



Welsh, Stephen Stewart (2017) Applications of single-pixel imaging. PhD thesis.

<http://theses.gla.ac.uk/8369/>

Copyright and moral rights for this work are retained by the author

A copy can be downloaded for personal non-commercial research or study, without prior permission or charge

This work cannot be reproduced or quoted extensively from without first obtaining permission in writing from the author

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the author

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given

Enlighten:Theses
<http://theses.gla.ac.uk/>
theses@gla.ac.uk

Applications of Single-Pixel Imaging

Stephen Stewart Welsh, B.Sc.

School of Physics and Astronomy,
University of Glasgow

Presented as a thesis for the degree of Ph.D.
in the University of Glasgow, University Avenue,

Glasgow G12 8QQ

© S. Welsh, 2014

August 14, 2017

Contents

Acknowledgements	xii
Abstract	xvi
Preface	xix
Summary	1
1 History of Single-Pixel Imaging	1
1.1 Introduction	1
1.2 First Single Pixel Imaging Systems	3
1.2.1 Single Pixel Scanning	3
1.2.2 Applications of Single-Pixel Scanning Systems	6
1.2.3 Alternate Scanning Systems	9
1.3 History of Ghost Imaging	9

1.3.1	Quantum Origin	9
1.3.2	Classical ghost imaging	10
1.3.3	Computational ghost imaging	12
1.3.4	Normalized ghost imaging	13
1.4	History of Single-Pixel Cameras	15
1.4.1	Single-pixel and ghost imaging cameras	15
1.4.2	Reconstruction Timescales	16
1.5	Summary	18
2	Experimental Methods	20
2.1	Spatial Light Modulation	20
2.2	Structured illumination and Coded Aperture methods	22
2.3	Pattern Design	23
2.3.1	Random	23
2.3.2	Orthogonal	29
2.4	Image Reconstruction	30
2.4.1	Pattern and signal	30
2.4.2	Generalised Signal and Noise Estimates	34

2.5	Compressed Sensing	45
2.5.1	Ill conditioned inverse reconstruction	45
2.5.2	Compressed Sensing Results	46
2.6	Summary	48
3	Multi-wavelength single-pixel imaging	51
3.1	Introduction to multi-wavelength imaging	51
3.2	Simultaneous and sequential illumination	56
3.3	Multi-wavelength single pixel imaging	61
3.3.1	RGB imaging	61
3.3.2	Fluorescence imaging	62
3.3.3	Shortwave infrared imaging	67
3.3.4	Fast multi-wavelength imaging	69
3.4	Discussion	78
4	Polarisation sensitive single-pixel imaging	80
4.1	Introduction to polarisation	80
4.1.1	Polarisation states	81
4.1.2	Jones calculus	85

4.1.3	Stokes' Parameters	88
4.2	Polarisation imaging with single-pixel cameras	93
4.2.1	Stokes' parameter image reconstruction	95
4.3	Experimental Method	98
4.4	Results	102
4.4.1	Angle of Polarization	105
4.4.2	Degree of Linear Polarization	106
4.5	Video rate single-pixel polarization imaging	109
4.6	Summary	110
5	Current Investigations with single-pixel detectors and future work	112
5.1	Long Range and Astronomical Observation	112
5.2	Design and Experimental Setup	114
5.2.1	Preliminary Results	116
5.3	Additional Future Applications	119
5.3.1	METAMATERIALS for Optics	119
5.3.2	Image enhancement via spectral image masking	123
5.4	The Future of Single Pixel imaging	125

6	Conclusions	127
6.1	Summary	127
6.2	Multichannel Imaging	128
6.3	Polarisation SPC imaging	129
6.3.1	Long Range SPC	130
6.4	SPC Future And Final Statement	131
A	Image Analysis	133
A.1	SNR Estimates	133
	Bibliography	147

List of Figures

1.1	Simplified diagram of a laser raster scanning system	5
1.2	Diagram depicting the basic operation of a confocal microscope	7
1.3	Diagram showing a simplified classical ghost imaging setup . . .	11
1.4	Diagram showing a simplified computational ghost imaging setup	13
1.5	Diagram showing a simplified normalized ghost imaging setup .	14
1.6	Comparison diagram highlighting the reversibility between ghost imaging and a single-pixel camera	16
2.1	A simple diagram showing the internal configuration of a LCOSLM chip	22
2.2	Comparison of simulated results of a single-pixel camera system using a standard iterative reconstruction algorithm	25
2.3	Signal to noise ratio estimates for each spot size when the number of measurements remained constant	27

2.4	Visual representation of the data ‘library’ accrued on a computer after acquisition is complete	32
2.5	A visual representation of the pattern weighting process	34
2.6	A visual representation of the reconstruction process for a 2×2 array and 1-2 pin head objects	39
2.7	Power law comparisons for testing pattern sets which contain either no spatial overlap or a random distribution, assuming a flat uniform object	41
2.8	Result of an iterative GI algorithm on a library of weighted pattern data taken from a standard normalized ghost imaging system for increasing number of patterns samples	44
2.9	Iterative TGI algorithm reconstructions for increasing measurement values	47
2.10	Iterative look at CGI algorithm reconstructions for increasing measurement values	48
3.1	Experimental setup of light commander based computational imaging system	53
3.2	Diagram of the pattern projection sequence constrained within a 24 bit-plane projection cycle	56
3.3	A visual representation of the RGBMYC patterns produced when attempting sequential multi-wavelength imaging	59

3.4	Reconstructions from data obtained from three spectrally filtered PD's for RGB colour channels	62
3.5	Jablonski diagram detailing the absorbance, the non-radiative portion of the electron decay and the fluorescence as the electron returns to its ground state	63
3.6	Results of florescence imaging on the Rubik's cube test object with 265k iterations per colour channel	66
3.7	Non-visible result of the Rubik's cube test object with 850nm illumination	67
3.8	Experimental setup of light commander based computational imaging system	70
3.9	Graph demonstrating the linear relationship between the acquisition time in seconds verses the total pixel count of the final reconstructed image	71
3.10	Full-colour iterative reconstruction of a 3D scene	72
3.11	Comparison of coloured image reconstructions	74
3.12	Compressive vs Iterative algorithm performance comparison	75
3.13	SSI calculations for 2 iterative and 2 compressive results	77
4.1	Planewave propagation for different polarisation states	81
4.2	Figure describing Left-Handed Elliptical Polarisation and Right-Handed Elliptical Polarisation	84

4.3	The Polarisation Ellipse and Poincare sphere for the Stokes' parameters	89
4.4	SNR of the detectors in the polarisation experiment	96
4.5	Experimental SPC systems for characterising linear polarisation states of diffuse and specular reflections	99
4.6	Stokes' parameter images S_0 , S_1 and S_2 reconstructed in colour (RGB), for which the object is a stressed clear plastic sheet which has had a triangular portion removed	103
4.7	Iterative reconstructions of the first three Stokes parameter images S_0 , S_1 and S_2 used in the RGB combination images	104
4.8	Reconstructed images of a corner cube obtained by the polarization sensitive single pixel camera prototype for various incident linear polarization illumination	105
4.9	Stokes parameter image reconstructions for a small cactus plant	107
4.10	A selection of S_1 images acquired at different intervals for scenes exhibiting dynamic behaviour	109
5.1	Original design of the attachment component for mounting a single pixel imaging system onto a telescope	114
5.2	Updated and final design model for the telescope single pixel camera mount	115
5.3	Photographs of the DMD and 3D Printed mounting unit	117

5.4	Initial reconstruction result from a single pixel camera mounted on a telescope	118
5.5	A series of SWIR images captured with the single pixel camera system at near video rate reconstruction rates of approx. $10Hz$	119
5.6	METATOY viusual summary of principle	120
5.7	Simulated METATOY result showing deformation of 3D objects	121
5.8	Simulated METATOY result for view rotation	122
5.9	The simultaneous multi-wavelength single pixel imaging system modified for image enhancement by spectral image masking . . .	124
A.1	Signal to noise ratio estimation made by thresholding signal and background ROI in the RGB image data	134
A.2	Signal to noise ratio estimates for the fluroescent SPC results . .	135
A.3	Signal to noise ratio estimates for a selection of varying M value reconstructions featuring both iterative and compressive results	136
A.4	Signal to noise ratio estimates for the RGB and RED color channels of the Perspex sheet	137
A.5	Signal to noise ratio estimates for the 0 degree linear polarisation illumination of the cactus	138
A.6	Signal to noise ratio estimates for the 90 degree linear polarisation illumination of the cactus	139

List of Tables

2.1	Generalised Signal, Noise and S:N Estimate	35
4.1	Jones Matrices	88
4.2	SNR Estimates for Fig. 4.6 and Fig.4.7 at 0°	105
4.3	SNR Estimates for Fig. 4.9	108

Acknowledgements

I would like to thank my supervisors Miles Padgett and Johannes Courtial for the many guided discussions and suggestions which led to new experimental proposals and subsequent publications. I would also like to thank the tireless efforts of my colleagues Dr Matt Edgar and Dr Baoqing Sun for their continued support during the last three and half years without which the systems described within would not be where they are today. Additionally, thanks goes to a now ex member of the Optics group at Glasgow, Dr Richard Bowman. who provided his advanced computer programming knowledge freely resulting in rapid advancement in underlying software essential for these systems to operate successfully.

“The quickest way to make money at photography is to sell your camera.”

- Yousuf Karsh, photographer for Winston Churchill

Abstract

In this body of work, several single-pixel imaging applications are presented, based on structured light manipulation via a Digital Micromirror Device (DMD) and a single element photodetector (PD). This is commonly known as computational single-pixel imaging, and is achieved by using the measurements made by the PD to weight a series of projected structured light-fields. This indicates the strength of correlation between each light-field, and some object or scene placed in its propagation path. After many iterations the ensemble average of the weighted structured light-field converges to the object.

Historically, computational single-pixel imaging has suffered from long image acquisition times and low resolution. Inhibiting the ability of physical systems from competing with conventional imaging in any form. Advances in computer and DMD technology have opened new avenues of research for this novel imaging technique. These advances have been utilised in this work by creating fast acquisition demonstrator systems, which have real world applications, such as multi-wavelength, polarisation, and long-range imaging.

Several PDs were added to allow for simultaneous measurement of multiple images in the desired application. For multi-wavelength, RGB and white light illumination was spectrally filtered on three detectors to create full-colour images. While conversely the same multi-detector approach allowed for simulta-

neous measurement of orthogonal linear polarisation states essential to Stokes' parameter image reconstruction. Differential projection of the structured light-fields further allowed for the single-pixel camera to compensate from some sources of real world noise, such as background illumination.

This work demonstrates an evolution of the single-pixel camera. From a system capable of only imaging simple, binary transmissive objects twice per hour and constrained to an optical bench, to a semi portable camera, capable of multiple frames per second 2D reconstructions of 3D scenes over a range of 20 kilometres. These improvements in capability cement the idea that the single-pixel camera is now a viable alternate imaging technology.

Preface

Chapter 1: History of Single-Pixel Imaging

This introduction chapter will chronicle the very basics and origins of single element imaging, outlining the differences between projection and recording devices which use full field capture and those systems which sample only a portion of the object at a time. It will describe applications of these early single element ‘scanning’ systems before detailing the history of the primary type of single pixel camera described in this work. There will be a progression of systems relating to the chronological order of development of this technology and the advances each system has brought.

Chapter 2: Experimental Methods

Here the practical implementation and associated challenges with building and operating various single pixel imaging systems will be discussed along with the basic processes inherent to all single pixel imaging systems for image reconstruction. A discussion of pattern implementation will be undertaken as pattern selection and projection or encoding is the most important aspect of creating a library of data which can be reconstructed into a final image, with the correct encoding and corresponding measurement there could be no image. A brief introduction to compressed sensing in the context of single pixel camera systems and image reconstruction methods will be given along with simulated results demonstrating the various previously mentioned algorithms in action.

Chapter 3: Multi-wavelength single-pixel imaging

This chapter will focus on the first main application of a single pixel imaging system, multi-wavelength imaging producing full-colour images. This implementation of a single pixel imaging system was the first to use a digital micromirror array where the patterns are projected at high speed (approx 1kHz) onto a real 3D scene in order to create 2D images in full colour. There are multiple approaches to multi-wavelength imaging from multi-coloured pattern projection with spatially filtered detectors, a simultaneous full-colour imaging system, to sequential reconstruction using mono-chrome patterns with a single detector, building up the multi-wavelength image in a non temporally changing object. Real results from the corresponding publication are given including uncompressed and compressed results. An addendum to the work which explored basic fluorescence imaging will also be presented.

Chapter 4: Polarisation sensitive single-pixel imaging

This chapter outlines a second major application of a similar system to the one presented previously. However, this system has been made polarisation sensitive and is capable of reconstructing polarisation images corresponding to the Stoke's parameter images, along with images detailing Angle of Polarisation and Degree of polarisation. Furthermore, the system has been improved to operate in a superfast mode (approx 24kHz) which allow near video rate real time reconstruction of polarisation sensitive images.

Chapter 5: Current investigations with single-pixel detectors and future work

Current and future work is presented focusing on two distinct avenues of research, Firstly, an investigation was started in which long range single pixel imaging through a telescopic device was attempted. Investigations into visible and Infrared comparisons were made to test variations image quality due to

differing wavelength propagation through the atmosphere with the intention that this work lead to a suitable camera system capable of imaging distant and potentially astronomical objects in the non-visible. A diverging investigation into optical metamaterials is demonstrated along with a proposed future single-pixel investigation using wavelength ‘masking to achieve improved image quality.

Chapter 6 Conclusions

The conclusion will discuss findings from previous applications and cover the main reasons that a completely different imaging approach can still be applied in each situation and the advantages a single pixel system has over a traditional camera based on a detector array.

Summary

Camera devices and imaging systems have allowed unique insight and measurement of a vast array of the physical properties of light. Ingenuity in camera technology has allowed the capture of images from a vast portion of the electromagnetic spectrum far beyond what the human eye can acquire alone, albeit at a generally much increased cost in non-visible light regions. Alternate imaging systems do exist to the well-known standard of imaging based on digital detector arrays, including the overlapping fields of ghost imaging and single pixel cameras. The basic principles behind the classical interpretation of such imaging systems are outlined in this work as well as multiple demonstrations of this technology in different imaging applications. Advancements in image quality and acquisition rate can be observed as work progresses from initial simplistic applications to more complex versions of the single pixel camera system.

Single pixel cameras now offer a portable device which is applicable to a wide band of the electromagnetic spectrum, without the large variations in cost normally associated with cameras out with the visible portion of the spectrum. With Digital Micromirror Device (DMD) based systems generally covering $0.2 - 2.5\mu m$ regions, this system has great potential as a truly hyper spectral camera system. As a new alternate method of imaging, it can help build confidence in the technique if it can be shown that the single-pixel approach

is robust enough to be applied across a wide variety of well-known imaging applications. It is the goal of this body of work to demonstrate that various incarnations of a prototype single-pixel camera system can have real world applications, replicating the results of standard imaging systems while still having the inherent advantages of a single-pixel approach in relation to primarily the extended operational range associated with a single detector or series of single detectors verses and coated DMD with a broadband spectral response, including Short Wave Infra Red (SWIR).

Early demonstrations of what was known as ‘ghost imaging’ were demonstrated utilising parametric down conversion to produce entangled photons and subsequent image reconstruction from two detectors, one with no spatial resolution observing the photons interacting with an object, and one with spatial resolution but without any object present. However, this was not seen as competitor to standard imaging systems, but rather a quirk of the underlying quantum nature of the paired photons needed to produce the image reconstruction effect. This was due in part to the fact that acquisition times and image quality were no comparison to film or digital detector array devices of the time. However, even after a classical based method was shown based on a random pattern projection, technological limitations have been the primary reason why there has not been a much wider adoption of this technology in the past decade. For high resolution, high speed image reconstruction in a classical based system, the computational load becomes large and the underlying analogue or digital spatial modulation systems for light fields were restricted in the speed at which they could be modulated. Such were these limitations that only a few years ago publications in this field consisted of only binary transmissive object reconstruction, as opposed to 2D reconstruction of real 3D scenes, with resolutions typically on the order of 1000-4000 pixels per image and reconstruction times exceeding 20 minutes set primarily by restrictive projection timescales of

less than 25Hz. Additionally, since these systems were usually based on a laser for illumination, published work focused primarily on single wavelength visible light reconstruction where visible spectrum lasers were more readily available than non-visible counterparts.

