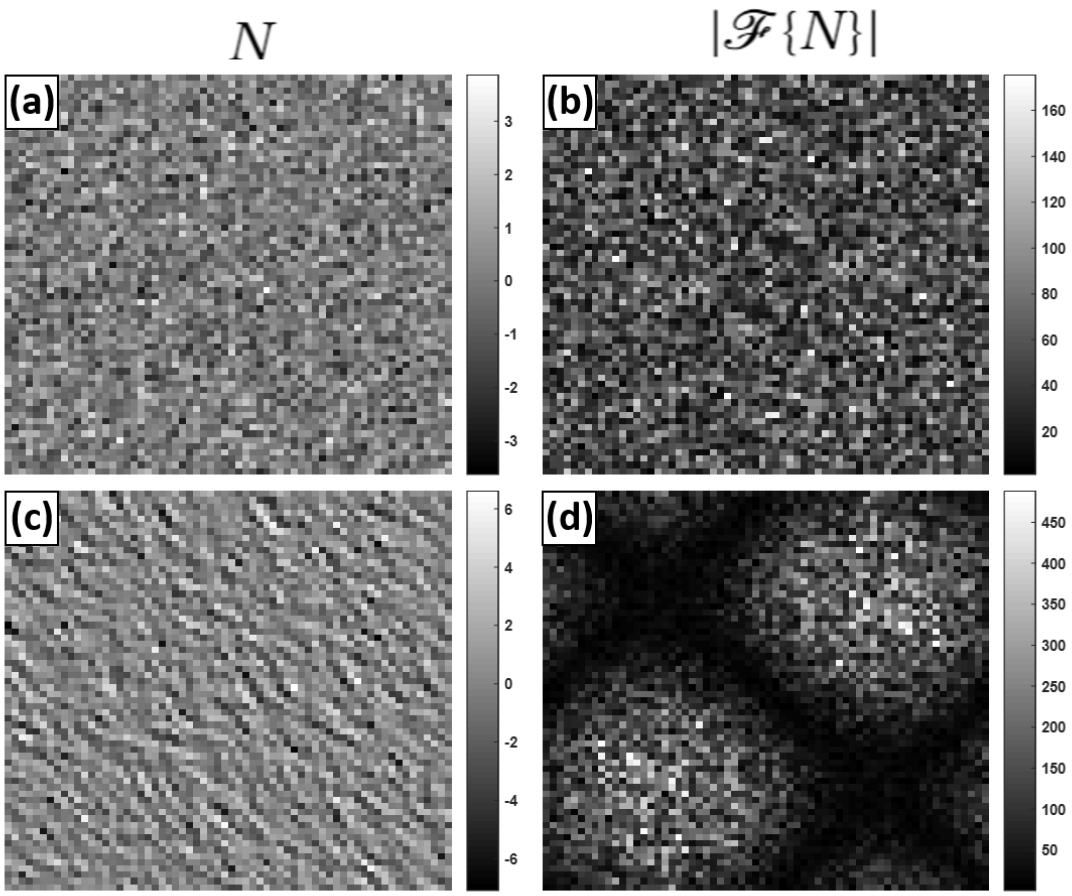


**Figure 6.14:** Spatially correlated noise after post-processing. (a) Normally distributed random noise with zero mean. (b) Fourier Transform of (a). (c) $K$ convolved with (a). (d) Fourier Transform of (c).

convolution theorem

$$\mathscr{F}\{K * N\} = \mathscr{F}\{K\} \cdot \mathscr{F}\{N\}, \tag{6.18}$$

where $\mathscr{F}$ denotes the Fourier transform. We can see then that modified spatial distribution of the noise will be determined by Fourier transform of the filter kernel $K$. This is shown for the edge-detection filter (eq.6.10) in (d).

Correlated background noise has the appearance of containing useful information and will be more likely to be mistaken for structure in the final image, than its uncorrelated (white) counterpart (compare (a) and (c)). Thus we can see that even in situations where the SNRs of the two methods are equivalent, there exists a subtle difference between the results.

### 6.2.2 Compressed Sensing

In this section we briefly compare the results of traditional post-processed ghost imaging with our novel basis-processing technique when used in conjunction with compressed sensing methods. The central concept behind compressed sensing is using some known or assumed prior knowledge of the object to be imaged to reduce the number of measurements needed to reconstruct it. This prior knowledge comes in the form of the sparsity of the coefficients describing the object when it is represented in an appropriate basis. A signal is considered sparse when most of its elements are equal to zero, or in practice are close enough to zero.

Natural images are generally sparse in the Discrete Cosine Transform (DCT) basis, which is a relative of the Fourier basis that uses only cosine functions of varying wavenumbers and operates on real valued signals. This basis is commonly used for the compression of images for this reason [136]. Fig.6.15 shows a histogram of the DCT coefficients of the boat test image seen in Fig.6.4(a). Noting the logarithmic y-scale, we see that the vast majority of the DCT basis coefficients have a value close to zero. During image compression, only a fraction of the highest value coefficients are retained, reducing the file size substantially. If the original image was sufficiently sparse in this basis the compressed image can be almost indistinguishable from the original.

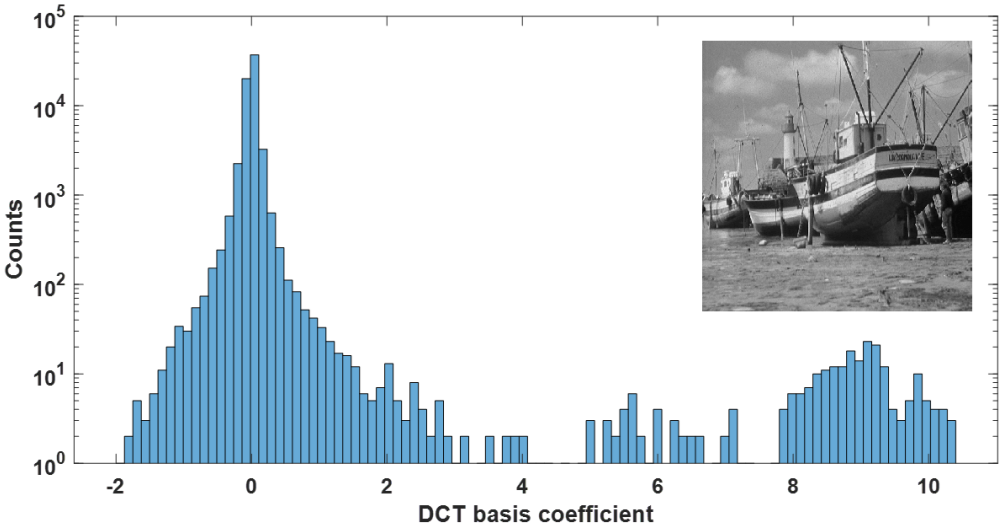


**Figure 6.15:** Sparsity of a natural image in the Discrete Cosine Transform basis. Histogram of the DCT basis coefficients of the 128 × 128 resolution image of boats inset [138].