During this time, these early demonstrations were primarily undertaken not to show a competitive imaging system, but rather, to further the debate of whether ghost imaging was solely a quantum or classical phenomenon. In the past few years, computer systems and spatial light modulation devices have improved, allowing us to apply this technique as a serious alternate candidate for non-visible imaging. We can now utilise broadband light sources which extend into the Short-wave Infrared region (SWIR) for multiple wavelength reconstruction. Using three colour illumination (Red, Green, and Blue (RGB)) we can construct ‘full-colour’ images of reflective objects rather than the binary transmissive test targets of previous demonstrations of this technique. The target wavelength(s) can be selected based on the need of the observation, simply by adjusting either illumination wavelength or the spectral filter of the broadband source. The single-pixel camera can also demonstrate sensitivity to additional properties of light beyond wavelength, such as polarisation. Given the appropriate setup, it has been demonstrated that polarisation information such as Stoke’s parameters, Degree of Polarisation (DoP) and Angle of Polarisation (AoP) can be acquired using only single pixel detectors. The images reconstructed with a polarisation sensitive single pixel camera are similar to the information obtained with a Stoke’s imaging CCD based setup, however there are some inherent benefits to using the single pixel method over the CCD based system such as perfect pixel alignment and immunity to depolarizing components in the imaging path which could interfere with received polarisation states. In traditional polarisation camera systems, multiple CCD cameras call for near perfect pixel alignment to measure simultaneous multi-

ple polarisation state images. The inherent properties of a single pixel camera require no such alignment within the single element detectors.

Additionally, the system has been adapted to be compatible with a commercially available telescope. Employing the use of simultaneous visible and SWIR detectors real time images of objects up to 20km away have been successfully acquired. The aim of this application is to demonstrate the benefit of having a system capable of multi-wavelength reconstruction outside the visible spectrum when applied to imaging over large distances. While atmospheric turbulence is a dynamic effect, approximately changing at a rate of $50Hz$, longer wavelength light rays will be scattered less during propagation over long distances. It is therefore possible to see an increase in the quality of image obtained simply by observing in a longer wavelength band. This is also true when applied to astronomical objects. In astronomy, a celestial object's spectrum is the primary source of information from elemental composition, to distance due to the Doppler effect. Having access to simultaneous imaging at many wavelengths has the potential to offer an interesting alternative for spectrum imaging over a much larger portion of the EM spectrum. This can be specifically targeted at atmospheric 'windows' of high transmissivity in the SWIR region allowing access to useful emission lines on a smaller equipment scale and budget than is currently possible. Finally, a divergent optical metamaterial investigation is demonstrated which may have future applications in single-pixel systems alongside a proposed experiment based on wavelength 'masking' to improve image quality in high transmissive images.

Chapter 1

History of Single-Pixel Imaging

1.1 Introduction

The fundamental goal of any imaging device is the detection and spatial mapping of a light field, with this representation compromised to some resolution set primarily by the detection medium. Modern digital imaging devices use a detector array suited to a specific waveband, positioned behind an appropriate optic and placed in the imaging plane of that optic to achieve spatial mapping and detection of incoming light fields. Where these older analogue technologies utilise photographic film in place of a digital array. However, other alternate imaging methods exist, for example, the overlapping fields of single-pixel cameras, ghost imaging and computational imaging. These alternate imaging techniques all rely on unique spatial correlation methods, sharing certain aspects of their image reconstruction techniques, and use a single-pixel detector. Over the past two decades extensive work in these alternate imaging systems has approached a level in which real world applications and potentially competitive imaging systems can be demonstrated. This is primarily due to

the development of faster, more powerful computers and improved underlying spatial light mapping technologies such as Spatial Light Modulators (SLM) and Digital Micro-Mirror Devices (DMD). This work describes these systems as applied to real world imaging applications, demonstrating the ability for this alternate imaging approach to successfully meet the requirements of a chosen application and explore what possible benefits these systems may have over more traditional imaging devices.

Every traditional camera system has a finite operational bandwidth. That is, it will only respond to a small portion of the electromagnetic (EM) spectrum for which that camera system was designed. This is primarily due to the bandwidth response of the detection medium. Therefore, a highly varied and generally expensive developmental approach is needed in making cameras responsive to different portions of the EM spectrum. This is especially the case when moving out with the visible ($400 - 700nm$), approximately \$5 per visible megapixel, and \$50000 per megapixel in the $1\mu m - 2.5\mu m$ range.

One of the largest potential advantages offered by this single-pixel correlation approach to imaging lies in the relatively large portions of the electromagnetic spectrum in which a single device can be capable of measuring. Traditional imaging systems are completely distinct from each other and generally cannot be switched from one bandwidth to another. Thus, a separate camera is needed for every distinct portion of the EM spectrum to be measured. In addition, not all portions of the EM spectrum have a corresponding detector with spatial resolution due to either the difficulty or practicality of implementing such a device. The advantage of using a device which can produce full 2D representations of an object or scene from a single-pixel is that detectors utilising the appropriate medium for each corresponding portion of the EM spectrum can be added to the same basic system. This comes with the assumption that system

can still perform the necessary spatial correlation required for image reconstruction. Therefore, the goal of this alternate approach is to produce a device capable of many simultaneous real-world applications due to its potential to be a fully hyper-spectral imaging system. While that goal has not been fully realised yet, this work demonstrates a variety of prototype systems. As each improvement over the previous incarnation of the prototype is demonstrated, this technology becomes closer to a truly competitive and viable alternative to a traditional imaging system. This work features an exploration of SPC capabilities, including simultaneous multi-channel reconstruction, fundamental light properties, and the robustness of the system at long and short ranges. Each implementation of the SPC system demonstrates the adaptable applicability of the technique, where the performance of traditional detector array based camera systems is used as a metric to determine the advantages, if any, of a SPC approach.

1.2 First Single Pixel Imaging Systems

1.2.1 Single Pixel Scanning

There is a close common history between the development of camera and projection systems[1, 2]. Original imaging and projection technologies took the form of a pinhole camera. This remarkably simple device provided both functions, filtering the image of a bright light field and subsequently projecting that image onto a screen, although each image had no way to be stored and thus it could only ever show a live representation of the bright incoming light field. Developments into photographic film allowed the first true camera, the camera obscura, to capture the projected image as a photograph. This was achieved

by placing a plate covered with photosensitive powder on the main table of the camera obscura, leading to the first ‘photographs’. Therefore, function of the pinhole camera was a simple image projector, and it was only the inclusion of the photosensitive screen to ‘store’ the projected image that altered its function to that of a ‘camera’. Essentially all devices capable of full field image projection become cameras if or only when an appropriate method to store the full field is present. Measuring or displaying a single spatial component of the light-field field sequentially over time is a fundamentally different approach to imaging or projection. It would be with the development of the first television screens that would bring a true ‘single element’ projection device and with it, the first device which could be clearly split into the category of a projector or a camera. While the underlying principles would remain the same with this ‘non full-field approach to projection, each device had to use different components and arrangements in order to be put into the categories of camera or projector.

The Cathode Ray-Tube (CRT) combined an electron gun and phosphor front screen to create a device capable of producing a raster image of any encoded pattern sent to the device. This was the first analogue single pixel projection device where magnetic deflection of the electrons caused fluorescence in the display screen at appropriate coordinates, to create a light intensity representative of the encoded image. A complete still image is built up in a sequence of (usually) horizontal lines, electrons creating a single row of illumination before being deflected to the original starting position with a small vertical decrease in the electron path’s deflection to strike the front screen in the correct position of the next row. The process could be completed fast enough and repeated at a rate whereby the apparent effect of motion can be observed frame to frame. We define such a frame as the combination of the illumination effect from the first projection, to the projection that completes the image, in the case of the CRT specifically, a new frame would cause the deflection path to reset to the

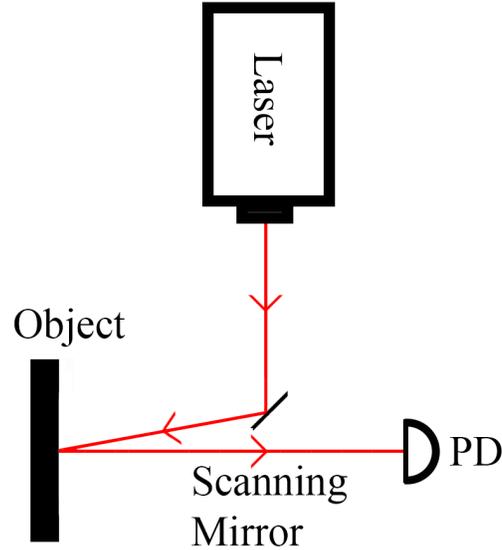


Figure 1.1: A simplified diagram of a laser raster scanning system. Here the laser illumination is directed by a scanning mirror which moves over the surface of the object in a series of horizontal lines. The photodetector (PD) measures the returned intensity for every ‘scanned’ coordinate in sequence, where each value of the detector then becomes the intensity of a pixel in the final reconstructed image.

first coordinate.

A single-pixel raster scanning camera is an optical reversal of the CRT device with altered components for detection rather than projection. A simplified laser raster scanning imaging system is shown in Fig. 1.1, where the electron gun has been replaced with a suitable photo sensitive detector(PD). The PD in this case has an active area of around 2cm^2 allowing ample reflected light collection, without collecting all the back scattered light. In an active raster scanning system, illumination from the laser is directed by a scanning mirror onto a series of horizontally arranged coordinates. The PD then measures the reflected light intensity, which is then the value used to represent that pixel

in a final reconstructed image, made from measuring the reflected signal from all coordinates and arranging the values appropriately. Since each coordinate must be measured to create a complete image, this sets a very strict limit on the number of measurements required for image reconstruction where the number of measurements, M , must equal the number of pixels in the final image, N . Additionally, in a raster scanning system if M exceeds N , this is simply oversampling of the object. This can only be useful if repeated scanning of coordinates generated average values for the measurements, which could reduce noise from the detector by finding an average value rather than a single value. However, for each additional set of identical coordinates which are detected, the time to measure all coordinates is multiplied by the additional detection cycles which becomes highly inefficient. The $M = N$ hard measurement number limit may seem trivial when applied only in the context of this type of imaging system, however it becomes an important standard of comparison for other types of single-pixel imaging systems described later in this work.

1.2.2 Applications of Single-Pixel Scanning Systems

Single element scanning systems are widely used in an extensive array of fields and applications, with one such field that of microscopy. There are several incarnations of scanning microscopes, using both photons, in the case of Confocal Laser scanning microscopy (CLSM) and Near-field scanning optical microscopes (NSOM), and electrons used in a scanning electron microscope. For microscopes using light as their illumination method, the single element scanning component of the microscope provides improved contrast and in the case of a NSOM microscope, the ability to overcome the diffraction limit from a given illumination wavelength by exploiting properties of the light waves in the near field.

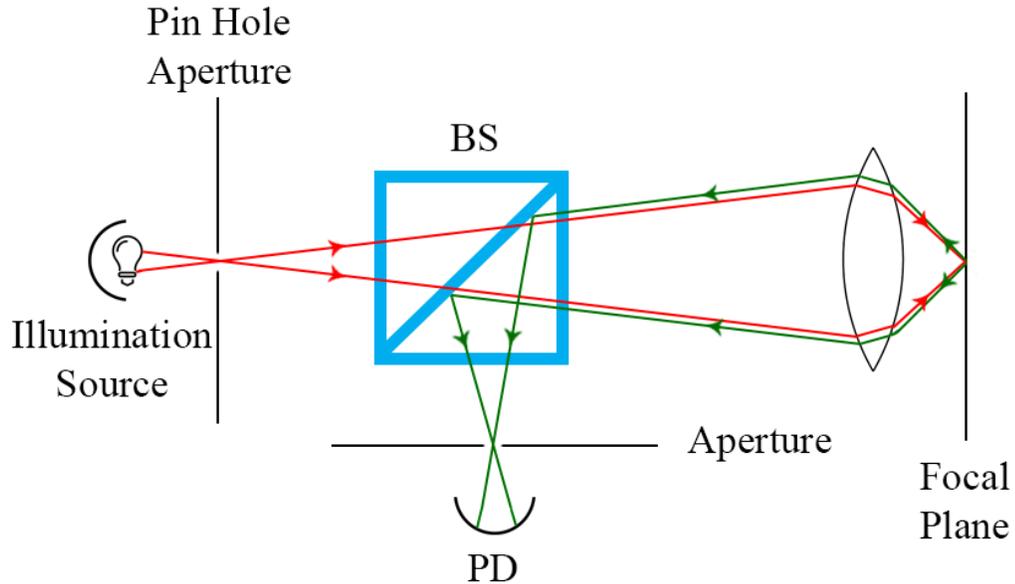


Figure 1.2: A diagram depicting the basic operation of a confocal microscope. Here the inclusion of the pinhole aperture and illumination source which can be scanned across the focal plane via one or two servo controlled mirrors serve to only illuminate objects located on the focal plane, eliminating background light from anything outside this plane, increasing contrast of the sample. The beam splitter (BS) only allows light after reflection from the focal plane to be sent down the detection path to the Photo Detector (PD).

Figure 1.2 demonstrates the well-known basic operation of a CLSM [3, 4]. The image reconstruction process works exactly like a standard raster scanning camera, where each pixel is arranged into its corresponding row and the image is built from the correct arrangement of these rows. However, in this system a 3D image of the object can be constructed if an extended object is moved through the focal plane. This corresponds to the rate of the scanning mirrors and detector, creating 2D slices which can be further arranged into a 3D representation. Without the use of the pin hole aperture, a detector array with high

spatial resolution could capture the entire image in one measurement however, light from all layers outside the focal plane in extended objects, or background out of focus light in 2D objects will drastically reduce the quality and contrast of the final image. The advantage of introducing a scanning single pixel approach to this type of visible microscopy is therefore a direct enhancement to image quality over full field measurement microscopes.

Another application of single-pixel imaging systems pertains to gas imaging [5]. All molecular gases have an electromagnetic (EM) spectrum absorption profile distinct to each specific gas. Once this absorption profile is known, illuminating a molecular gas cloud with multiple wavelengths of light can allow detection of an otherwise transparent and thus invisible substance. A simple gas spectrometer works by collecting a light field which has propagated through the gas. A prism splits this incoming light field into its constituent wavelengths which will have been attenuated by the gas. Since the spatial position of light after propagation through a prism is dependent on its wavelength, and not the original position from within the molecular gas, it is important to use a mask and scanning mirror system to only capture light from a known specific coordinate in the gas cloud. This allows imaging of the gas to be produced from raster scanning the results of the spectrometer readings, where each reading, and relative values from the attenuated wavelengths, give information relating to the density and concentration of the gas. Using a scanner laser to illuminate the gas at a specific point and depth is useful for precise image reconstruction. A thin depth of gas is selected so that this is an approximate of a 2D plane of gas, even though there is some small 3D integrated depth, this is effectively flattened in the final reconstruction. The laser would then alternate between two illumination wavelengths, one wavelength would be chosen as the wavelength of a strong absorption line in the target gas, and the second would correspond to that absorption line with a small offset. The differential of the

two results in this case would produce the relative value for the intensity to be used in the reconstruction of the gas image.

1.2.3 Alternate Scanning Systems

A variant on the single-pixel raster scanning system is the line scanning camera. This works by having a row of pixels, rather than a single pixel, where the illumination source creates a focused thin line of illumination which either scans over the object from top to bottom, or more commonly, the object is passed through the light column. These cameras are still utilised for various industry applications. In high speed applications, uniformly illuminating a 2D area for long periods of time can be challenging, whereas illuminating a column of light and passing the object over this column at high speed can offer high resolution, fast imaging for products or equipment. This is useful where such objects can be extended in size or diameter and illumination is only required in a fixed area.

1.3 History of Ghost Imaging

1.3.1 Quantum Origin

Initially, ghost imaging relied on the use of two correlated light fields and two photo-detectors to produce an image[6]; one detector with no spatial resolution used to collect the light field which had interacted with a binary transmissive object, and a second detector with high spatial resolution which was only activated upon registration of a detection by the first detector. However, the path in which the detector with spatial resolution is placed had no object. Neither

detector alone is therefore capable of imaging the object, only a combined coincident measurement can produce an image from such a system.