The key to compressed sensing is to exploit this sparsity during the measurement or 'sensing' process, thereby reducing the required number of measurements. Once one knows

a basis in which their object is sparse, the trick is not to image in this sparse basis, where the chance of finding the non-zero coefficients with a reduced number of measurements is poor, but instead to choose an illumination basis that is mutually incoherent with it. The mutual coherence of two bases is proportional to the maximum value of the spatial overlap between any combination of patterns from each basis [142]. By choosing an illumination basis with low mutual coherence with the sparse basis one evenly samples the sparse basis coefficients with each measurement.

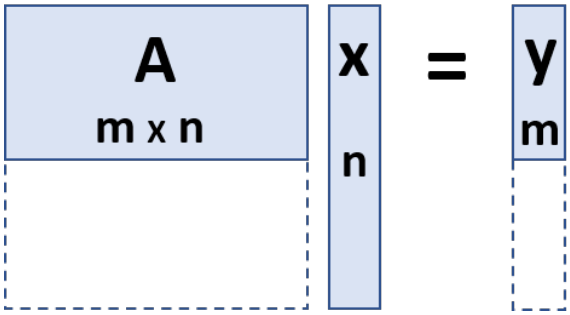


**Figure 6.16:** Graphical illustration of the matrix equation describing compressed sensing. $\boldsymbol{A}$ is the matrix of illumination patterns, $\boldsymbol{x}$ is the object / signal to be reconstructed and $\boldsymbol{y}$ contains the measurement coefficients

The underdetermined system of linear equations that describes the compressed measurement process is shown graphically in Fig.6.16. $\boldsymbol{A}$ is the matrix of illumination patterns, $\boldsymbol{x}$ is the object / signal to be reconstructed and $\boldsymbol{y}$ contains the measurement coefficients. As the measurement is compressed, $m < n$, $\boldsymbol{A}$ is no longer invertible and infinite solutions for $\boldsymbol{x}$ exist for the relation $\boldsymbol{Ax} = \boldsymbol{y}$.

With the knowledge of constraints on the object $\boldsymbol{x}$ such as sparsity in some basis, the correct solution can be determined. The sparsity condition is technically satisfied by finding the solution for $\boldsymbol{x}$ that has minimal $L^0$ - norm. In practice however it has been shown that in most cases the simpler problem of minimising the $L^1$ - norm is equivalent [142]. Here we choose to use a popular algorithm referred to as the Iterative Shrinking-Thresholding (ISTA) algorithm [143].

Starting from some informed guess for $\boldsymbol{x}$ or a zero-valued initial condition, the improved solution for $\boldsymbol{x}$ at the next iteration is calculated as [143]

$$\boldsymbol{x}_{k+1} = \mathcal{T}\{\boldsymbol{x}_k - \delta(\boldsymbol{A}^T(\boldsymbol{A}\boldsymbol{x}_k - \boldsymbol{y}))\}, \tag{6.19}$$

where $\delta$ is an adjustable parameter. The thresholding step denoted $\mathcal{T}\{..\}$ is

$$\mathcal{T}\{\boldsymbol{x}_i\} = \mathbf{sgn}(\boldsymbol{x}_i)(|\boldsymbol{x}_i| - \alpha)_+, \tag{6.20}$$

where $\alpha$ is the other adjustable parameter and $(..)_+$ denotes retention of only the positive values. **sgn** denotes the sign function, which has value +1, 0 or -1 if the real input is positive, zero or negative respectively. The shrinkage / soft-thresholding step increases the sparsity of the current solution for $\boldsymbol{x}$ by shifting its coefficients towards zero with each iteration. By adjusting the values of the parameters $\alpha$ and $\delta$ one can dictate how fast and with what weighting the thresholding process and information from measurements $\boldsymbol{y}$ are applied.

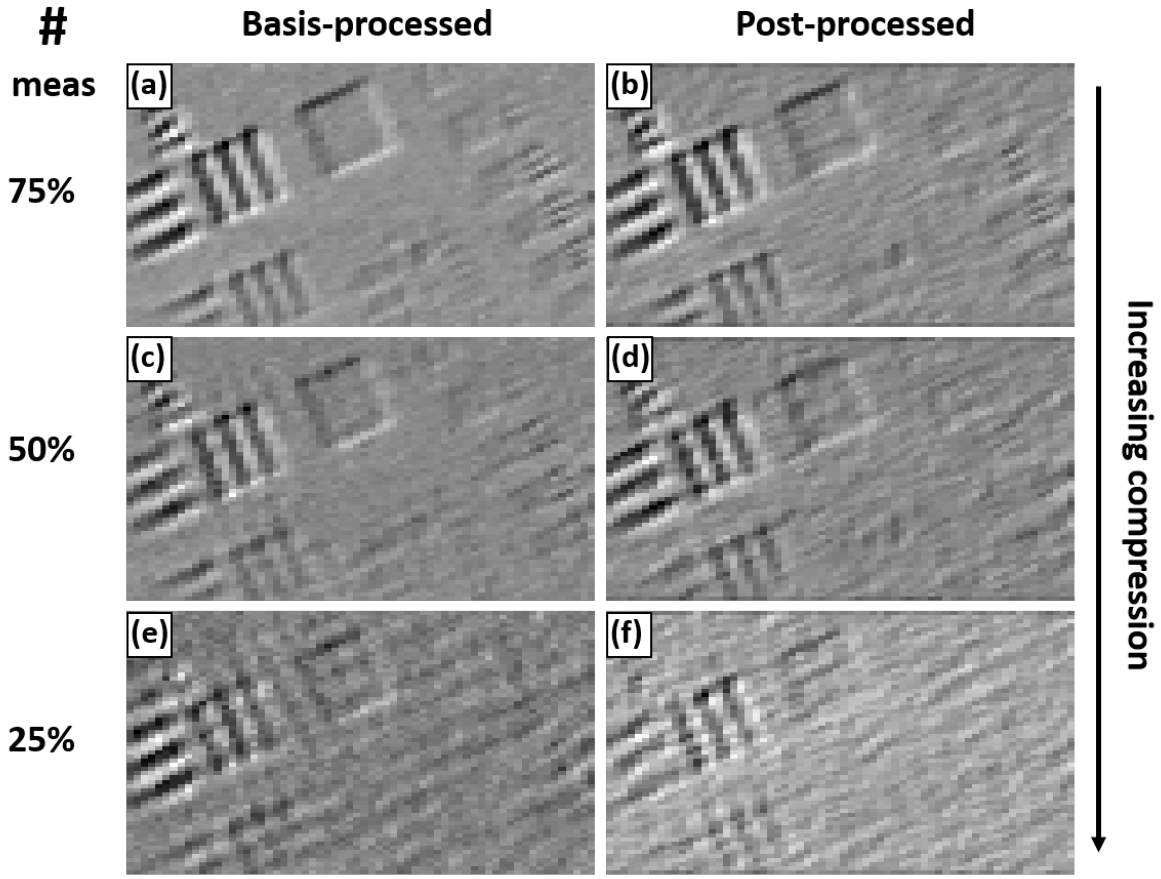


**Figure 6.17:** Comparison of image reconstructions from basis-processed and post-processed experimental measurements with the number of measurements used reducing from top to bottom. Percentage of the original $N^2$ measurements as indicated by # on the left. Compressed measurements have been reconstructed using sparsity in the DCT basis as a prior, following the ISTA approach (eq.6.19).

We have applied the ISTA reconstruction algorithm to subsets of the measurements taken previously, those shown in Fig.6.8(a) and (b), to compare the results for the basis and post-processing methods. The measurements were taken at $64 \times 64$ resolution with an exposure time of 220 ms and the power normalisation step included. The illumination basis $\boldsymbol{A}$, which is treated as the canonical basis in both cases is first transformed to the DCT basis, where the object is assumed to be sparse. A subset of the illumination patterns and measured coefficients, according to the desired level of compression, are taken at random and supplied to the ISTA algorithm as $\boldsymbol{A}$ and $\boldsymbol{y}$ as in eq.6.19.

The results are presented in Fig.6.17 for degrees of compression given by the % of measurements used as shown. For the majority of cases the optimal parameters (determined by trial and error) of the ISTA algorithm were the same, with $\delta = 0.75$ and $\alpha = 0.01 \times \max(\boldsymbol{y})$. The exceptions were in (e) $\alpha = 0.05 \times \max(\boldsymbol{y})$ and in (f) $\delta = 0.6$. The number of iterations was 1000 in all cases.

We can see in Fig.6.17 that the basis-processing method yields the better reconstructions in each case. Recall however that the SNRs for basis-processing were already higher in the absence of compressed sensing (Fig.**??**). A no conclusions can be drawn as to whether compressed sensing directly in the modified basis is advantageous, and this remains an open avenue for future work.

## 6.3 In the Context of Other Work

Improving the choice of illumination basis in response to prior or acquired knowledge of the object to be reconstructed is a principle that has been applied in several forms to computational imaging. In so-called principle component analysis the illumination wavefront is tailored to match the principle components of the object [144, 145]. In other adaptive imaging works the spatial resolution is varied in response to the detection of edges, which can be located via the wavelet transform [146, 147].

Less common is the modification of the illumination basis in response to the intended processing of the scene or 'task at hand' [148]. A recently proposed Learned Integrated Sensing Pipeline addresses the use of this additional prior knowledge in the context of machine learning and thus shares some similar principles to this work [148]. Finally we note some limited similarities with a microscanning imaging approach which combines shifted low resolution images to produce a high resolution image with improved SNR that has effectively been convolved with a smoothing kernel [149].

## 6.4 Conclusions

We have compared the images obtained using our novel ghost imaging technique, termed 'basis-processing', with the conventional post-processing method. We have presented the results for a simple edge-detection kernel as our processing operation. We find that if care is taken to make the comparison as fair as possible, the resulting SNRs from the two methods are equivalent. In this case a subtle advantage to the the basis-processing method is that the background noise remains 'white' with no potentially misleading spatial correlations introduced by the filtering step. Our simulations indicate that in the presence of constant background signal in the main detector (PD2) that is not properly removed, basis-processing is advantageous (Fig.6.9). We have also seen that if one doesn't normalise to fluctuations in the light source power with an additional bucket detector, basis-processing can produce higher SNRs (Fig.6.12).

The simple analytical treatment presented in eq.6.13 and eq.6.16 allows the prediction of the SNR changes for arbitrary filters. Finding a filter that universally yields higher SNRs when basis-processed requires that the noise increase from the filtering step, given by the square root of the filter energy (eq.6.12), outweighs the cost in measurement time associated with the need to increase the number of projected patterns. One such example would averaging filters where the filter energy is high but no additional measurements are needed, or simple Gaussian smoothing filters. However it does not seem likely these types of filters would be desirable in a ghost imaging experiment, where resolution is fought for at the expense of increased measurement duration.

Finally we have demonstrated that basis-processed measurements are compatible with compressed sensing techniques exploiting DCT basis sparsity, with no additional stand-out advantage or disadvantage apparent when compared to post-processing a reconstructed compressed measurement. One avenue for further investigation would be to explore alternative bases in which the basis-processed signal might exhibit greater sparsity, as these processed images may not share the characteristics of 'natural' images which make the DCT basis favourable.

# 7 Summary and Outlook

In the first results chapter of this thesis, Chapter 2, we investigated the lower size limits at which parabolic reflectors retain their ability to generate highly directional emission. Through finite element method simulations we determined that for the paraboloids in question the cut-off is just above one wavelength vertex - opening distance. Below this size the directed power is noticeably reduced with respect to the geometrical optics solution. Just above this size the directed power is comparable to and may even exceed that obtained in the geometrical optics regime.

The amount of power directed into a given solid angle is, for paraboloids of size comparable to the wavelength, a delicate interplay between changes in the local density of optical states and the increasing emission beam divergence as the antenna size is reduced. Thus the optimal position of an emitter to achieve the maximum directed power has been shown to deviate considerably from the macroscale case.

There are numerous ways in which the simulations could be extended. For the sake of generality we have chosen to model the reflector as a perfect mirror and not to include the refractive index of the luminescent material and the dielectric interfaces it may introduce. If one was simulating an experiment the choice of material for the mirror, index and distribution of the luminescent material would need to be accounted for. There is also the freedom to pursue paraboloids of differing aspect ratios and extents, as well as deviations from the ideal paraboloid, such as spherical reflectors, which may open new avenues for fabrication.