These early demonstrations of the ghost imaging technique utilized the apparent entanglement arising from spontaneous parametric down-conversion to produce entangled photons, an approach often referred to as quantum ghost imaging[7, 8]. In these early experiments, only the spatial entanglement of paired photons traveling down different paths could explain the image formed in the path with no object. The term ‘ghost’ was coined in this instance to refer to absence of an object in the path in which the image is captured. This early approach was a slow imaging technique on the scale of several minutes or more to ‘build up’ a reconstruction from individual detections of paired photons and required the use of binary transmissive objects only. At this stage in the development of an alternate imaging system, this approach was considered as simply an observation of the underlying quantum effect and not as a competitive imaging system to camera film or digital detector arrays. Many of the discussions in this new emerging field were focused solely on why this imaging technique worked, and if it was an inherent quantum effect [9, 10]

1.3.2 Classical ghost imaging

The presence of entanglement led many to interpret ghost imaging as fundamentally quantum behaviour. However, there have since been demonstrations using a pseudothermal light source[11, 12, 13, 14, 15], whereby laser light is propagated through a ground glass diffuser to produce a speckle field, after which a beam splitter makes a correlated copy of the field as shown in Fig. 1.3, a technique commonly termed classical ghost imaging. For both classical and quantum ghost imaging approaches, the field-of-view, spatial resolution,

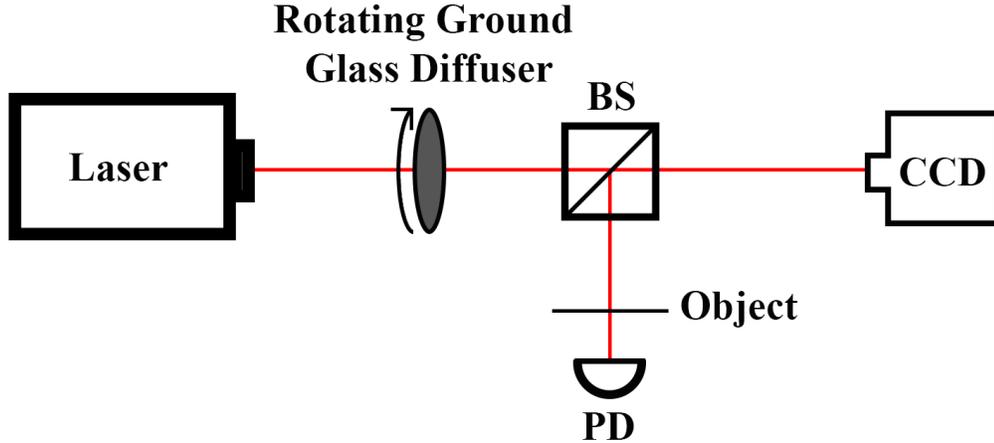


Figure 1.3: *Diagram showing a simplified classical ghost imaging setup. Laser speckle is created from the ground glass diffuser when the beam passes through. Subsequently a copy of the speckle is made and sent down a secondary path where it encounters a binary transmissive object. The light which successfully propagated through the object is then collected by a detector with no spatial resolution, a photo-diode (PD) in this case. The other copy of the speckle field propagates down another path, encounters no object as is recorded in full by a second detector with high spatial resolution, shown here as a CCD camera.*

contrast and signal-to-noise ratio of an image can be described by semi classical photo detection theory [16]. This proves that GI is not inherently quantum.

Classical ghost imaging was the first to expand the observational targets beyond that of binary transmissive objects. Since the single value measurement of a pattern after transmission through an object, quantifies the correlation between pattern and object, any greyscale or binary transmissive object is now a viable target. See Chapter 2 for experimental reconstruction methods.

1.3.3 Computational ghost imaging

The need for a beam splitter in the classical ghost imaging approach can be removed by using a SLM, which generates a programmable light field to illuminate the object. Controlling both the intensity and phase, allows the intensity structure to be calculated at any plane and stored in computer memory rather than measured on a detector with high spatial resolution. This simplified experimental approach is known as *computational ghost imaging* [17] and can be performed on systems with and without lenses between the source and object. When a ghost imaging system is employed, any phase information becomes redundant, since only the intensity correlations of the light reflected or transmitted from an object are used to produce an image.

The removal of the detector with spatial resolution (the CCD) is key to understanding why computational ghost imaging consistently outperforms classical ghost imaging [18]. Any given CCD or detector array has a quantum efficiency value < 1 , which compromises the measurement of the pattern to some degree. Since the value measured by the PD is associated to a specific pattern, any deviation from ‘true’ values or misrepresentation will cause a reduction in overall image quality. By removing the CCD entirely, this eliminates quantum efficiency entirely from affecting the SNR on the measurement of the pattern, since it is simply known and stored in the computer.

Computational ghost imaging forms the basis for the applications demonstrated in this body of work, since it provided a simplified and improved version of the classical technique. In addition, it was a platform to develop prototype systems which could be made into portable imaging systems, not constrained to an optical bench.

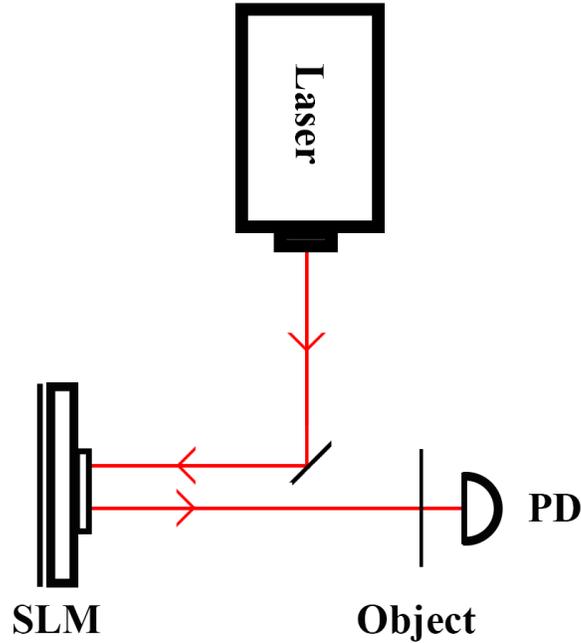


Figure 1.4: *Diagram showing a simplified computational ghost imaging setup. Pseudo-laser speckle is generated by the Spatial Light Modulator (SLM) removing the need for a beam splitter and second detector with spatial resolution. The patterns are pre-generated on a computer and it is these patterns which are correlated with the values measured by the PD.*

1.3.4 Normalized ghost imaging

Standard computational ghost imaging improves over the classical ghost imaging system, in effect due to the removal of the CCD and thus removing the quantum efficiency of that device from contributing factors of Signal to Noise Ratio (SNR) reduction. Reintroducing the reference path back into the system under a specific circumstance can make further improvement to the images[18]. This is achieved by introducing a second PD in the reference path which samples the full intensity of the field, providing a normalization factor the measurement made by the original PD. This normalization factor helps account

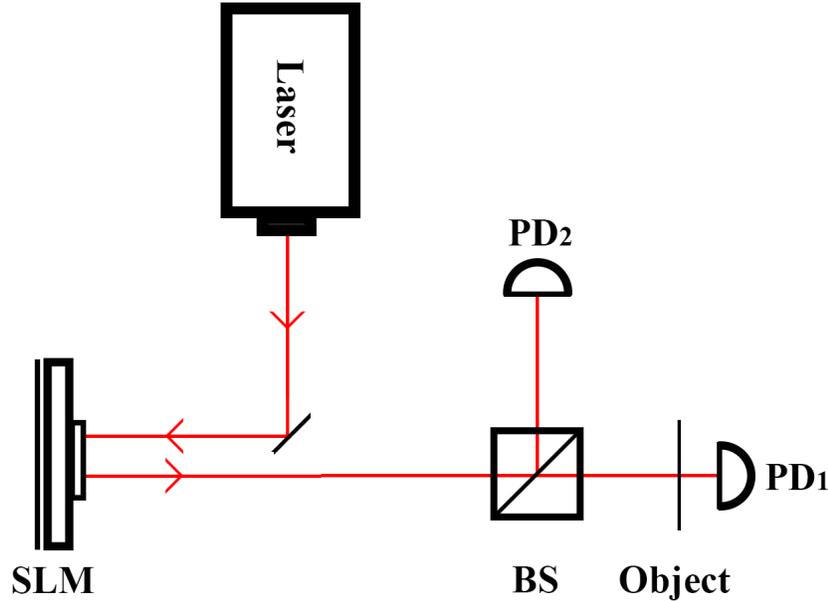


Figure 1.5: *Diagram showing a simplified normalized ghost imaging setup. compared to standard computational ghost imaging the BS has been reintroduced and a second single element detector (PD2) has been added. This extra detector measures the full field intensity of the pseudo-speckle pattern before it encounters the object thus providing a normalization factor to the signals measured on the original photodetector (PD1). This can be shown to give a baseline improvement over standard computational ghost imaging.*

for time varying noise sources such as fluctuating laser illumination. This fundamentally improves upon standard computational ghost imaging.

It is also possible to perform pseudo normalized ghost imaging in cases where data was obtained via a standard computational ghost imaging setup with a single detector or where a second detector is not available. In this instance, since known patterns which are stored on a computer are used, an estimate of the full field intensity can be made and used as the normalization factor to achieve a similar result to direct observation of the full field intensity.

1.4 History of Single-Pixel Cameras

1.4.1 Single-pixel and ghost imaging cameras

A general optical rule in any imaging system is that it is reversible, this also holds true in correlated single-element detector imaging. By switching the position of the illumination source and PD as shown in Fig. 1.6 b) the reverse process can still be shown to be a viable imaging technique. Within the community this version of correlated imaging is known as a single-pixel camera [19, 20, 21].

There are several further distinctions to be made when comparing ghost imaging and single-pixel cameras in relation to the types of observations best suited to each system, which is primarily due to the light-field encoding occurring before or after interaction with the object. Notably, single-pixel cameras can make use of background illumination sources, which allows for example, observation of objects which are self-illuminating or too distant for practical illumination by a nearby light source. While in contrast, ghost imaging systems are therefore excellent candidates as an imaging approach for objects which absorb certain illumination wavelengths, such as molecular gases or fluorescent material where active stimulation is needed to produce those effects. Therefore, the choice of which mode to select is largely application dependent. Any object able to be active raster scanned, can and should be measured via encoded light field-projection or ghost imaging. Conversely, any object which must be passive raster scanned must also be measured in an encoding masking, or single-pixel camera regime.

In the case of using a spatial mapping device such as a DMD which has two input/output angles, it is entirely possible to construct a system with a light

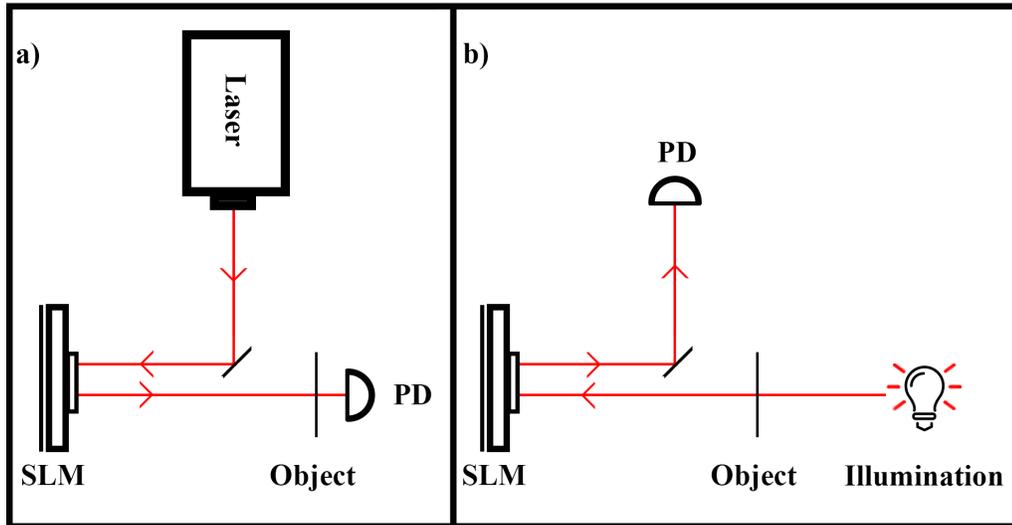


Figure 1.6: a) Copy of Figure 1.4, a simplified diagram showing a computational ghost imaging setup. b) A simplified diagram displaying a single-pixel camera setup. From comparison with a) it is clear this is a reverse process. The measurements and computationally stored patterns can be identical in both systems, only the position of the illumination source and PD has changed.

source at one input angle, and a PD at the other output angle to allow switching of computational imaging and single-pixel camera ‘modes’. This is extremely useful in creating a camera system capable of observing any target and selecting the best mode based on object properties or desired imaging application. This ‘on the fly’ reversibility showcases the versatility offered by this technique over conventional imaging systems.

1.4.2 Reconstruction Timescales

Figure 1.6 demonstrates the two iterations of computational single-pixel imaging side by side, clearly demonstrating the ‘soft’ reversibility of this technique

(Light source and detector positions must be switched). Both these techniques utilise purely classical sources and require no quantum effects for image reconstruction. Historically, the system shown in diagram a) always utilised either an analogue rotating ground glass diffuser or a digital SLM to produce speckle. In the first case, the advantage was true random speckle patterns however, this was offset completely by the degradation in image quality due to the need for a CCD device to measure the speckle patterns produced.

Digitising the experiment by including a SLM allowed the removal of the CCD. While a computer controlled SLM could only approximate randomness in the patterns it produced, there was a net gain of final image reconstruction quality. This became the preferred method for all classical ghost imaging experiments however, the use of SLM technology puts a hard limit on the reconstruction time due to relatively low pattern projection rates, between 10 and 100Hz. Where 100Hz was far more technically complex and computationally heavy, and thus most experiments did not attempt to use temporally dynamic objects, opting instead for static transmissive objects while clearly stating that the results obtained were over extremely long timescales ($\geq 20min$). This effectively stopped the technology from competing with standard imaging devices capable of high resolution video rates, as it could not compete with standard single image exposure times of a normal CCD.

Recent advances in projector technology have allowed the use of Digital Micro-Mirror Devices(DMD) to be used in computational single-pixel imaging experiments in new applications such as 3D imaging [22]. The use of the DMD vs the SLM does introduce a non-coherent propagation of the pattern since the SLM modulates the phase of the wave with a small attenuation of the intensity while the DMD is an intensity modulation only. This requires the use of a lens in the case of a system using a DMD to focus the pattern on a plane,

introducing a new limitation into the system, but one which any traditional imaging device must also adhere to. However, it is the potential speed of DMD devices which make them extremely attractive to computational single-pixel imaging experiments. A standard digital light projector(DLP) containing a commercial DMD already displays full colour RGB images at 60Hz. With appropriate pattern selection and encoding, the binary pattern projection rate is 1440Hz. This alone reduces image reconstruction times to the minute scale for high resolution reconstructions(128×128), or several seconds at low (32×32) resolutions.

The most recent DMD developments include devices which can be pre-loaded with a limited set of patterns, which sets the maximum resolution, and can be run at extremely high speeds approaching 22.7kHz. These devices offer the potential to produce more than a single image in a second, giving rise to real time reconstructions and constitute the current upper limit of single-pixel image reconstruction timescales.

1.5 Summary

Capturing an image of an object, while that object does not lie in the path of the measurement device is the core principle of the original quantum ghost imaging experiments. This is achieved by considering some correlation effect. Time correlation between the two entangled photons in the original quantum experiments allow the CCD to only capture those photons which would have passed through the binary transmissive object, if it had been there. The convolution of a transmitted pattern and transmissive object, create an attenuated light field, if measured by a bucket detector can assign a weight to the original pattern for how well that pattern happened to match the object. This causes

the average pattern to converge over many iterations to the desired object. While fundamentally different approaches to imaging where there is no detector with spatial resolution in the object path, they both share this idea of correlation.

Improvements to the classical version of the technique allow the camera to be removed entirely, using computationally stored patterns, and some device to spatially shape the light-field to approximate the stored pattern. Technological advances projection technology allows for much higher projection rates of these computed patterns, reducing the overall acquisition time to obtain an image. It is this linear evolution of the SPC approach which creates the potential for a real-world system that can be considered a viable alternate imaging solution.

Chapter 2

Experimental Methods

2.1 Spatial Light Modulation

There are only two criteria in the reconstruction of an image using the correlation between coded light fields and a single-pixel detector, the ability to shape a light field, and record the summed value of that light field with a detector. This imaging criteria is also invariant with detection bandwidth, or the number of simultaneous wavelengths being observed, if each wavelength can be shaped and recorded with the same degree of precision. Therefore, the selection of the encoding technology used to spatially shape light fields within a single-pixel camera(SPC) determines exactly how that system operates and what real-world applications it is suitable for.

Spatial Light Modulators(SLM) are normally used in the context of computer controlled light modulation devices [23, 24, 25] however, a wider definition exists whereby an SLM is any optical device which imposes a spatially varying modulation on a light field, for example an overhead projector transparency is

an example of non-controllable, fixed “SLM” . Some of the first Classic Ghost Imaging (CGI) experiments [14, 15] used a non-controllable “SLM” in the form of a rotating ground glass diffuser producing true random laser speckle. Since they had no control of the spatial shaping of the light field they instead used a high-resolution CCD to measure and temporally correlate measured light fields with signals. This was a source of image degradation in the system as the CCD image of the light field is only an approximation of the actual random laser speckle, compromised by the intrinsic quantum efficiency of each pixel within the CCD.

The ‘second generation’ of CGI experiments [26] could remove the CCD in the system by using a computer controlled Spatial Light Modulator (SLM) to produce pseudo random speckle instead of true laser speckle. Since the pseudo random patterns were stored on a computer, these patterns could be used in the processing stage of single-pixel imaging reconstruction, instead of measured patterns from a CCD. This relied on the assumption that what was produced by the SLM matched the stored pattern in relative intensity, which was corrected for by a known intensity correlation curve for any given SLM. Sources of noise in the system were now contained to fluctuating power within the laser, and the noise profile of the chosen single pixel detector.

SLM’s can modulate the phase or amplitude of an incident light field and have a wide variety of commercial and research uses. There are two main categories of SLM, optically addressed (OASLM) and electronically addressed SLM (EASLM).