In Chapter 3 we explained the reasoning behind our choice of system for the desired directional phosphor component. This included discounting techniques that would likely restrict the thickness of a phosphor layer to of the order hundreds of nm. This was on the grounds of the difficulty this would pose for absorbing a high enough fraction of the pump light to produce white output. The chosen strategy was to pursue fluorescent optical fibres as a means of generating bright white sources with small cross-sectional areas.

The efficiency of fluorescent fibres for this purpose depends heavily on their ability to guide the fluorescence generated within them. This ability, quantified by the trapping fraction, was a major theme in the assessment of fluorescent fibres in Chapter 4. We saw that a simplified analytical treatment neglecting off-axis rays, which is commonplace for scintillating fibres, could underestimate the trapping fraction by as much as 25-35% of the total fluorescence generated.

The radial excitation position dependent trapping fraction was investigated through sim-

ulation and experiment and demonstrated for high quality monocrystalline YAG:$Ce^{3+}$ cylindrical rods to vary approximately half as much as predicted for pristine fibres (Fig.4.30). The demonstrated presence of this additional trapping power in cylindrical guides could facilitate mass manufacturing via conventional fibre drawing techniques, where otherwise the more challenging square cross-section may have been deemed necessary.

Prototype fluorescent fibres (rods) are measured to produce up to 150 lumens of white emission from their target output facet per Watt of blue laser pump, due to the contribution of off-axis rays. Simple improvements to the prototype design and implementation are suggested to increase the efficiency beyond the 200 lm/W target specified in Section 1.4.1.

The strength of the fibre approach is to allow a high fluorescence output within a small cross-sectional area, which enables more efficient collimation following the law of etendue. Consequentially, the prototypes are demonstrated without modifications to provide highly directional emission. The task of designing the optimum outcoupling and directing optic for the fibre end facet is left as a future application (and manufacturing feasibility) specific engineering challenge. One novel approach to this problem, the 'modified ball lens', has been investigated here through numerical simulations (Section 4.3.4).

We believe that fluorescent optical fibres are a strong candidate solution for meeting the full list of functional requirements outlined in Section 1.4.1. The next steps in assessing their suitability will depend upon manufacturing restrictions and application specifics, which are beyond the scope of this initial investigation.

In the final results chapter of this thesis, Chapter 6, we presented a novel technique in the field of ghost imaging, which is applicable to computational imaging in general. We have shown that by modifying the illumination basis such that the reconstructed image is already 'post-processed', the amplification of detector noise by the post-processing step can be avoided. In addition no undesirable spatial correlations in the noise are introduced in this basis-processing method. This technique can be applied to any operation which can be formatted as a matrix multiplication with the original image, which includes the many image filtering operations performed by convolution with a filter matrix.

The cost to this method is that the illumination basis becomes more complex, and may require a linear combination of binary patterns to implement. In the example we explore in detail, a simple edge-detection kernel, the reduction in noise achieved is in principle exactly cancelled by the need to perform twice as many measurements. However, we discuss several examples; such as when no additional detector is available to normalise to fluctuations in the source power, and when the measurements contain a non-zero background signal, in which basis-processing is shown to be advantageous.

Finally we tested the basis-processing and traditional post-processing approaches when combined with compressed sensing techniques. The relative visual quality of the compressed images for each method is in line with the uncompressed case and neither method shows a noticeably higher compatibility with compressed sensing. It is thought that with the right choice of sparse reconstruction basis for the basis-processing approach, measuring directly in the 'edge-modified' basis might prove favourable. The results of this chapter are the subject of ongoing research.

As the industry sponsors of this work Dyson have provided the following statement regarding the status of this work from their perspective:
*"Dyson received the work presented in this thesis and are satisfied by and excited about the results. The internal review was passed and the project closed. The Lighting Research team are currently working with the internal legal team to identify what details are outstanding to compile a full proposal for patenting. These outstanding pieces do not concern this work and are taking place internally."*

# 8 Publications and Intentions

- The work on wavelength-scale parabolic reflectors presented in Chapter 2 of this thesis has been published:
  Penketh, H., Bertolotti, J., & Barnes, W. L. (2019). Optimal position of an emitter in a wavelength-scale parabolic reflector. Applied Optics, 58(29), 7957.
  https://doi.org/10.1364/ao.58.007957
- The work in Chapter 4 concerning fluorescent fibres (and supporting Chapters 3 and 5) is of direct commercial interest to the project's sponsors and is being assessed for potential intellectual property as discussed in the previous section.
- The Ghost Imaging work in Chapter 6 is being prepared for publication.

# References

[1] Richard Feynman, Robert Leighton, and Matthew Sands. *The Feynman Lectures on Physics (New Millenium Edition): Volume 1*. New York: Basic Books, 2010.

[2] F. M. Clydesdale and E. M. Ahmed. “Colorimetry — methodology and applications”. In: *C R C Critical Reviews in Food Science and Nutrition* 10.3 (1978), pp. 243–301.

[3] Colour and Vision Research Laboratory UCL. *Judd Vos 2-deg XYZ Color Matching Functions*. `https://web.archive.org/web/20071227013005/http://www.cvrl.org/database/data/cmfs/ciexyzjv.txt`. Accessed: 2020-09-14.

[4] *Wikimedia Commons - Planckian Locus*. `https://commons.wikimedia.org/wiki/File:PlanckianLocus.png`. Accessed: 2020-09-14.

[5] E.Fred Schubert. *Light-Emitting Diodes*. 2nd. Cambridge University Press, 2006, p. 217.

[6] John Wilson and John Hawkes. *Optoelectronics - An Introduction*. Third. London: Prentice-Hall, 1998.

[7] Nicola Trivellin et al. “Laser-based lighting: Experimental analysis and perspectives”. In: *Materials* 10.10 (2017).

[8] Jaehee Cho, E. Fred Schubert, and Jong Kyu Kim. “Efficiency droop in light-emitting diodes: Challenges and counter measures”. In: *Laser and Photonics Reviews* 7.3 (2013), pp. 408–421.

[9] Giovanni Verzellesi et al. “Efficiency droop in InGaN/GaN blue light-emitting diodes: Physical mechanisms and remedies”. In: *Journal of Applied Physics* 114.7 (2013).

[10] Kristin A. Denault et al. “Efficient and stable laser-driven white lighting”. In: *AIP Advances* 3.7 (2013).

[11] Felipe Aquino, Wojciech M. Jadwisienczak, and Faiz Rahman. “Effect of laser speckle on light from laser diode-pumped phosphor-converted light sources”. In: *Applied Optics* 56.2 (2017), p. 278.

[12] Anant A Setlur. “Phosphors for LED-based Solid-State Lighting”. In: *Electrochemical Society Interface* 18.4 (2009), pp. 32–36.

[13] Michael R. Krames et al. “Status and future of high-power light-emitting diodes for solid-state lighting”. In: *IEEE/OSA Journal of Display Technology* 3.2 (2007), pp. 160–175.

[14] Hyun Kyoung Yang, Hyeon Mi Noh, and Jung Hyun Jeong. “Low temperature synthesis and luminescence investigations of YAG:Ce, Eu nanocomposite powder for warm white light-emitting diode”. In: *Solid State Sciences* 27 (2014), pp. 43–46.

[15] Ralph R. Jacobs, William F. Krupke, and Marvin J. Weber. “Measurement of excited-state-absorption loss for Ce3+ in Y 3Al5O12 and implications for tunable 5d-4f rare-earth lasers”. In: *Applied Physics Letters* 33.5 (1978), pp. 410–412.

[16] Nathan C. George, Kristin A. Denault, and Ram Seshadri. “Phosphors for Solid-State White Lighting”. In: *Annual Review of Materials Research* 43.1 (2013), pp. 481–501.

[17] W. J. Miniscalco, J. M. Pellegrino, and W. M. Yen. “Measurements of excited-state absorption in Ce3+: YAG”. In: *Journal of Applied Physics* 49.12 (1978), pp. 6109–6111.

[18] Stelian Arjoca et al. “Temperature dependence of Ce:YAG single-crystal phosphors for high-brightness white LEDs/LDs”. In: *Materials Research Express* 2.5 (2015), p. 055503.

[19] Crytur (company). *Crytur - Materials - YAG Ce*. `https://www.crytur.cz/materials/yagce/`. Accessed: 2020-09-14.

[20] Shunsuke Fujita et al. “YAG glass-ceramic phosphor for white LED (I): background and development”. In: *Fifth International Conference on Solid State Lighting* 5941.September 2005 (2005), p. 594111.

[21] Jimmy Wang et al. “High thermal stability of phosphor-converted white light-emitting diodes employing Ce:YAG-doped glass”. In: *IEEE Journal on Selected Topics in Quantum Electronics* 17.3 (2011), pp. 741–746.

[22] Jin Seok Lee et al. “Smart design to resolve spectral overlapping of phosphor-in-glass for high-powered remote-type white light-emitting devices”. In: *Optics Letters* 39.4 (2014), p. 762.

[23] Rui Zhang et al. “A new-generation color converter for high-power white LED: Transparent Ce3+: YAG phosphor-in-glass”. In: *Laser and Photonics Reviews* 8.1 (2014), pp. 158–164.

[24] Setsuhisa Tanabe et al. “YAG glass-ceramic phosphor for white LED (II): luminescence characteristics”. In: *Fifth International Conference on Solid State Lighting* 5941.September 2005 (2005), p. 594112.

[25] Daqin Chen et al. “Advances in transparent glass-ceramic phosphors for white light-emitting diodes-A review”. In: *Journal of the European Ceramic Society* 35.3 (2015), pp. 859–869.

[26] P. Siyushev et al. “Coherent properties of single rare-earth spin qubits”. In: *Nature Communications* 5.May (2014), pp. 1–6.

[27] Robert W. Boyd. *Radiometry and the Detection of Optical Radiation*. John Wiley and Sons, 1989.

[28] Y. S. Lin, R. S. Liu, and B.-M. Cheng. "Investigation of the Luminescent Properties of Tb[sup 3+]-Substituted YAG:Ce, Gd Phosphors". In: *Journal of The Electrochemical Society* 152.6 (2005), J41.

[29] Chang-yu Shen et al. "White LED Based on YAG : Ce , Gd Phosphor and CdSe-ZnS Core/Shell Quantum Dots". In: *IEEE Photonics Technology Letters* 22.12 (2010), pp. 884–886.

[30] Song Hu et al. "Transparent YAG:Ce ceramics for WLEDs with high CRI: Ce3+ concentration and sample thickness effects". In: *Ceramics International* 42.6 (2016), pp. 6935–6941.

[31] Shaowei Feng et al. "Spectrum regulation of YAG:Ce transparent ceramics with Pr, Cr doping for white light emitting diodes application". In: *Journal of the European Ceramic Society* 37.10 (2017), pp. 3403–3409.

[32] Qiping Du et al. "Massive red-shifting of Ce3+ emission by Mg2+ and Si4+ doping of YAG:Ce transparent ceramic phosphors". In: *Journal of Materials Chemistry C* 6.45 (2018), pp. 12200–12205.

[33] Miles J. Padgett and Robert W. Boyd. "An introduction to ghost imaging: Quantum and classical". In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 375.2099 (2017).

[34] Jeffrey H. Shapiro. "Computational ghost imaging". In: *Physical Review A - Atomic, Molecular, and Optical Physics* 78.6 (2008), pp. 1–4.

[35] Olympus Lifescience. *Introduction to Light Emitting Diodes*. `https://www.olympus-lifescience.com/en/microscope-resource/primer/lightandcolor/ledsintro/`. Accessed: 2020-09-14.

[36] C. A. Balanis. *Antenna Theory*. 3rd ed. Hoboken, New Jersey: Wiley, 2005.

[37] William L Barnes, Simon A R Horsley, and Willem L Vos. "Classical antennae, quantum emitters, and densities of optical states". In: *Journal of Optics* (2020).

[38] Sinan Karaveli and Rashid Zia. "Spectral tuning by selective enhancement of electric and magnetic dipole emission". In: *Physical Review Letters* 106.19 (2011), pp. 1–4.

[39] Mark Kasperczyk et al. "Excitation of magnetic dipole transitions at optical frequencies". In: *Physical Review Letters* 114.16 (2015), pp. 1–5.

[40] David J Griffiths. *Introduction to Electrodynamics*. 3rd. Pearson Education, Inc., 2008, pp. 444–449.

[41] Piers Andrew. “Molecular fluorescence near metallic interfaces Submitted by Piers Andrew in the Faculty of Science”. PhD thesis. University of Exeter, 1998.

[42] K. H. Drexhage. “Influence of a dielectric interface on fluorescence decay time”. In: *Journal of Luminescence* 1-2.C (1970), pp. 693–701.

[43] Enrico Fermi. “Quantum theory of radiation”. In: *Reviews of Modern Physics* 4.1 (1932), pp. 87–132.

[44] W. L. Barnes. “Fluorescence near interfaces: the role of photonic mode density”. In: *Journal of Modern Optics* 45.4 (1998), pp. 661–699.