The primary SLM devices employed in single-pixel imaging systems is a Liquid crystal on silicon-SLM (LCOS-SLM) and is a reflection modulation device capable of phase or amplitude modulation. These are the ideal system as they offer full wave-front control of an incident laser beam and are an example of

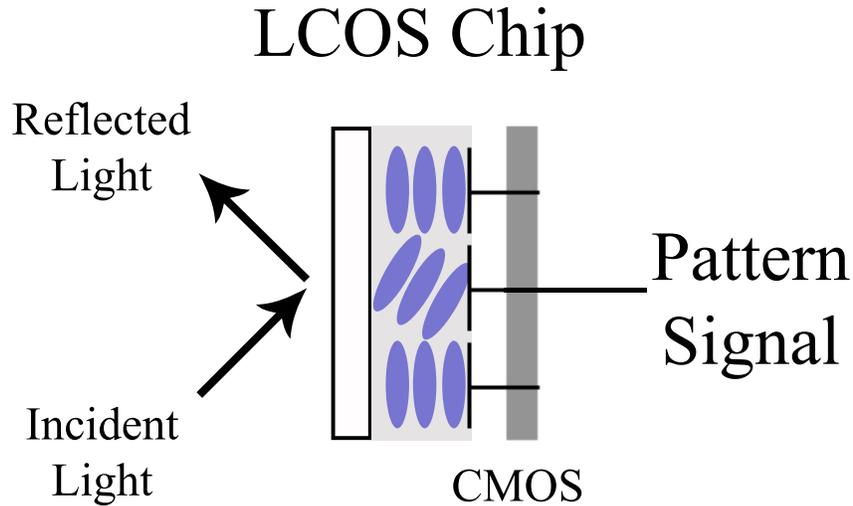


Figure 2.1: A simple diagram showing the internal configuration of a LCOSLM chip. This is an example of a phase modulating reflective device using an incident light from a laser to illuminate the chip. The single-pixel camera system is then aligned to the output path of the SLM. Each pixel in the SLM device is controlled via a computer connection which transmits the appropriate pattern.

an EASLM. Figure 2.1 describes the operation of a LCOSLM in a simplified diagram. The system works in reflection only, and in the case of single-pixel camera experiments, an optical bench mounted laser usually acts as the illumination source.

2.2 Structured illumination and Coded Aperture methods

Structured Illumination(SI) describes any single-pixel system which pre-encodes the light-field with pattern information, and transmits that encoded onto a ob-

ject. This total intensity is then measured by the single-pixel detector (SPD), creating a ‘score’ for the pattern to be used later in reconstruction. Conversely Coded Aperture (CA) describes any system with allows non-encoded light to interact or propagate through the intended target object first, and is then subsequently encoded with pattern information. This post propagation encoded pattern is then measured by the SPD, creating an identical weight compared to SI for the same object, pattern, and setup. This is true for any object for which both SI and CA are applicable and where any incident light field undergoes reflection only. The distance between the object and the system, and the properties of the intended imaging target, determines the applicability of either SI or CA methods when encoding the light-source with an appropriate pattern. For example, if an object range exceeds a reasonable distance to allow for free-space projection, CA must be employed. However, for certain applications such as in measuring absorption spectra, SI must be employed. Therefore, the choice between these two methods can sometimes be arbitrary but in most cases, is application dependent.

2.3 Pattern Design

2.3.1 Random

Using pre-generated patterns as the basis for reconstruction allows for the simplification and improvement of traditional systems with non-controllable SLM components. This can be considered a truly random pattern, where the pattern is generated by a sufficiently complex interaction with a moving diffuser. This completely analogue process produces laser speckle and by taking a 2D snapshot of this speckle pattern in a transverse plane, a random pattern can

be measured. Indeed, the non-controllable SLM in traditional ghost-imaging experiments was chosen precisely because it generated these truly random laser speckle patterns. To illustrate why this was advantageous, we choose to not have any a priori knowledge of what the object or its properties can be, allowing a GI or SPC system to reconstruct an image from any object. If we had instead selected some ‘ruleset’ for our object to follow, we would have a much easier set of choices to make in terms of pattern selection, but would be constrained to only image objects which fulfilled our pre-condition(s). For example, if we say that our object must be a circle with some unknown radius, r and known centre. and for which the maximum radius of this circle $r_{max} \leq FoV$, where FoV is our imaging Field of View, we could encode circles of varying radii with some increment size, Δr . From the signals returned by the SPD, we would know the exact reconstruction of our object for where the signal reached its maxima. However, we could only ever image circles in the above example.

When we choose to not impose any pre-conditions on our object, we must attempt to also choose the best possible encoding pattern set to maximise the information gained after each projection. We can test a random pattern generated by a non-controllable SLM, such as the rotating ground glass diffuser in traditional experiments, to determine if this results in successful image reconstruction the ‘best’ choice for information per pattern given some image reconstruction algorithm.

Figure 2.2 displays a set of example patterns for varying spot sizes in pseudo speckle patterns along with the simulated reconstruction of a cropped greyscale 2D image. These results, taken from early simulations, demonstrates that as expected, random pattern sets can produce positive image reconstruction, with a demonstration on how spot size affects detail representation in the

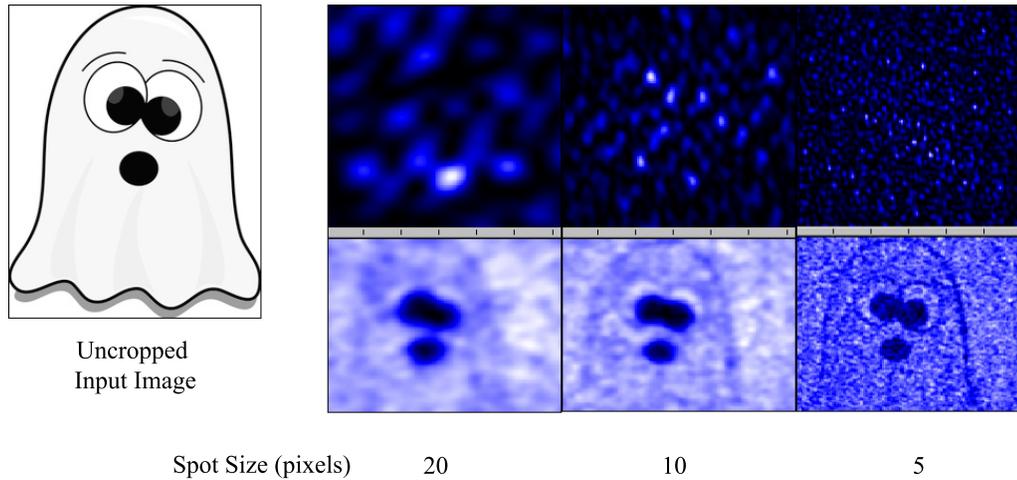


Figure 2.2: *Comparison of simulated results of a single-pixel camera system using a standard iterative reconstruction algorithm. Pseudo speckle pattern sets were chosen with varying average spot size where spot size is the diameter of the spot in pixels. The input image was cropped to fit the resolution of the pattern sets (128×128). It can be clearly demonstrated that no feature smaller than the average spot size can be resolved in the final reconstruction.*

reconstruction. Here as we move to a smaller spot size, there is an increase in the total number of patterns required for successful image reconstruction, but the ability to reconstruct finer detail in the image is also increased. This happens because features which were previously smaller than the average spot size, did not affect the ‘weighting’ or score of the pattern currently being projected onto the target. However, with a finer spot size, we increase the resolution in the final image. This was achieved within a Labview simulation program, where an image of a greyscale object was loaded into the program and multiplied by a series of pseudo-random patterns. These patterns were generated as simulated laser speckle, and not random binaries. The average spot size was a control value in the program which determined the pixel size of

the average ‘spot’ in the 128×128 images. These patterns were then summed to produce a simulated measurement that could otherwise be obtained by a PD taking a measurement in a real system. This provided a weighting factor for the original pattern. Through the standard iterative algorithm, reconstructions were made. A series of arbitrarily increasing average spot sizes were used to produce a set of reconstructions using an identical input image. This allowed a initial qualitative comparison in image quality to be made between each reconstruction, where the only value altered was the ‘grain’ of the pseudo speckle patterns used.

Figure 2.3 demonstrates the effect of SNR in a SPC based on changing the resolution scale of the patterns. In this case we are using pseudo-speckle patterns which have an average spot size which can be increased or decreased. As pseudo-speckle size in this example decreases, the final reconstruction appears to a human observer to become sharper, or more ‘in focus’. Since each speckle can be considered a ‘probe’ for testing a spatial component of the object, as these probes get smaller, the effective resolving power of the pattern increases. This is analogous to increasing the number of pixels in a 2D sensor array in a standard camera. We are observing less pixilation of the object as we move to the finer scale starting at a) to b) and finally at the finest scale shown in c). However, the estimated SNR for each image also decreases. This occurs because although we have a smaller minimum testable feature size, we also have a greater total number of features, or potential spatial components to test, yet we maintained a constant number of measurements for each reconstruction. In other testing regimes, such as using random binary patterns, it is instead the pixel size which sets the minimum testable feature.

This results in experimentally having to balance increasing resolution with greatly increasing acquisition time. Generally, to give good results at high

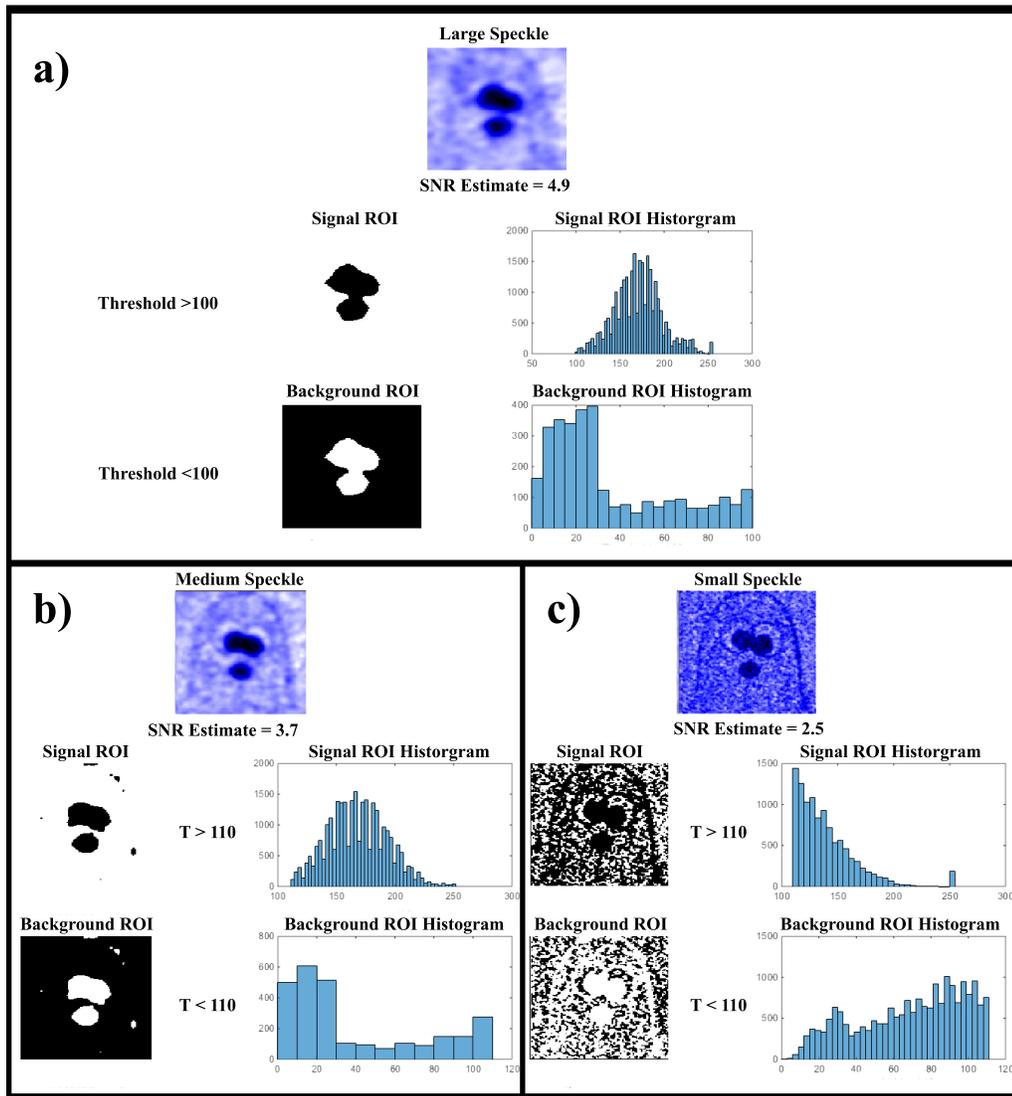


Figure 2.3: *Signal to noise ratio estimates for each spot size when the number of measurements remained constant. SNR decreases as effective resolution in the final reconstruction increases.*

resolution, many more patterns are required than at low resolutions. On the order of $M \gg N$ where M is the number of measurements made and N is the number of pixels in the final reconstruction.

SNR Estimation method for single images

The SNR estimate made here, and in all subsequent images which were quantitatively analysed, is based on Region Of Interest (ROI) selection. This method is an established technique [27, 28] for estimating SNR based on pixel values for single images, when no reference image is available. Firstly, all images are converted to an 8-bit greyscale for this process. As can be seen in Fig. 2.3 a threshold value, T , was manually selected for high pixel values and another T value was selected for low pixel values. When all pixel values are viewed in a histogram, there are many pixels which form a pseudo-Gaussian ‘peak’ in the low region, which indicates where most of the background pixels lie, with the FWHM of this gaussian dependent on the standard deviation of these low value pixels. The background threshold should be set to include this peak, but not include higher values which may be signal pixels. For good contrast images, this process is simple, with the background peak easy to define.

For the signal threshold, a similar process is employed by observing a histogram of the pixels. In good contrast images, there should be many pixels above a certain pixel value, this is most extreme in the case of a knife edge object illuminated with a uniform source. The exact transmissivity and dynamic range of the object will affect this signal ‘wall’ so it can sometimes be more challenging to define depending on the object and SNR in the image. Generally, selecting a high value threshold which displays few or no pixels in known background regions is an appropriate approximation. This leads to selecting the middle of the pseudo-Gaussian peak associated with higher value pixels. This is a general procedure to follow, which can be hampered as true SNR decreases. Therefore, for estimates which are close to 2 or below 2, this method is not robust and a reference image must be employed to get a valid estimate of the SNR.

After finding the threshold levels, we now calculate the standard deviation, σ , of the background ROI pixels. This allows us to find the variance of σ and subsequently the square root of the mean variances defines the Root Mean Square Noise (RMS Noise). Taking the mean value from the signal ROI, and subtracting the mean value from the background ROI defines the net signal. The net signal divided by RMS Noise is the SNR estimate. A complete list of Figures describing the image analysis process for images in which an SNR estimate was made can be found in Appendix A.

2.3.2 Orthogonal

Regardless of the generation method, any truly random pattern will inherently have overlap in spatial information between any other pattern. This overlap reduces the quality of a reconstruction based on a random pattern set, as the contribution to the final reconstruction arises from the various spatial ‘testing’ done by each pattern when correlated with its measured returned signal. Creating a pattern set with zero overlap and thus complete orthogonality from one pattern to the next, naturally maximises the contribution of any one pattern. A common orthogonal sampling approach makes use of the Hadamard basis [29], from which each pattern can be derived and where the number of patterns is the same as the number of pixels to be sampled. The basic building block of any Hadamard matrix, with a dimension length of 2, is

$$H_2 = \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \quad (2.1)$$

For any matrix of order 2^k there exists such a Hadamard [30], defined as

$$H_{2^k} = \begin{bmatrix} H_{2^{k-1}} & H_{2^{k-1}} \\ H_{2^{k-1}} & -H_{2^{k-1}} \end{bmatrix}, \quad (2.2)$$

which is known as a Sylvester constructed Hadamard matrix. In this way the Kronecker product, $H_2 \otimes H_{2^{k-1}}$, can produce any square Hadamard matrix and hence the number of elements in a row equates to the number of rows of the matrix. Each pattern that illuminates the object, $P(x, y)$, corresponds to one row of the large matrix that can be reshaped and up-sampled onto the spatial encoding device.

2.4 Image Reconstruction

2.4.1 Pattern and signal

A successful reconstruction as shown in Figure 2.2 requires the sampling of many random or computer-generated patterns onto a target object. Each pattern tests a set of possible spatial frequencies within the object, giving some unknowable fraction of information about each pixel per pattern. The results of this test are given by the numerical value obtained from the single element detector, usually a photodiode (PD). Therefore, each measured bucket signal becomes a weighting factor for how well a random or pseudo-random pattern matched an object. The total number of random patterns, M , needed to obtain reasonable quality reconstructions is generally greater than the resolution, N , of any individual pattern e.g a 32x32 reconstruction would require > 1024 patterns.