[45] S. Haroche. “Cavity Quantum Electrodynamics”. In: *Fundamental Systems in Quantum Optics*. Amsterdam, The Netherlands: North-Holland, 1992. Chap. 13, pp. 768–940.

[46] Ho-seok Ee et al. “Design of polarization-selective light emitters using one-dimensional metal grating mirror”. In: *Optics Express* 19.2 (2011), pp. 1609–1616.

[47] M A Lieb and A J Meixner. “A high numerical aperture parabolic mirror as imaging device for confocal microscopy”. In: *Optics Express* 8.7 (2001), pp. 458–474.

[48] A. Drechsler et al. “Confocal microscopy with a high numerical aperture parabolic mirror”. In: *Optics Express* 9.12 (2001), p. 637.

[49] J. Stadler et al. “Tighter focusing with a parabolic mirror”. In: *Optics Letters* 33.7 (2008), pp. 681–683.

[50] W. P. Ambrose, Th Basché, and W. E. Moerner. “Detection and spectroscopy of single pentacene molecules in a p-terphenyl crystal by means of fluorescence excitation”. In: *The Journal of Chemical Physics* 95.10 (1991), pp. 7150–7163.

[51] L. Fleury et al. “Fluorescence spectra of single pentacene molecules in p-terphenyl at 1.7 K”. In: *Chemical Physics Letters* 236.1-2 (1995), pp. 87–95.

[52] Robert Maiwald et al. “Collecting more than half the fluorescence photons from a single ion”. In: *Physical Review A - Atomic, Molecular, and Optical Physics* 86.4 (2012), pp. 1–5.

[53] David T Schoen et al. “The Planar Parabolic Optical Antenna”. In: *Nano Letters* 13.1 (2013), pp. 188–193.

[54] Noel H. Wan et al. “Efficient Extraction of Light from a Nitrogen-Vacancy Center in a Diamond Parabolic Reflector”. In: *Nano Letters* 18.5 (2018), pp. 2787–2793.

[55] Sergii Morozov et al. “Metal - Dielectric Parabolic Antenna for Directing Single Photons”. In: *Nano Letters* 18 (2018), pp. 3060–3065.

[56] Sergii Morozov et al. “Objective-free excitation of quantum emitters with a laser-written micro parabolic mirror”. In: *APL Photonics* 5.7 (2020).

[57] COMSOL. *COMSOL Multiphysics Programming Reference Manual*. 5.4. COMSOL, Inc, 2018.

[58] J.A. Straton and I.J. Chu. “Diffraction Theory of Electromagnetic Waves”. In: *Physical Review* 56.1 (1939), p. 99.

[59] Comsol. *The RF Module User Guide*. 2012.

[60] D. K. Cheng. *Field and Wave Electromagentics*. 2nd ed. Boston: Addison-Wesley, 1991.

[61] R. K. Wangness. *Electromagnetic Fields*. 2nd ed. Hoboken, New Jersey: Wiley, 1986.

[62] Zlatko Živkovi et al. “Radiation pattern and impedance of a quarter wavelength monopole antenna above a finite ground plane”. In: *Software, Telecommunications and Computer Networks (SoftCOM), 2012 20th International Conference on Software, Telecommunications and Computer Networks (SoftCOM), 2012 20th International Conference* (2012), pp. 1–5.

[63] Harry Penketh, Jacopo Bertolotti, and William L. Barnes. “Optimal position of an emitter in a wavelength-scale parabolic reflector”. In: *Applied Optics* 58.29 (2019), p. 7957.

[64] Patrick Campbell and Martin A. Green. “Light trapping properties of pyramidally textured surfaces”. In: *Journal of Applied Physics* 62.1 (1987), pp. 243–249.

[65] P. Papet et al. “Pyramidal texturing of silicon solar cell with TMAH chemical anisotropic etching”. In: *Solar Energy Materials and Solar Cells* 90.15 (2006), pp. 2319–2328.

[66] Bennett W. Schneider et al. “Pyramidal surface textures for light trapping and antireflection in perovskite-on-silicon tandem solar cells”. In: *Optics Express* 22.S6 (2014), A1422.

[67] T. H. Taminiau et al. “Optical antennas direct single-molecule emission”. In: *Nature Photonics* 2.4 (2008), pp. 234–237.

[68] Alberto G Curto et al. “Unidirectional Emission of a Quantum Dot Coupled to a Nanoantenna”. In: *Science* 329.August (2010), pp. 930–933.

[69] L. Sortino et al. “Enhanced light-matter interaction in an atomically thin semiconductor coupled with dielectric nano-antennas”. In: *Nature Communications* 10.1 (2019).

[70] Gabriel Lozano et al. “Plasmonics for solid-state lighting: enhanced excitation and directional emission of highly efficient light sources”. In: *Light: Science & Applications* 2.5 (2013), e66.

[71] Gabriel Lozano et al. “Tailor-made directional emission in nanoimprinted plasmonic-based light-emitting devices.” In: *Nanoscale* July (2014), pp. 9223–9229.

[72] Vincenzo Giannini et al. “Plasmonic nanoantennas: Fundamentals and their use in controlling the radiative properties of nanoemitters”. In: *Chemical Reviews* 111.6 (2011), pp. 3888–3912.

[73] Aleksandr Vaskin et al. “Light-emitting metasurfaces”. In: *Nanophotonics* 8.7 (2019), pp. 1151–1198.

[74] Michele Gaio et al. “Modal Coupling of Single Photon Emitters Within Nanofiber Waveguides”. In: *ACS Nano* 10.6 (2016), pp. 6125–6130.

[75] Yasuhisa Inada et al. “Resonantly Enhanced Emission from a Luminescent Nanostructured Waveguide”. In: *Scientific Reports* 6 (2016), p. 34396.

[76] S. Murai et al. “Hybrid plasmonic-photonic modes in diffractive arrays of nanoparticles coupled to light-emitting optical waveguides”. In: *Optics Express* 21.4 (2013), p. 4250.

[77] Remote Phosphor and Light Source. “D ATA S H E E T Intematix ChromaLit”. In: ().

[78] Beer. “Bestimmung der Absorption des rothen Lichts in farbigen Flüssigkeiten”. In: *Annalen der Physik* 162.5 (1852), pp. 78–88.

[79] J.D. Jr. Ingle and S.R. Crouch. *Spectrochemical analysis*. Jan. 1988.

[80] Crytur (company). *Crytur homepage*. `https://www.crytur.cz/`. Accessed: 2020-09-14.

[81] W. L. Barnes, A. Dereux, and T. W. Ebbesen. “Surface plasmon subwavelength optics”. In: *Nature* 424.6950 (2003), pp. 824–30.

[82] S. Nishiura et al. “Properties of transparent Ce:YAG ceramic phosphors for white LED”. In: *Optical Materials* 33.5 (2011), pp. 688–691.

[83] Y Guyot ACP Rocha, LHC Andrade, SM Lima, AM Farias, AC Bento, ML Baesso and G Boulon. “Tunable color temperature of Ce3+/Eu2+, 3+ co- doped low silica aluminosilicate glasses for white lighting”. In: *Optics Express* 20.9 (2012), pp. 10034–10041.

[84] M. Clara Gonçalves, Luís F. Santos, and Rui M. Almeida. “Rare-earth-doped transparent glass ceramics”. In: *Comptes Rendus Chimie* 5.12 (2002), pp. 845–854.

[85] Stijn Roelandt et al. “Incoupling and outcoupling of light from a luminescent rod using a compound parabolic concentrator”. In: *Optical Engineering* 54.5 (2015), p. 055101.

[86] Jian Xu et al. “Investigation of laser-induced luminescence saturation in a single-crystal YAG:Ce phosphor: Towards unique architecture, high saturation threshold, and high-brightness laser-driven white lighting”. In: *Journal of Luminescence* 212.February (2019), pp. 279–285.

[87] Saint-Gobain Crystals. *YAG Material Datasheet*. `https://www.crystals.saint-gobain.com/sites/imdf.crystals.com/files/documents/yag-material-data-sheet_69775.pdf`. Accessed: 2020-09-14.

[88] Bradley S Jones. “Investigation of YAG:Ce Scintillating Fiber properties using scilicon photomultipliers”. In: *THESIS* (2011).

[89] C. Joram, G. Haefeli, and B. Leverington. “Scintillating Fibre Tracking at High Luminosity Colliders”. In: *Journal of Instrumentation* 10.08 (2015), pp. C08005–C08005.

[90] Jonathan D. Weiss. “Trapping efficiency of fluorescent optical fibers”. In: *Optical Engineering* 54.2 (2015), p. 027101.

[91] Dick K. G. de Boer et al. “LED-based projection source based on luminescent concentration”. In: *Optics, Photonics and Digital Technologies for Imaging Applications IV* 9896.November (2016), p. 989606.

[92] Max J. Riedl. *Optical Design: Applying the Fundamentals*. Bellingham Washington: SPIE, 2009.

[93] T. Vettenburg, S. A. R. Horsley, and J. Bertolotti. “Calculating coherent light-wave propagation in large heterogeneous media”. In: *Optics Express* 27.9 (2019), p. 11946.

[94] I. Koutromanos. *Fundamentals of Finite Element Analysis: Linear Finite Element Analysis*. and Sons, John Wiley, 2018.

[95] R. G. Newton. *Scattering Theory of Waves and Particles*. Springer, 2013.

[96] Gerwin Osnabrugge, Saroch Leedumrongwatthanakun, and Ivo M. Vellekoop. “A convergent Born series for solving the inhomogeneous Helmholtz equation in arbitrarily large media”. In: *Journal of Computational Physics* 322 (2016), pp. 113–124.

[97] Benjamin Krüger, Thomas Brenner, and Alwin Kienle. “Solution of the inhomogeneous Maxwell’s equations using a Born series”. In: *Optics Express* 25.21 (2017), p. 25165.

[98] Tom Vettenburg. *MacroMax - Microscopic Maxwell Solver, GitHub*. https://github.com/tttom/MacroMax. Accessed: 2020-09-14.

[99] M. J.A. de Dood et al. “Local optical density of states in Si02 spherical microcavities: Theory and experiment”. In: *Physical Review A - Atomic, Molecular, and Optical Physics* 64.3 (2001), p. 7.

[100] Intematix (company). *Datasheet NYAG4454-L*. http://www.intematix.com/uploads/phosphor-datasheets/NYAG4454-L.pdf. Accessed: 2020-09-14.

[101] Norland (company). *Norland Optical Adhesive 164*. https://www.norlandprod.com/adhesives/NOA164.html. Accessed: 2020-09-14.

[102] Digi-Key Electronics (company). *Intematix Chromalit Round*. https://www.digikey.com/en/product-highlight/i/intematix/chromalit-round. Accessed: 2020-09-14.

[103] Anastasiya Latynina et al. “Properties of Czochralski grown Ce,Gd:Y3Al5O12 single crystal for white light-emitting diode”. In: *Journal of Alloys and Compounds* 553 (2013), pp. 89–92.

[104] Ocean Optics (company). *OceanView Installation and Operation Manual*. https://www.oceaninsight.com/globalassets/catalog-blocks-and-images/manuals--instruction-old-logo/software/oceanviewio.pdf. Accessed: 2020-09-14.

[105] Joanne Zwinkels et al. “Mise en pratique for the definition of the candela and associated derived units for photometric and radiometric quantities in the International System of Units (SI)”. In: *Metrologia* 53.3 (2016), pp. 1–14.

[106] Colour and Vision Research Laboratory UCL. *Judd-Vos 2-deg photopic luminosity curve*. https://web.archive.org/web/20081228083025/http://www.cvrl.org/database/text/lum/vljv.htm. Accessed: 2020-09-14.

[107] *IES Method for Evaluating Light Source Color Rendition*. Illuminating Engineering Society, 2018.

[108] CIE. *CIE 15: Colorimetry - Technical Report*. 3rd. 2004.

[109] Yoshi Ohno. “Practical use and calculation of CCT and Duv”. In: *LEUKOS - Journal of Illuminating Engineering Society of North America* 10.1 (2014), pp. 47–55.

[110] Julio Chaves. *Introduction to Nonimaging Optics*. CRC Press, 2015.

[111] Edmund Optics. *Edmund Optics - Products - Compound Parabolic Concentrators (CPCs)*. https://www.edmundoptics.co.uk/f/compound-parabolic-concentrators-cpcs-3213/13944/. Accessed: 2020-09-14.

[112] Yen Sheng Lin et al. “High-luminance white-light point source using Ce,Sm: YAG double-clad crystal fiber”. In: *IEEE Photonics Technology Letters* 22.20 (2010), pp. 1494–1496.

[113] E. Mihóková et al. “Luminescence and scintillation properties of YAG:Ce single crystal and optical ceramics”. In: *Journal of Luminescence* 126.1 (2007), pp. 77–80.

[114] Crytur (company). *Crytur - Products - Scintillation Arrays*. https://www.crytur.cz/products/scintillation-arrays/. Accessed: 2020-09-14.