This creates a ‘library’ of numerical values which have corresponding patterns associated with them due to the synchronous nature of the measurements, visually represented in Fig. 2.4. It is from these two sets of information that a computer algorithm can then calculate the shape of the original object. As a simple analogy, lets imagine taking randomly generated binary patterns

and printing them out on acetate transparent paper. By having an estimate of the ‘value’ of each pattern in terms of how well did the spatial feature distribution happen to, by complete chance, match the object observed. If one were to overlay our printed non-weighted patterns, we would observe simply a random distribution of pixels. However, if these patterns were weighted by the measurements made by the photodetector so that a higher measurement value corresponded to a relatively darker, and lower opacity print, we would find that as more weighted patterns were added, spatial features corresponding to the object would begin to be observed.

Whenever a pixel was ‘turned on’ which corresponded to a bright spot in the object, a high fractional change would be observed in the measurement photodetector. For every pattern that contained this important pixel, that pattern would increase its weight or score by some degree. Over many iterations these patterns with this important pixel turned on would show a darker spatial component at this pixels location. Conversely, any pattern with this pixel turned off, might be weighted less, producing a fainter, higher opacity prints on the transparent acetate. Thus, when stacking the transparent sheets, only those high scoring patterns could affect the combined image to a significant degree. Since these high scoring patterns have a high degree of correlation with the object, the average pattern converges to the object observed by the measurement photodetector.

Definition of signal

Equation 2.3 defines the signal, S , obtained from the detector in the system with no spatial resolution. Each signal is in effect the summed multiplication of the encoded light field, $I(x, y)$, and the transmission function of the object

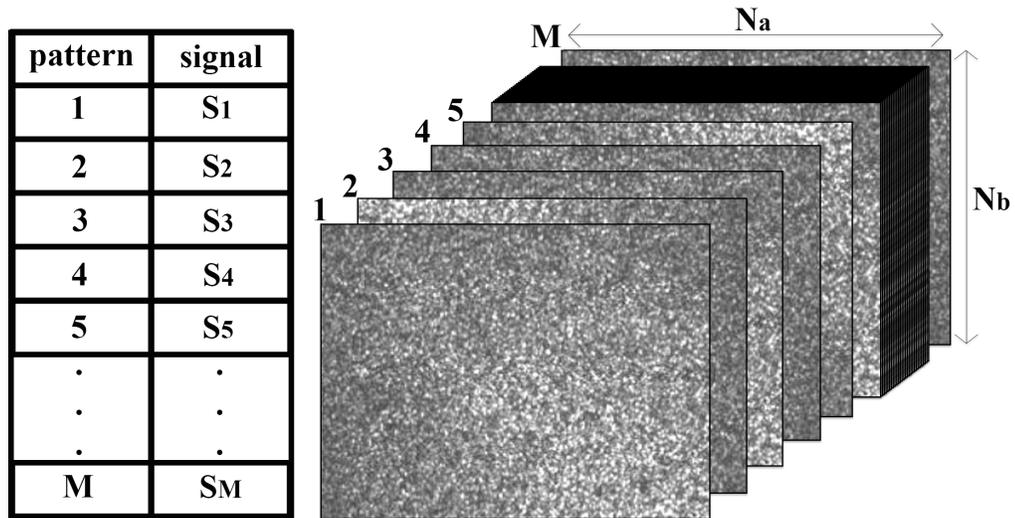


Figure 2.4: *Visual representation of the data ‘library’ accrued on a computer after acquisition is complete. Each averaged signal has its own corresponding pattern. For a good quality image to be reconstructed from this data, $M \gg N$, where M is the total number of measurements made, and thus total number of patterns sampled and $N = N_a \times N_b$ is the resolution of each individual pattern, and also the resolution of the resulting reconstruction once an appropriate algorithm has been applied to this data.*

placed in the object path, $T(x, y)$.

$$S = \int I(x, y)T(x, y)dxdy, \quad (2.3)$$

During the data recording process, many dozens of signals can be measured for each individual encoded light field since the digital to analogue converter (DAC), which is used to capture signals from the PD, can record data several orders of magnitude faster than a SLM based system can display patterns. The average of these many signal measurements becomes a better representation of the true signal value for any given pattern, than using any one signal by itself. For example, attempting to use only a single measurement of the PD as

a weighting factor for a specific pattern generally results in an extremely poor or no reconstruction result, depending on the transmissivity of the original object. However, the situation where one would be forced to do this never generally arises since DAC measurement rates far exceed the projection rates of any SLM.

Weighting the pattern

The first step in image reconstruction from the stored library of data is achieved by assigning a weighting factor to each individual pattern based on the strength of signal measured by the PD. A visualization of a sample of weighted pattern contributions, and the corresponding normalized signal values is shown in Fig. 2.5. Here an identical sample of four patterns is selected for three arbitrary, but distinct objects. However, due to differences in the spatial and/or spectral frequencies of each unknown object, identical patterns can score very different weighting factors during discrete runs between individual objects. This can be observed by the distinct opacity differences in each column of table b). Fig 2.5 describes the correlation between a signal measured on the PD and a visual representation of the weight of the corresponding pattern. This will be an important visualization to refer to when pattern reconstruction methods are introduced in the next section. Here only a small sample of patterns are visually depicted however, in any single set of data, many hundreds or thousands of patterns are used to produce an accurate representation of the object.

Figure 2.5 is an illustration of the weighting process only. It is designed to visually represent this process during live image reconstruction when an appropriate algorithm is applied. A real dataset would contain many more patterns for the resolution shown before successful image reconstruction could take place.

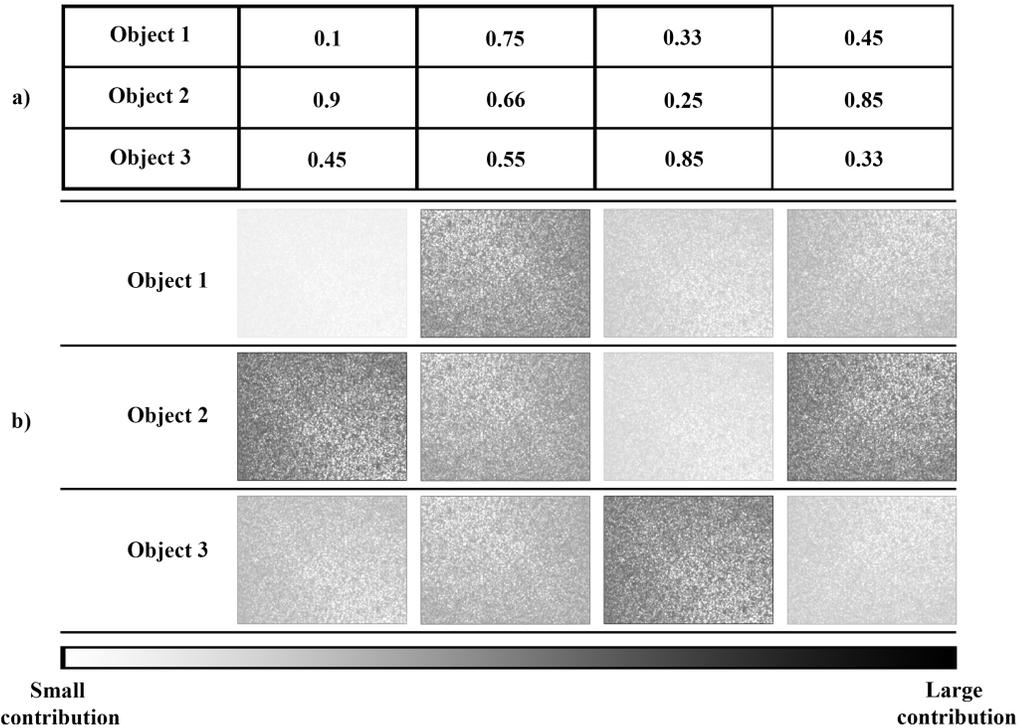


Figure 2.5: a) Table of a sample of normalized signals measured by the PD, for three different objects in discrete data acquisition runs, where the normalization value is set simply by the highest intensity returned signal present in the data library. b) Visual representation of a sample of ‘weighted’ patterns for the three objects listed in table a). The contribution or ‘weight’ of each individual image is shown by its relative opacity.

2.4.2 Generalised Signal and Noise Estimates

How noise affects various SPC regimes is an active research topic, however a generalised estimate can be made when dealing in either a fully detector limited or laser power limited regime. Comparison of these generalised estimates can allow selection of the best potential type of SPC, including active and passive raster scan system.

Table 2.1: Generalised Signal, Noise and S:N Estimate

Single Pixel Camera		Passive Raster Scan	Active Raster Scan
$\frac{P}{2\sqrt{N}}$	The signal	$\frac{P}{N}$	P
σ	Noise in the detector	σ	σ
$\alpha\frac{P}{2}$	Noise in the illumination	$\alpha\frac{P}{N}$	αP
$\frac{P}{2\sigma\sqrt{N}}$	$\frac{S}{N}$ when Detector Limited	$\frac{P}{\sigma N}$	$\frac{P}{\sigma}$
$\frac{1}{\alpha\sqrt{N}}$	$\frac{S}{N}$ when Signal Limited	$\frac{1}{\alpha}$	$\frac{1}{\alpha}$

Table 2.1 shows a compact comparison chart for Signal, Noise and $S : N$ for a SPC, Active Raster Scanning(RSa) and Passive Raster Scanning(RSp) systems. The estimate for signal in a RSa system is simply the laser power, and in the RSp this laser power is spread over each pixel in the passive system, thus $\frac{P}{N}$. For a SPC which uses some subset of pixels to display a full field pattern. The laser or illumination power is spread over the pixels on the SLM, but scales with $2\sqrt{N}$. For a given detector we have standard average noise value denoted σ . The noise component within the illumination scales with laser power, but also a proportionality constant, denoted α .

Thus, dividing the S by N cells for a Detector or Signal limited regime gives the best case scenario for each type of system. An active raster scan, where the laser power is not distributed over many pixels, has the lowest inherent noise based scaling with only the ratio between laser power and detector noise when detector limited. This is perhaps unsurprising, however there might be an argument to make for the SPC when comparing to passive raster scanning as noise scaling with \sqrt{N} vs N beats RSp. There should be a caveat noted here that the size of the detector might ‘cancel’ out the advantage from the \sqrt{N} scale factor, but this would have to be experimentally tested, and would constitute novel research to find the optimal detector size for use with a SPC or RSp system.

It should also be noted that the RSp Signal limited $S : N$ appears to beat out SPC, however in practice, you will never be signal limited within the RSp regime. It is of course not just the direct S:N estimate that could dictate the choice of one system over another. SPC is inherently compressible and lends its self to noise reduction computational processes whereas RS generally does not. The advantage offered by a RSa system would need to be experimentally tested against an appropriate SPC to determine if compressibility overcomes the generally lower S:N in SPC compared to RSa.

The true ‘noise’ profile of any SPC does however depend on the actual contributions from the PD, the interaction with the SLM surface and the chosen light source. When measured and treated collectively, this has now been experimentally demonstrated [31] to give a fair assessment of the SNR and overall performance of any given SPC system, while it should be noted that some approximates have been made in terms of quantification of the laser light source.

The reconstructions in this work were obtained over several years from various LABVIEW output programs. In most cases, only the final reconstruction is available for quantitative analysis. This has resulted in estimations being made using only a single image, with no reference image available. No actual SNR measurement or calculation can be made on one single image, however for good and moderate contrast images, the estimation method for single images does provide a valid result for the SNR. The results obtained are effectively a measure of the total collective contribution of all noise sources within the system. Changing one will affect the final SNR of the reconstruction. The generalised examples described in Table 2.1 allow easy characterisation of the noise contributions in an SPC by setting one factor to be fully limiting. In this way, it is possible to make a choice of system based on knowing some basic parameters of the SPC, such as total illumination intensity, the type of

detector and its spectral response to the input illumination wavelength. This should enable a user to select the regime which will most likely match the conditions in which the SPC camera system will be operating, or to opt for a raster scanning approach instead if SPC performance is not expected to be adequate.

Pattern and Transmissivity contribution

An interesting aspect of the SPC is that by changing the construction method of your pattern basis, or using a fast SLM to reconstruct a temporally dynamic object, the SNR of the final reconstruction can change. This is of course in addition to the effect on the final SNR by the collective noise contributions of the PD, the surface interaction with the SLM and the chosen light source. Combined this leads to an extremely complex set of parameters which determine the final SNR of the reconstruction.

To simplify things let us consider the case of an active raster scanning system, RSa. Each coordinate position with which an incident laser beam illuminates the target object becomes a pixel in the final image. Each pixel must be measured at least once to fully reconstruct the object, setting a hard limit on the number of measurements to be made where $M = N$, M is the number of measurements and N is the number of pixels in the final image.

In any SPC regime, we instead test spatial components within the object simultaneously, and have a number of different configurations of patterns, $P_i(x, y)$. The more positive reflections we receive from the individual pixels or speckle spots in these patterns, the higher the score it will receive for any given object. Let us assume we have some prior knowledge of our object, in this case, our object is a bright pinhead. No matter the resolution scale we chose to observe

this object, only one pixel will ever correlate to the pin heads position. If we use a 2×2 SLM fixed to a 50/50 black and white pattern ratio, there will only be 6 possible pattern combinations using random binary configurations in which to test our object.

As we cycle patterns we will receive either a positive result from the PD or no result. There can be no other output of the PD in this special case. For convenience, let us also assume the PD is a perfect detector which will give us integer values for measured intensity. For any of our 6 patterns, 3 of them will have an ‘on’ pixel which correlates to the position of the pinhead. Since our object is so small it will never lie in a mid-position between pixels. With these simple parameters, the minimum number of patterns needed for reconstruction with prior knowledge of the object is 1-4, with this variance dependent on ‘lucky’ imaging; getting a positive hit by chance, where you are guaranteed to get a positive hit on the fourth pattern if you have had 3 unlucky measurements. This is equivalent to the performance of RSa.

Without such prior knowledge i.e not knowing if there may potentially be multiple pin heads, we need 3-6, with the SNR of the final 2×2 result improving after each iteration. You will therefore equally beat the performance of RSa the same number of times as you fail to match it. However, tested multiple times with a random arrangement of the 6 patterns, there will be more runs in which the object will be partially reconstructed before RSa finds the location of the pin head. In these cases the SNR will be lower than the final value of 3, as shown in Fig. 2.6. Thus, with SPC it may be possible to extract the position of the pin head before 4 iterations, even if you happen to only get 2 out of three possible ‘lucky’ patterns. Whereas RSa is an ‘all or nothing’ approach based on luck for this special case.

In the case of two pin heads, the image can be fully reconstructed potentially

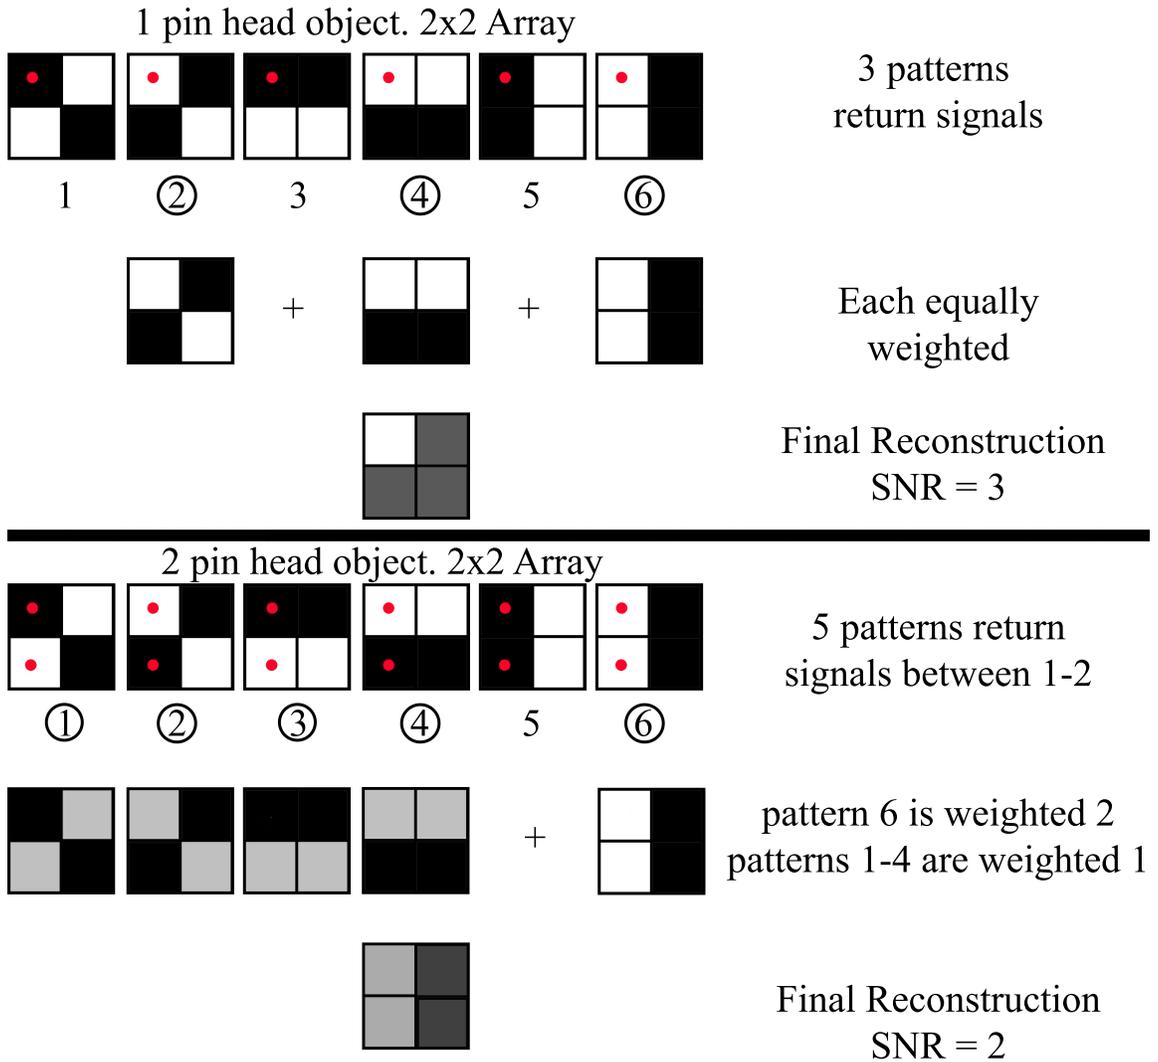


Figure 2.6: A visual representation of the reconstruction process for a 2×2 arrays and 1-2 pin head objects, note that in the case of 2 pin heads, both heads are assumed to be uniform brightness. Red dots indicate pin head true positions.

after 5 patterns, since this is the number of patterns which will return a positive hit of strength 1 or 2. This depends on whether the specific pattern had one pixel or two pixels correlating to positions of the two pin heads. As shown

in Fig 2.6, the final SNR for a single pin head will be higher than two pin heads, even with the same pattern basis. Thus, as you increase the number of spatial components in the scene, thereby increasing its total transmissivity, you decrease the SNR in the final image. This was clearly demonstrated in our publication [18]. Note that since a detector has been set to return a single integer for each positive correlating pixel, the value of the SNR may seem low in comparison to using at an 8-bit scale image (0-256 dynamic range). This is simply relative to the low numbers we have chosen to represent our signal with, therefore to describe SNR on varying scales, it is often better to calculate the SNR in Decibels(dB). For the remainder of the work, all images are treated on a 8-bit scale thus a simple ratio is sufficient. A brighter pin head will also return a higher signal which will be reflected in the result, but for this simplistic example, we have not considered non-uniform objects.