[115] Chun-nien Liu et al. “Fabrication and Characteristics of Ce-Doped Fiber for High-Resolution OCT Source”. In: 26.15 (2014), pp. 1499–1502.

[116] Dick K. G. de Boer, Dominique Bruls, and Henri Jagt. “High-brightness source based on luminescent concentration”. In: *Optics Express* 24.14 (2016), A1069.

[117] Juna Sathian et al. “Solid-state source of intense yellow light based on a Ce:YAG luminescent concentrator”. In: *Optics Express* 25.12 (2017), p. 13714.

[118] Hemant Kumar Raut et al. “Anti-reflective coatings: A critical, in-depth review”. In: *Energy and Environmental Science* 4.10 (2011), pp. 3779–3804.

[119] T. B. Pittman et al. “Optical imaging by means of two-photon quantum entanglement”. In: *Physical Review A* 52.5 (1995).

[120] Ryan S. Bennink, Sean J. Bentley, and Robert W. Boyd. “Two-Photon” Coincidence Imaging with a Classical Source”. In: *Physical Review Letters* 89.11 (2002), pp. 9–12.

[121] A. Gatti et al. “Ghost imaging with thermal light: Comparing entanglement and classical correlation”. In: *Physical Review Letters* 93.9 (2004), pp. 1–4.

[122] F. Ferri et al. “High-resolution ghost image and ghost diffraction experiments with thermal light”. In: *Physical Review Letters* 94.18 (2005), pp. 2–5.

[123] Alejandra Valencia et al. “Two-photon imaging with thermal light”. In: *Physical Review Letters* 94.6 (2005), pp. 1–4.

[124] Ron Meyers, Keith S. Deacon, and Yanhua Shih. “Ghost-imaging experiment by measuring reflected photons”. In: *Physical Review A - Atomic, Molecular, and Optical Physics* 77.4 (2008), pp. 1–4.

[125] Alba M. Paniagua-Diaz et al. “Blind ghost imaging”. In: *Optica* 6.4 (2019), p. 460.

[126] N. Fayard et al. “Intensity correlations between reflected and transmitted speckle patterns”. In: *Physical Review A - Atomic, Molecular, and Optical Physics* 92.3 (2015), pp. 1–9.

[127] I. Starshynov et al. “Non-Gaussian Correlations between Reflected and Transmitted Intensity Patterns Emerging from Opaque Disordered Media”. In: *Physical Review X* 8.2 (2018), p. 21041.

[128] Vialux (company). *Vialux - Products - SuperSpeed V-Modules*. `https://www.vialux.de/en/superspeed-specification.html`. Accessed: 2020-09-14.

[129] Alba M. Paniagua-Diaz. “Light in scattering media : active control and the exploration of intensity correlations”. PhD thesis. University of Exeter, 2018.

[130] Wai-Hon Lee. “Binary Synthetic Holograms”. In: *Applied Optics* 13.7 (1974), p. 1677.

[131] Catherine Forbes et al. *Statistical Distributions: Fourth Edition*. Fourth. Hoboken, New Jersey: John Wiley and Sons, 2010.

[132] J. W. Daniel et al. “Reorthogonalization and Stable Algorithms for Updating the Gram-Schmidt QR Factorization”. In: *Mathematics of Computation* 30.136 (1976), p. 772.

[133] David Ceddia and David M. Paganin. “Random-matrix bases, ghost imaging, and x-ray phase contrast computational ghost imaging”. In: *Physical Review A* 97.6 (2018).

[134] J.J. Sylvester. “ LX. Thoughts on inverse orthogonal matrices, simultaneous sign successions, and tessellated pavements in two or more colours, with applications to Newton’s rule, ornamental tile-work, and the theory of numbers ”. In: *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 34.232 (1867), pp. 461–475.

[135] Thorlabs (company). *Thorlabs - Products - Resolution Test Targets*. `https://www.thorlabs.com/newgrouppage9.cfm?objectgroup_id=4338`. Accessed: 2020-09-14.

[136] Wilhelm Burger and Mark J. Burge. *Digital Image Processing - An Algorithmic Introduction Using Java*. Ed. by David Gries and Fred B. Schneider. Second. London: Springer, 2008.

[137] Linda G. Shapiro and George C. Stockman. *Computer Vision*. New Jersey: Prentice-Hall, 2001, pp. 279–325.

[138] University of Wisconsin-Madison. *Public Domain Test Images*. `https://homepages.cae.wisc.edu/~ece533/images/boat.png`. Accessed: 2020-09-14.

[139] Aravind Vasudevan, Andrew Anderson, and David Gregg. "Parallel Multi Channel convolution using General Matrix Multiplication". In: *Proceedings of the International Conference on Application-Specific Systems, Architectures and Processors* (2017), pp. 19–24.

[140] *Convolution as Multiplication, GitHub*. `https://github.com/alisaaalehi/convolution_as_multiplication`. Accessed: 2020-09-14.

[141] Mathews Jacob and Michael Unser. "Design of steerable filters for feature detection using Canny-like criteria". In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 26.8 (2004), pp. 1007–1019.

[142] Michael Elad. *Sparse and Redundant Representations: From Theory to Applications in Signal and Image Processing*. New York: Springer, 2010.

[143] Amir Beck and Marc Teboulle. "A Fast Iterative Shrinkage-Thresholding Algorithm". In: *Society for Industrial and Applied Mathematics Journal on Imaging Sciences* 2.1 (2009), pp. 183–202.

[144] Mark A Neifeld and Premchandra Shankar. "Feature-specific imaging". In: *Applied Optics* 42.17 (2003), pp. 3379–3389.

[145] Min Liang et al. "Reconfigurable array design to realize principal component analysis (PCA)-based microwave compressive sensing imaging system". In: *IEEE Antennas and Wireless Propagation Letters* 14 (2015), pp. 1039–1042.

[146] Marc Aßmann and Manfred Bayer. "Compressive adaptive computational ghost imaging". In: *Scientific Reports* 3 (2013), pp. 1–5.

[147] David B. Phillips et al. "Adaptive foveated single-pixel imaging with dynamic supersampling". In: *Science Advances* 3.4 (2017), pp. 1–11.

[148] Philipp del Hougne et al. "Learned Integrated Sensing Pipeline: Reconfigurable Metasurface Transceivers as Trainable Physical Layer in an Artificial Neural Network". In: *Advanced Science* 7.3 (2020), pp. 1–8.

[149] Ming-Jie Sun et al. "Improving the signal-to-noise ratio of single-pixel imaging using digital microscanning". In: *Optics Express* 24.10 (2016), p. 10476.