This makes SPC potentially a very good choice of system for finding the position of a few spatial components, often beating RSa in any case where there are a relatively few number of non-zero spatial components compared to the total number of pixels in the final reconstruction. The assumption has also been made that patterns can be displayed at an equal rate to the measurements being made by the raster scan.

Higher Resolution Scales

As spatial components in the object, and resolution required increases, the interaction on the result becomes far more complex. However, to maximise the potential contribution to the final image by each pattern, it is always better to ensure that there is no spatial overlap in your pattern set. Random binary patterns for example, always include overlap in the spatial components from pattern to pattern, due to the random nature of the pattern construction,

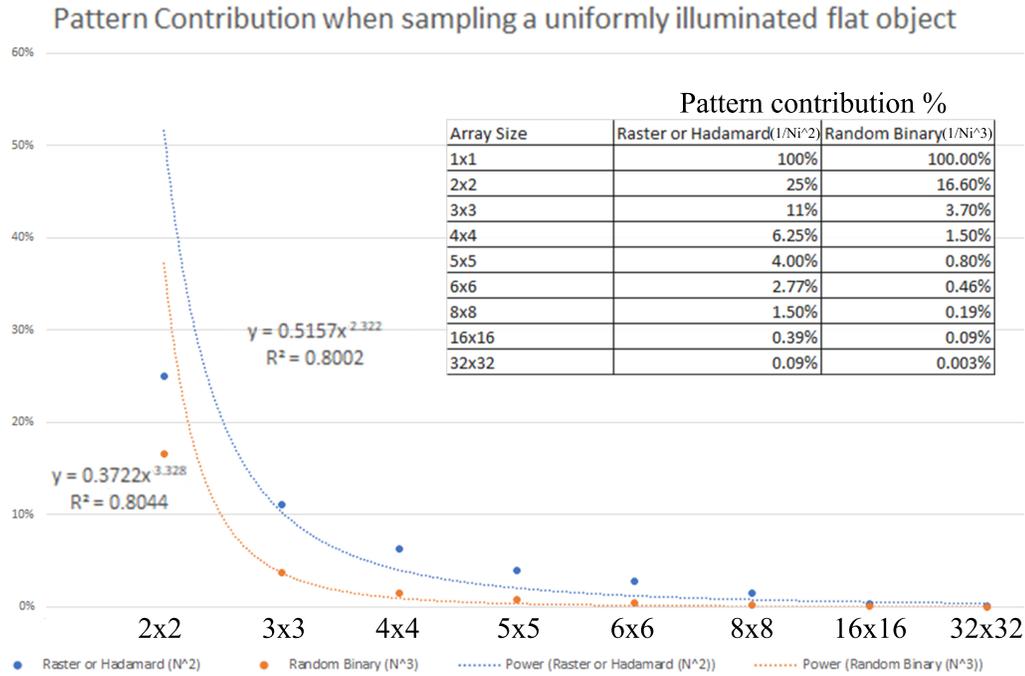


Figure 2.7: *Power law comparisons for testing pattern sets which contain either no spatial overlap or a random distribution, assuming a flat uniform object. Having no overlap means that pattern contributions vary as $\frac{1}{N_i^2}$ while having overlap causes the contribution to drop by approx. $\frac{1}{N_i^3}$*

while RSA never has any spatial overlap. For a range of pattern resolutions from 2×2 to 32×32 Fig.2.7 demonstrates the contribution each pattern makes to the result. This is only true when illuminating a uniform object which fully fills the scene, if the object was non-uniform, each pattern would contribute based on the individual spatial components which happened to match the object. We compare RSA with a randomly generated binary pattern basis and Hadamard patterns. Patterns based on Hadamard matrices match the performance of RSA, in terms of pattern contribution percentage. This is due to having the unique property of zero spatial overlap within their set, but only exist for square $N_i \times N_i$ arrays where N_i is the dimension of the array.

When raster scanning or using a Hadamard basis and excluding real noise, each pattern contributes exactly $\frac{1}{N_i^2}$, and thus the reconstruction is perfectly retrieved after N_i^2 measurements, for square arrays. In the case of random binary distributions, there exists many Degrees of Freedom (DoF) of total possible patterns scaling exponentially with larger resolution arrays.

Generally, it is found in experimental tests that random binary distributions scale on the order of $M > N_i^3$ to get a high SNR image from iterative reconstruction. This gives an indication that the true pattern contribution lies between $\frac{1}{N_i^3}$ and $\frac{1}{DoF}$, allowing confidence that at least each pattern is contributing more than the equivalent of turning a single pixel on and off at a time, the minimum possible contribution.

However, the order of magnitude gain in pattern contribution by simply eliminating spatial overlap in the pattern set is very significant in improving the overall performance of a SPC system. Thus, results which were obtained later in the work all utilise such non-spatially overlapping patterns.

Iterative Reconstruction

The iterative process generally allows for ‘live’ viewing of the reconstruction with incremental improvements frame by frame. If patterns are created randomly during the actual data acquisition process instead of a finite pre-prepared set, this type of reconstruction best suits image reconstruction where time constraints do not apply, as one can simply continue to sample until the image quality is at the desired level. Additionally, the computational load is minimised in such algorithms since only the current refined estimate and the next weighted pattern in the sequence is stored in the computer at any given time. Such algorithms were the primary reconstruction method of classical

and computational ghost imaging systems since acquisition times could exceed > 20 minutes for sufficient image quality.

The iterative algorithm generally employed in most systems is known as a traditional GI algorithm, defined by

$$O(x, y) = \langle (S - \langle S \rangle)(I(x, y) - \langle I(x, y) \rangle) \rangle, \quad (2.4)$$

where $O(x, y)$ is the estimate of the scene and $\langle \dots \rangle$ denotes an ensemble average. In the case where the average signal and average pattern tends to zero Eq. 2.4 can be rewritten into a more intuitive format as

$$O(x, y) = \frac{1}{M} \sum_{i=1}^M S_i I_i(x, y), \quad (2.5)$$

When a second detector is present which measures the full field intensity of projected patterns or an estimate of this full field intensity can be made, a normalization process can be applied to the measured signals which can account for some time varying sources of noise within a real system. This can be summarised as

$$O(x, y) = \frac{1}{M} \sum_{i=1}^M \frac{S_i}{R_i} I_i(x, y), \quad (2.6)$$

Where R_i is the reference signal measured or calculated for each individual pattern.

In general, it is found that after many samples, $M \approx N$ (Nyquist limit), an image of the transmissive object is reconstructed. An iterative sequence is demonstrated in Fig. 2.8. This result was obtained on our primary computational single-pixel imaging system based on the Texas Instruments Light Commander. A single Thorlabs DET36a Si PD was used to capture the 1024 measurements for a set of 1024 random binary patterns. Due to overlap in

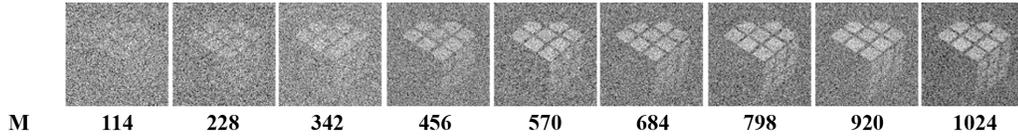


Figure 2.8: *Result of an iterative GI algorithm on a library of weighted pattern data taken from a standard normalized ghost imaging system for increasing number of patterns samples (increasing M) when observing a Rubik cube object. Each image is 32×32 within the final image reaching the Nyquist limit of 1024 measurements.*

spatial frequencies of the pseudo-random patterns used for this image reconstruction, further sampling beyond the Nyquist would lead to a reasonable increase in image quality compared to the final image in this sequence. Thus, Nyquist limit measurements using pseudo-random patterns do not necessarily produce the best possible quality image for any given system.

Inverse Reconstruction

For inversion algorithms we can reshape each 2D random binary pattern, which can be calculated as part of a pattern set intended to be used in a SPC camera system, as a 1D array, I_i , and produce a measurement matrix, \mathbf{I} containing all patterns, such that

$$\mathbf{I} = [I_1 \ I_2 \ \dots \ I_M]^T, \quad (2.7)$$

and vectorize the measured signals to

$$\mathbf{S} = [S \ S \ \dots \ S]^T, \quad (2.8)$$

such that the problem can be realised as a set of linear Eqs. given by

$$\mathbf{I}\mathbf{O} = \mathbf{S}, \quad (2.9)$$

where \mathbf{O} is a set of unknowns which when recovered, becomes $O(x, y)$, the estimate of the scene. In the case where $M \geq N$, this problem can be solved by least-squares methods. However, as the image resolution increases, the size of \mathbf{I} makes performing this reconstruction computationally intensive, while using fewer measurements than the resolution size to address this results in an ill-conditioned problem and the reconstructed image quality rapidly decreases.

While this technique has a much higher computational load than the iterative counterpart, it does form the basis for an approach to reducing the number of measurements to sub-Nyquist levels. Successfully solving the resultant ill-conditioned problem can provide much improved results for a given number of measurements when compared to the standard iterative GI approach.

2.5 Compressed Sensing

2.5.1 Ill conditioned inverse reconstruction

For images containing a relatively small number of pixels (< 4096 pixels) a fully Nyquist sampled object can be reconstructed using a simple matrix inversion algorithm however, images above this threshold are required for resolving fine detail in objects or for general high resolution imaging. Since the number of measurements, M , required for fully sampling an object scales with the square of the resolution of the final reconstruction, N , high resolution reconstruction quickly becomes too computationally and memory intensive to solve in a reasonable timeframe.

A possible solution to this problem would be to reduce the number of measurements to a sub-Nyquist level forming an ill-conditioned inverse problem.

In practice, it is found that even removing a small fraction of samples and solving using least squares methods destroys the reconstruction. To reduce the number of measurements required for faithful image reconstruction we can instead employ compressive sensing techniques[32, 33].

Compressed sensing takes advantage of ‘sparsity’ in natural images to reduce the important coefficients of an image, when that image is represented in an appropriate domain. Sparsity in natural images refers to the assumption that can be made to any image excluding those containing infinitely repeating detail i. e. excluding fractal images, in which the image can be represented by only a few non-zero coefficients in some transformed space. When any natural image undergoes a discrete cosine transform (DCT)[34] for example, the number of important coefficients which are non-zero are relatively few, this DCT is effectively a Fourier transform operation, but with imposed boundary conditions, making all frequencies a superposition of cosines instead of sin waves. A matrix inverse algorithm [35, 36] capable of handling ill-conditioned matrices can then ‘solve’ this matrix problem, with the solutions of this problem taking the values of the reconstruction. The strength of this approach is that only a small fraction of patterns are needed to fully reconstruct the image, as long as the assumption that the image is sparse in an appropriate domain held true.

2.5.2 Compressed Sensing Results

Figure 2.9 is a single-pixel camera result using the traditional iterative reconstruction method for an increasing number of measurements. This result was obtained from a standard SPC based on Thorlabs Light Commander projection system, using a binary transmissive target object. Pseudo-random binary patterns were used and the single PD was placed in reflection to capture most

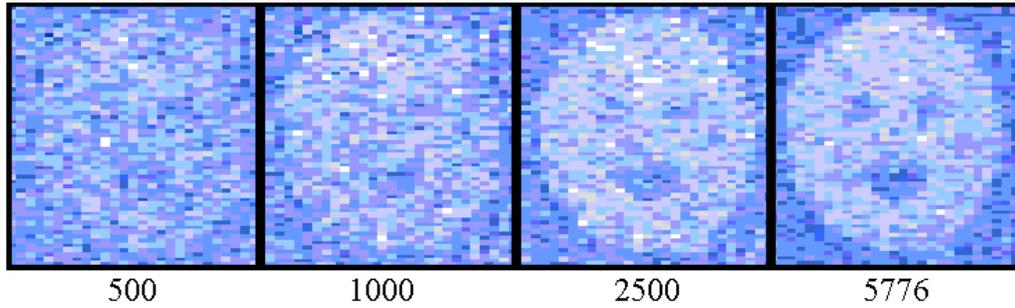


Figure 2.9: *Iterative TGI algorithm reconstructions for increasing measurement values for a 76×76 resolution image ($N = 5776$). Test object is a white card with cut-out sections resulting in a binary transmissive object. These results were obtained from a DMD based computational single pixel imaging system*

of the diffuse reflections from the object. This demonstrates the expected increase in image quality as the number of measurements increases. There is still a relatively low SNR however even when, M , matched the number of pixels in the final reconstruction, N . Using random patterns in a system such as this would normally require a greater number of M than N pixels. This is therefore a poor imaging technique when compared to a simplistic active raster scan imaging system, for example. Such a system can perfectly reconstruct any object at a fixed $M = N$, with noise in the image limited to real sources of noise from the detector or fluctuating laser power.

Figure 2.10 is a set of images which used identical base data to create the image sequence shown in this figure, as was used in Figure 2.9. However, this data has been processed by a compressed sensing algorithm [35, 36]. This is achieved by restructuring the already obtained data used to produce the iterative result in Fig 2.9 into a large $M * N$ array where each row in this array is a vectorised pattern. This large array is expressed in a sparse domain, in this case the DCT

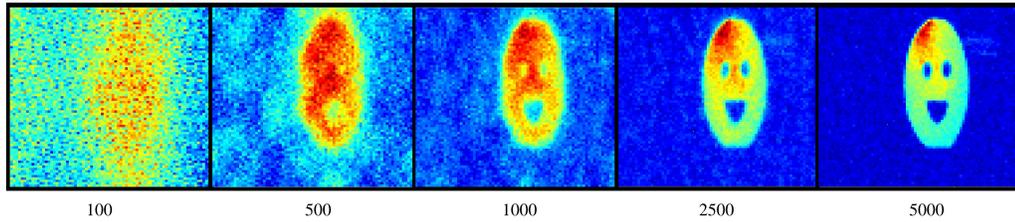


Figure 2.10: *Iterative look at CGI algorithm reconstructions for increasing measurement values for a 76×76 resolution image ($N = 5776$). Identical data as demonstrated in Fig. 2.9 with compressive techniques applied.*

transform function [34] was applied so that the matrix could be represented by a few non-zero coefficients. Subsequently, this sparse matrix, with associated measurements from the PD. The result is vastly improved image quality, even where $M \ll N$. Even at a 'half Nyquist' number of $M = 2500$ we observe that the image has almost been completely faithfully reconstructed. Taking advantage of sparsity in natural images allows better use of the information available, leading to increased image quality for a smaller value of M than with traditional methods, with the offset being an additional computational processing step.

2.6 Summary

At its core, SPC technology is based on the spatial shaping of a light field of a known computer pattern, and subsequent reflectance or transmittance measurement of that encoded light field after interaction with some object, with correlations between the two giving rise to a reconstructed image. Therefore, the performance and choice of the light source, the PD, the pattern encoding regime and the algorithm employed to reconstruct an image are the four crit-

ical ‘components’ which will determine the overall success and output of any system.

The algorithm however, can generally be applied as a post process to any SPC data and thus, only the first three parameters are needed to judge the ‘quality’ of data, while the fourth is the applied process which generates the final output. Which algorithm to use, whether it be a iterative, inverse or compressive is dependent on several factors. These include the number of measurements and the number of pixels in the final reconstruction which effect computational load. For example, having too large a matrix for the computer to store it in memory, results in being forced to use iterative reconstructions. These can handle any number of measurements or pixels in the final reconstruction, provided the final reconstruction is not so large a single pattern at a time and the current reconstruction cannot be stored in memory simultaneously. However, given current SPC standard resolutions mostly contained within the sub megapixel regime, this is not usually a problem for iterative algorithms. While a compressive algorithm can drastically beat the performance of an iterative or inverse algorithm, however, there is a larger computational load for higher resolutions.

SPC technology is also a reversible process in which ambient light can interact with an object and be encoded on and SLM, or an illumination source can be pre-encoded before interaction and ‘projected’ onto the object. Object parameters generally make this choice arbitrary or extremely well defined e.g. If the object is far away, or is astronomical in nature, it is not possible to ‘project’ encoded light fields onto such objects. While conversely, objects such as a molecular gas cloud require specific illumination targeted at a known absorption line so broadband ambient light is not suitable for encoding. This highlights SPC as an extremely adaptable technology for a wide range of tar-

gets and applications where the camera can effectively be ‘tuned’ by changing detection components or illumination mode, depending on the desired output or requirements set by target object parameters.

Chapter 3

Multi-wavelength single-pixel imaging

3.1 Introduction to multi-wavelength imaging

All materials and gases have unique spectral responses over the electromagnetic spectrum, and imaging devices sensitive to various parts of that spectrum reveal this well-defined behaviour[37]. Increasing the bandwidth of an imaging device can increase the total signal intensity, assuming there is a sufficiently broadband illumination source. However, selecting multiple small bandwidth regions in which to observe can reveal spectral properties of the object. Such information can be lost in broadband single channel images, where details might not be distinguishable from overlapping structure pertaining to another frequency. Therefore, multi-wavelength or spectral imaging can be a powerful tool for non-destructive material analysis.

Commonly, Red, Green and Blue (RGB) images are used in combination to

produce full-colour images. Since the human eye only contains receptors for these three channels, the combination of these colours gives rise to all other observable colours and thus the visible spectrum can be fully represented solely by RGB combination. In visible photography, it is always therefore desirable to be able to capture at least these three colour channels to properly represent any natural image. In non-visible imaging, the desired channels are always application dependent. An example of a non-visible multi-wavelength imaging application is two channel gas imaging in the infrared. A narrow-band image formed on a known absorption line of the molecular gas under observation would be acquired along with a second channel image at a slightly higher or lower frequency. The differential image of these two channels would then reveal the presence of the gas.

Each pixel within a visible light detector array assigns a value for a given spatial component of the scene entering the aperture, relative to the total intensity of light incident upon it. To achieve RGB colour imaging a standard digital imaging sensor such as a CCD or CMOS camera normally employs the use of a Bayer or similar multi-coloured filter to set groups of pixels for detection of a specific colour illumination bandwidth. On a high-resolution standard 18 Megapixel DSLR for example, this means 25% of the pixels detect red illumination, with a further 25% for blue and 50% detect green i.e. a 4.5/4.5/9 Megapixel breakdown in the number of detectors for each colour channel. Interpolation can then be employed to estimate the additional two colour channels not sampled on any given pixel so no reduction in resolution of the final image occurs. This is a form of simultaneous spectral imaging of the primary RGB colour channels which allows digital cameras to produce full colour images.

A unique difference in single pixel imaging systems is that the detection and

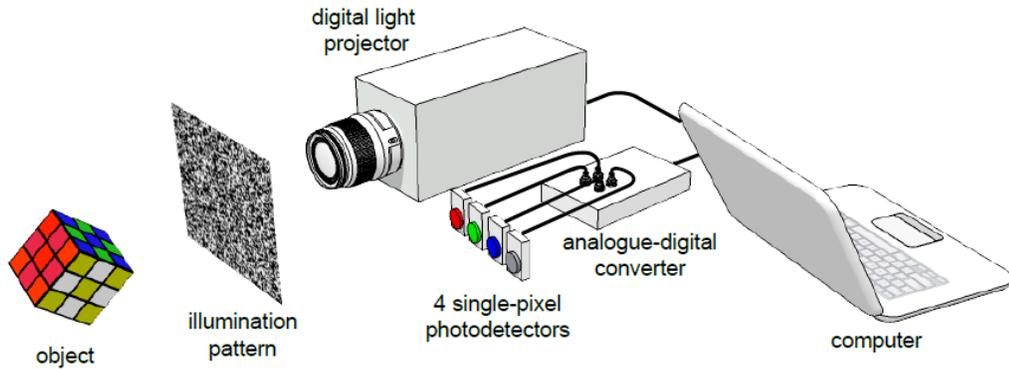


Figure 3.1: *Experimental setup of light commander based computational imaging system. A standard Rubik's cube puzzle toy is placed approximately 1m from the projector system, where the object was chosen for its multi-coloured faces. This system is capable of sequential or simultaneous multi-wavelength illumination allowing full colour scenes of the 3D object to be reconstructed in 2D.*

spatial resolution components of the imaging system have been ‘decoupled’. If a narrowband illumination source is chosen, or an appropriate filter is used in the case of broadband illumination, the detector’s measurements will be representative to the spectral response of the object for the chosen illumination bandwidth with the apparent illumination direction set by the detector position[38]. The weighted contribution for each individual pattern in any identical pattern set can be hugely disparate between the spectrally filtered and broadband case, when dealing with a multi-coloured object for example. This confirms that the differences measured in the PD come from the spatial components in the object which respond to white light vs single colour illumination. With the addition of multiple detectors, each filtered to a different illumination bandwidth, spectral imaging with single-pixel detectors becomes feasible.

In this experiment, we adapt a computational GI setup, as shown in Fig. 3.1, by replacing the SLM and laser with a digital light projector (DLP) (Texas Instruments Light Commander). The DLP contains a digital micro-mirror device (DMD) and four light emitting diodes: red (R), green (G), blue (B) and near-infrared (NIR), which can produce structured light fields. We utilize four different spectrally filtered single-pixel photo-detectors to measure the intensity of the reflected light at each of the emitted wavelengths, an analogue to digital converter to digitize the photo-detector signals and a computer to generate the random patterns and perform multi-wavelength reconstructions of the test object. In this demonstration, we utilize the built-in illumination sources of the DLP with frequencies in the extended visible band (400nm-850nm).

The patterns we use are randomly distributed binary patterns, having a fixed black and white ratio 1 : 1 which provides self-normalization in intensity output of the system. These fixed ratio patterns are then projected onto the object using a Nikon 50mm focal length lens. Four spatially separated single-pixel photo-detectors are positioned to monitor the back-scattered light. For every pattern projected, the corresponding object intensity is measured by each photo-detector, which is fed to a computer algorithm.

Normally during the data acquisition process, all external sources of light must be eliminated from the system since they would interfere with the recorded voltage from the SPDs, this requires ghost imaging to be done in the dark, however in this system we take advantage of the 24 bit-planes projected by the DLP, normally used for colour depth, to overcome such external noise sources. An OpenGL control program[39] ensured that every time the projector refreshed, a new pattern was displayed to maintain synchronisation within the system. Thus, the native 1024Hz binary projection speed of the structured

light mode of the Light commander could be utilised. The OpenGL program introduced flash planes of 100

Calibration of the Thorlabs DET36A visible Si detector was achieved by inserting a white diffuse sheet which is larger than the expected field of view. This produces a maximal signal for the detector when every pixel of the DMD is forced into an ‘on’ state. The amplification level of the PD should be set to the minimal setting for which there is no saturation in the measured signals. You can subsequently check that the smallest possible object is still detectable at this level by introducing a black sheet with a small cluster or area of white corresponding to approximately a single pixel in the final reconstruction. In practice for a static object, the minimal amplification setting is chosen to not produce saturated measurements for a given object and distance.

Figure 3.2 demonstrates a differential pattern projection sequence where each unique pattern is immediately followed by its inverse in a 24 bit-plane projection cycle. Such a 24 bit-plane cycle is normally associated with full colour RGB imaging where 8 bit-planes per colour channel combine to form the 24-bit image. By alternating between a binary pattern and its inverse in subsequent bit planes allows the demodulation of the measured signal at the frequency of the bit plane projection, which can isolate the back-reflected signal from light sources at other frequencies. It is possible to force 50% of the pattern to be black, and 50% to be white, while achieving a random distribution of pixels. This is accomplished within the Labview program which generates the patterns, based on a random number generator. The 50/50 constraint was introduced by randomly rearranging a 50/50 black and white static pattern, rather than randomly deciding whether an individual pixel was black or white as the output of the random number generator. Instead it is the index of the pixel which is randomized, producing a set distribution. Importantly, this use

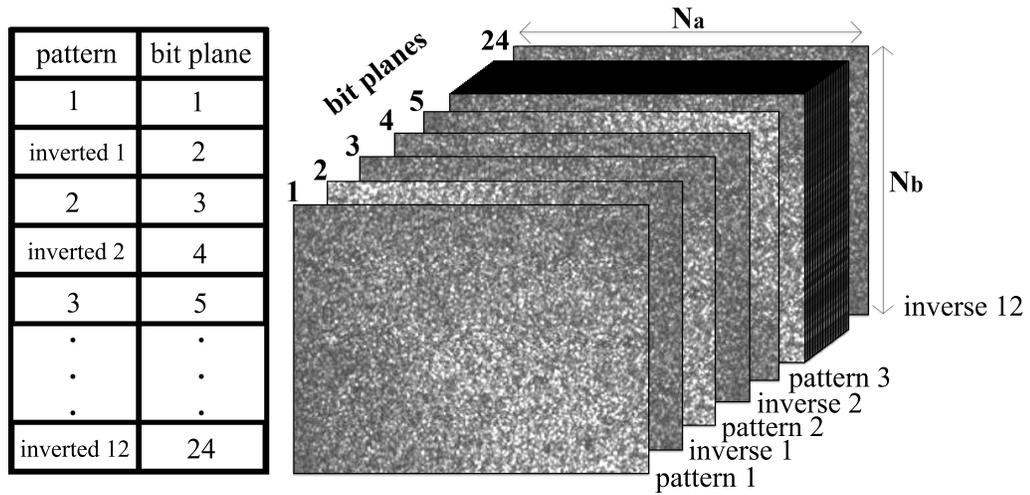


Figure 3.2: *Diagram of the pattern projection sequence constrained within a 24 bit-plane projection cycle. Each pattern is displayed followed by its inverse pattern, allowing a total of 12 unique patterns to be projected per cycle. The measured signal for each pattern and its inverse forms a differential signal which is paired with the corresponding unique pattern to allow image recovery. The differential pattern projection sequence allows isolation of only the signals corresponding to the patterns, and eliminating background light sources at other frequencies.*

of a unity black and white ratio pattern enables normalisation of the measured signals for each pattern, which has been shown to improve the SNR of the final reconstruction [18].

3.2 Simultaneous and sequential illumination

Several approaches can be undertaken to produce multi-wavelength single pixel reconstructions. The simplest approach is to perform time sequenced illumi-

nation of different wavelengths with an unfiltered SPD synchronized to the changes in the illumination sequence. The main advantage of this approach is that only a single detector is required, whereby the filter is simply changed for one corresponding to the next channel, leading to a simplistic and potentially cheaper system. Most standard DLP's project RGB illumination in sequence for the perception of white or full colour images to appear for use in video projection, this means that many DLP's are already producing illumination which is compatible with this approach and requires no advanced software or interface development of the DLP into the single pixel imaging system. This type of approach is however, always at a disadvantage in any application which the target object would be expected to exhibit temporal changes in position or spectral response with live biological samples being a primary example. Any sequential colour channel reconstruction could be excluded from any direct comparison between channels and thus this approach can only be applied in which the scene or object is known to exhibit no temporal changes on the full timescale between the first channel observed to the last. Additionally, the time required to complete a multi-wavelength reconstruction is a multiple of the number of desired channels and thus, is only practically applicable to situations which call for a reasonably small number of detectors.

Alternatively, another approach is to use simultaneous broadband illumination from all LEDs with multiple SPDs filtered for different frequencies. An individual detector must be assigned to each colour channel and be filtered to the desired bandwidth. This approach requires a more complex system in which data must be captured simultaneously from all detectors, and thus complexity will scale with the number of desired channels. The limitations of this technique are set by the number of available inputs which can be simultaneously read and processed by the computer, as well as by the optical topology which needs to enable multiple outputs for different regions of the spectrum. The

advantage to this approach is that any scene or object must only remain temporally stable for the duration of a single data acquisition period, in which a single set of unique patterns will be displayed. Additionally, extending this approach to include a vast array of filtered detectors responsive to a wide array of wavelengths would result in a hyper-spectral imaging system in which all information about an object or scene would be captured simultaneously, creating a complete ‘data cube’ for a single observation.

The approach that is described in this work is a hybrid of the two techniques, which uses time sequenced illumination, but with multiple filtered SPDs. This hybrid approach was chosen because it meant no hardware alteration of the DLP system was necessary to produce true white light, as the system naturally produces pseudo white light from sequential illumination of the RGB diodes. Since all investigations using this system only required a maximum of three detectors, the increase in time requirements for complete reconstructions was therefore not a hindering factor. This would triple the time needed for a single reconstruction however, this is acceptable when compared to the time requirement increases associated with an application which called for greater than 10 colour channels, for example. Multiple filtered detectors were used as opposed to a single detector since this meant no human intervention was necessary during ‘slow’ sequential illumination changes to replace the filter, or by extension, the assembly and programming of an automated filter wheel device.

An additional method of implementing the projection sequence in the hybrid method was that the system could be operated in a video projection mode producing Red, Green, Blue, Magenta, Cyan, Yellow (RGBMYC) (+Black and White) images, which could be projected at the native display rate. This would allow different pattern sets to be encoded within a single projected pattern with

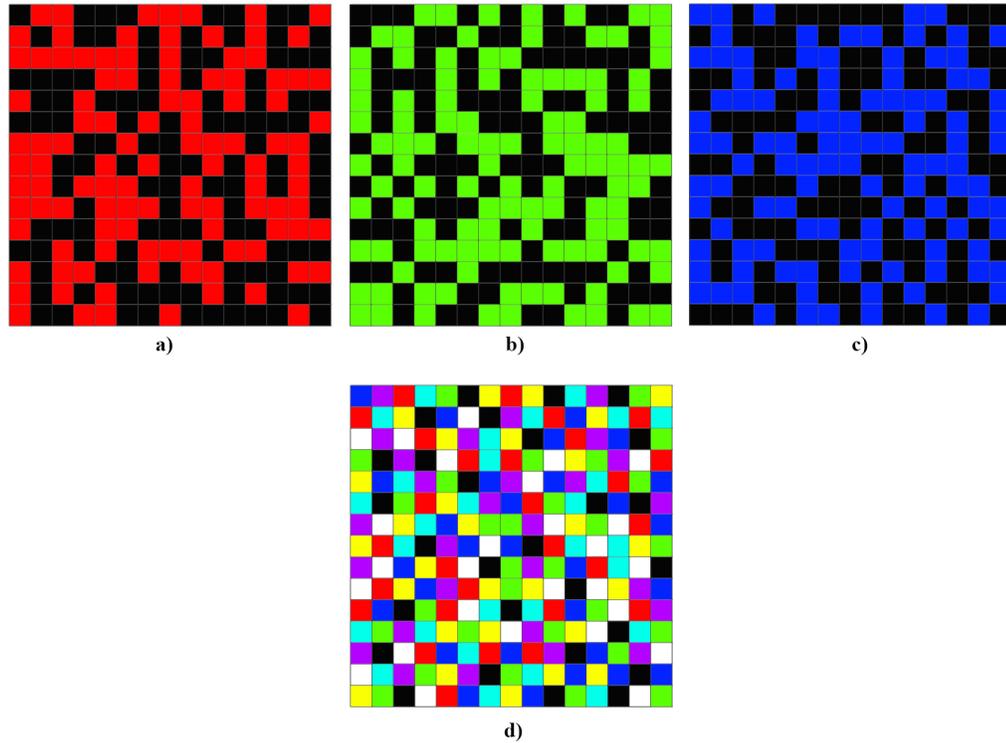


Figure 3.3: *A visual representation of the RGBMYC patterns produced when attempting sequential multi-wavelength imaging. a)-c) represent a single pattern produced for each colour channel where d) is the RGBMYC combination pattern when a)-c) are displayed quickly in sequence with 8 possible colour states depending on the breakdown of the colour combinations from each pattern. Each pattern is arbitrarily 15×15 in this example.*

spectral discrimination by filtered single-pixel detectors responding only to the filtered frequencies. This is an alternate method of image stacking and at the time of publication of the work described here, was an alternate method to utilize standard DLP's in a single pixel imaging system.

Fig.3.3 demonstrates the sequential pattern encoding process for use with the hybrid approach, which is operated in the native video projection mode of most

projector systems. This process produces patterns with 8 colour states which can be pre-recorded into a RGB colour sequence. This results in a ‘video’ sequence of RGBMYC patterns which could enable any standard 24 bit-plane 60Hz projector system to produce the encoded patterns needed for single-pixel reconstruction. Since most single-pixel camera systems of the time were constructed from any commercial digital projector operating in its native video mode, this created a ‘plug and play’ type system where a specific pattern set could be encoded as a video format file and run through any projector. For most SPC cameras this method was an alternate one only, and the motivation for implementing it derived from investigating the portability of SPC systems at the time. In contrast, in the results obtained in the main investigation into colour single-pixel imaging, all operation was undertaken in a specialist ‘structured light mode’ which is not found in standard projector systems. Structured light mode forces all bit planes to be equal in length and sets diodes to be active for the entire bit-plane cycle sequence, producing ‘white’ light for the entire sequence. In contrast, the video projection mode sets each colour diode on in a sequence covering the 8 bit-planes which correspond to a portion of the larger 24 bit-plane cycle. In summary, we are effectively only projecting one unique Red, Green and Blue pattern for every 24 bit-planes, with each pattern fixed ‘on’ for 8 of those bit-planes which correspond to the temporal illumination sequence of the RGB diodes. Therefore, we can only ever project at a native 60Hz video rate and each pattern projected is a RGBMYC pattern. Since the detectors are spectrally filtered, they ‘observe’ 8 bit-planes per 24 bit-planes of a single pattern, with the remaining 16 bit-planes returning no signal for a single spectrally filtered detector.

It should be noted that since the pattern projection rate was locked to 60Hz, there was no advanced synchronisation methods required to guarantee the correct assignment of signals to patterns. At this ‘slow’ rate, simply taking all

of the measurements obtained for one loop of the controlling Labview program, and making the assumption that all signals related to the current pattern the controlling program was instructing the DMD to project, would produce results with no apparent degradation.

3.3 Multi-wavelength single pixel imaging

3.3.1 RGB imaging

To investigate full colour single pixel imaging an object with a variety of colours, in this case a Rubiks cube, was chosen. We take advantage of the four colour illumination LEDs (R, G, B and NIR) for projecting binary patterns, and extract the signals from photo-detectors that have been filtered with bandpass coloured filters for each wavelength. These signals are used to reconstruct four separate images representing the four colour planes of reflected light measured. By overlaying the three R, G and B reconstructions we obtain a coloured combination image, identical to typical full colour imaging methods, as shown in Fig 3.4.

To produce the results in Fig. 3.4, the system was operated with a projection rate of 60Hz. Each pattern was illuminated by a single colour for $\frac{1}{3}$ of the total per pattern projection time, thus giving an equal share to R,G and B illuminations, creating a similar number of measurements per colour channel. The estimates of the SNR show a variance in each of the componenet color planes with the final RGB image taking an approximate average value in its estimate. Please refer to Appendix 1 for a detailed description of the thresholding levels used to make these estimates.

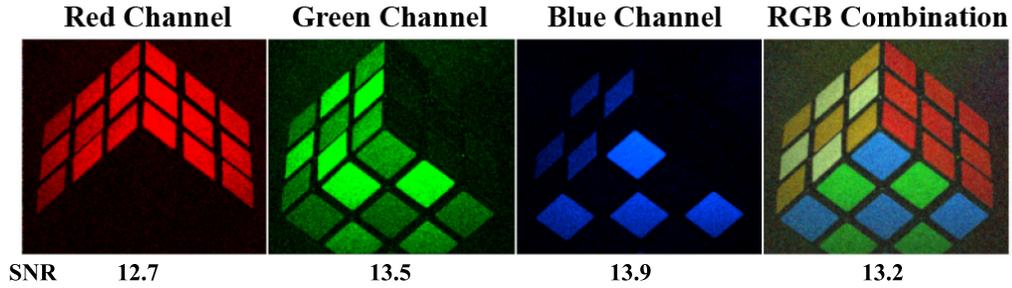


Figure 3.4: Reconstructions from data obtained from three spectrally filtered PD's for RGB colour channels from the system shown in Fig.3.1. The final image is a combination of the first three which gives the full colour RGB image of the Rubik's cube target object. These images were 128×128 (Nyquist=16384) are correspond to > 20 minutes of acquisition time with the projector running at the standard video rate of 60Hz. Estimates of the SNR in each colour plane and the RGB image were made.

3.3.2 Fluorescence imaging

Fluorescence is the emission of photons by certain minerals, dyes and paints when high energy incident light is absorbed by ground state electrons, re-emitted at a longer wavelength and lower energy when the excited electron returns to its ground state in a radiative decay. It is widely used in microscopy for a number of imaging applications [40, 41, 42].

Figure 3.5 describes the basic fluorescence process when under the correct chemical conditions, a material absorbs high energy incident light, producing a ground state electron excitation which moves through a two-step process back into its ground state. Firstly, the electron will transition into an allowed lower energy state, without radiating a photon. Secondly, the electron will transition back into a ground state, but will emit a photon of lower energy

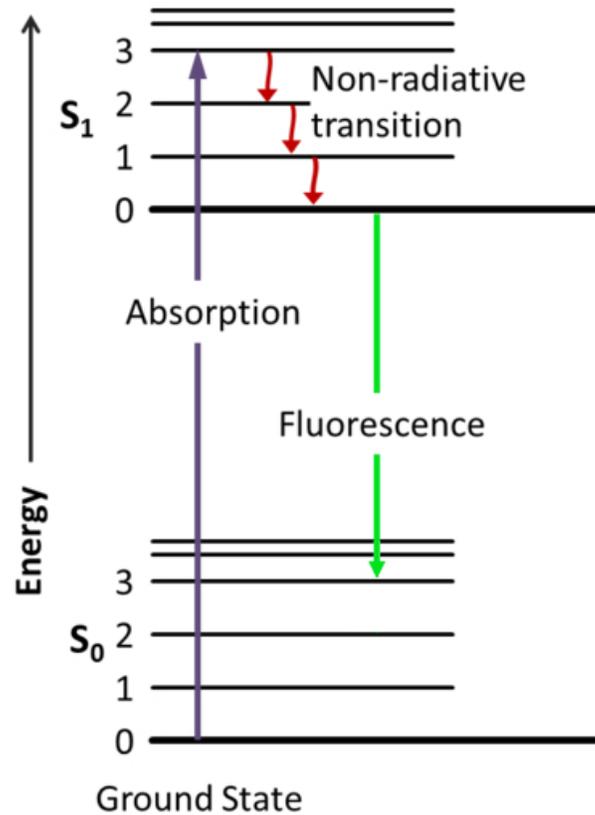


Figure 3.5: *Jablonski diagram detailing the absorbance, the non-radiative portion of the electron decay and the fluorescence as the electron returns to its ground state. Used with public domain permission.*

than the initial incident light. The lower energy photon is a result of some energy being transferred during the time the electron is in an excited state to kinetic energy, and is known as vibrational relaxation. The simple relationship between S_0 and S_1 can be shown as follows

$$S_0 + h\nu = S_1 \quad (3.1)$$

Where S_0 is the ground state of the fluorescent material, S_1 is the first excited state, $h\nu$ is the photon energy and h and ν are Planck's constant and the photon frequency respectively. A ratio to describe the efficiency of the fluorescence reaction, known as the quantum efficiency, Φ is then simply

$$\Phi = \frac{h\nu_E}{h\nu_A} \quad (3.2)$$

Where $h\nu_E$ is the number of emitted photons observed and $h\nu_A$ is the total number of absorbed photons. This well-known phenomenon is used extensively for a multitude of purposes including industrial testing techniques and colour enhancement in microscopy of biological samples. Normally visible light fluorescence, most commonly produced by ultraviolet stimulation, can be easily observed using CCD cameras. However, a Bayer, or other spectral filter within a commercially available CCD camera can have poor colour discrimination and mainly functions in the visible spectrum. Thus, a single pixel imaging approach may be able to offer some distinct advantages over traditional imaging methods constrained to poor colour discrimination filters, since narrowband filters can readily be applied to large area single element detectors. Additionally, it would be a simple procedure of changing detector type and subsequent filtration to alter the system for the infrared and lower energy fluorescence observation. The system described in Fig. 3.1 was thus used to investigate a fluorescent imaging application by imaging the Rubik's cube object while coated in paints known to fluoresce under blue light illumination.

Fluorescence sensitive single pixel reconstruction is achieved by precise selection of emission wavelength of the projector for a filtered detector to receive stimulated emission in a specific bandwidth, by pre-filtering the broadband emission with a narrowband filter. This has been shown in a classical (non-computational) GI setup [43]. To demonstrate fluorescence single pixel imag-

ing, certain sections of the test object have been painted with two different fluorescent inks: red and green, where these paints had similar fluorescence quantum efficiency, Φ . For this investigation, the projector was operated in a single illumination mode, such that it projects monochromatic, random binary patterns. It is subsequently possible to either reconstruct true colour images consisting of the components of an object which are solely reflecting light, or those components which are emitting light by fluorescence. This allows us to separate these ‘components’ or produce various combination images as any application required. In the various colour planes of Fig 3.6 we observe both reflection and fluorescence spectral components of the object. This is shown for example in the Green channel of Fig 3.6, where both sections of the test object which were painted with fluorescent paint and those which would naturally reflect light of the same wavelength as the illumination wavelength is observed, which was the first time this had been shown in such a system[44].

Figure 3.6 shows reconstruction in all observed colour planes, which correspond to results obtained from their respective spectrally filtered detectors, while fully maintaining a single monochromatic blue illumination source during data acquisition. The results from the green and red channels with a positive result demonstrate fluorescent light has been emitted by the coatings on each section since no other source of light was present. A ‘control’ section was arbitrarily selected to receive no coating of green fluorescent paint to demonstrate that under blue illumination, there is no natural response from the test object in a non-blue channel. This control confirmed that we were observing a positive reconstruction in that segment of the test object due to the fluoresce effect, and not for any other reason, such as broader than expected spectral response from the object.

The response from the red paint was substantially lower in intensity than that

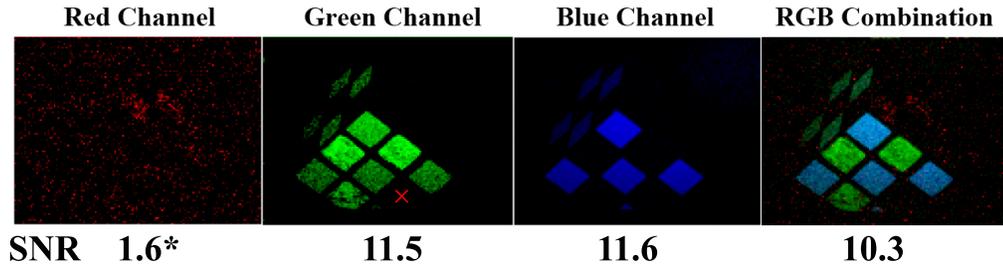


Figure 3.6: *Results of fluorescence imaging on the Rubik's cube test object with 265k iterations per colour channel. The illumination was set by activating only the blue LED while the target object was coated in various colours of florescent paint. In the test although only blue illumination was used, the florescent emission of green and to a small extent, the green and red filtered PDs observed red light. A combination image is shown with larger noise contributions in the red colour channel due to overall low intensity. Note that in the green channel image the 'missing' green square donated by the small red 'X' in that reconstruction was not coated in florescent paint, thus serving as a 'control' section.*

of the green paint and this is reflected by reduced SNR in the red channel. Any single-pixel imaging technique suffers from the effect that as signal intensity reaches comparable levels to intrinsic noise of the detector, the reconstruction will degrade until there is no visible reconstruction for a single colour channel in the case of multi-spectral imaging, this is an outcome of the separate normalisation process carried out on each individual colour channel. Global normalisation, whereby the maximal value over all colour channels is selected as the normalization factor, may improve specific examples of multi-channel reconstructions, where there is one poor channel compared to the others however, such a process would effectively nullify the poor channel, and overall noise would be reduced by simply applying the normalisations to similar qual-

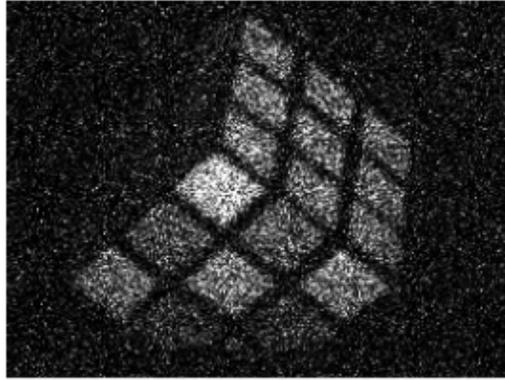
**SNR 5.2**

Figure 3.7: *Non-visible result of the Rubik's cube test object with 850nm illumination at 256×192 pixel resolution. This image was produced from a high sample number, approximately $5 \times$ Nyquist. The estimated SNR was lower than the individual visible colour channel results due to the lower overall spectral response of the object in the NIR region for the PD used.*

ity channels. This contrasts with traditional imaging methods where the image will simply appear darker as an object becomes 'fainter'. However, it is entirely possible to perform SNR analysis on reconstructions during the data acquisition process, thus a user defined SNR minimum can easily be set, below which a channel will simply be discarded, or a threshold value set to indicate no meaningful reconstruction was obtained.

3.3.3 Shortwave infrared imaging

The advantage offered by imaging with a single-pixel detector compared to a pixelated sensor when changing wavelength sensitivity is that only the expense of changing a single element detector is needed as opposed to an entire sensor with spatial resolution. Additionally, visible spectrum coated DMD's have an extended operational bandwidth when compared to most visible tuned projec-

tion or detection optical equipment, normally covering a region of $0.2\mu\text{m}-2\mu\text{m}$. Taking advantage of this broadband response to illumination of the DMD makes it an excellent choice to produce a broad spectrum or selective wavelength imaging system, with potential applications in hyper-spectral or absorption line imaging. Hyper-spectral imaging is simply any imaging device which can record many spectral components of a scene, generally simultaneously. Therefore, images produced from a hyper-spectral camera take the form of a 3D cube where each 2D plane represents a different measured frequency. In contrast, absorption line imaging is undertaken by having a camera system selectively measure a few frequencies of a material, usually a gas, where broadband light has transmitted through that material. Due to the properties of light interaction within mediums, certain wavelengths are readily absorbed into the substance. These frequencies correspond to that material's electron excitation states. Since specific wavelengths carry a discrete energy value, the energy values which matches the excitation energies of the material's electrons can be absorbed to a well-defined degree. This creates a 'fingerprint' like response when observing the transmitted light in a spectrometer, however in terms of imaging, it is then possible to select two image frequencies where one corresponds to an absorption wavelength, and the other does not. The differential image produced from these two images reveals the presence of any substance which may be colourless under normal illumination, but is highlighted by the large differential between the highly absorbing frequency, and the non-absorbing frequency.

The DLP used in this investigation comes with a fourth, non-visible near infrared (850nm) LED, an appropriate Near IR sensitive detector was added to demonstrate the proof of principle result that this system can image beyond the visible light spectrum. Using the same technique to obtain fluorescence images, the projector was placed in a single illumination mode and a sin-

gle photo-detector equipped with a NIR spectral filter was used to measure the reflected light from each projected binary random pattern only allowing light from the NIR diode to be transmitted to the detector. Fig.3.7 is only a single channel non-visible result, but demonstrates the ability to image in a non-visible region for which the DMD has an acceptable response to illumination. This suggests a potentially lower cost and ‘tuneable’ alternative to current non-visible imaging devices. DMD devices can range between \$500-\$8000 depending on operational speed, however it is estimated that current NIR cameras scale at a rate of \$50000 per megapixel (a 320×260 SWIR camera is currently approx \$14000), whereas a single-pixel imaging device, based on a DMD does not have this large cost scaling factor associated with it. Higher resolution DMD’s stay within the current price bracket whereas standard non-visible imaging cameras very quickly reach untenable costs for increases in resolution quality only.

3.3.4 Fast multi-wavelength imaging

The setup used in this experiment is shown in Fig. 3.8. The DLP provides spatially incoherent light which is imaged to the plane of the scene with a 55mm lens and is in effect an update to the system described in Fig. 3.1. For each incident pattern, the total intensity reflected from the scene is directed onto a composite dichroic beam splitter (X-Cube) using a large collection lens. The dichroic beam splitter is used to spectrally filter red, green and blue light towards different outputs and allow subsequent measurement on three unfiltered single-pixel photo-detectors. The computer records all three photodetector signals via an analogue-to-digital converter, providing image reconstruction using an appropriate algorithm. This is chosen based primarily on the desired resolution of the final image, high resolutions above 128×128 were not

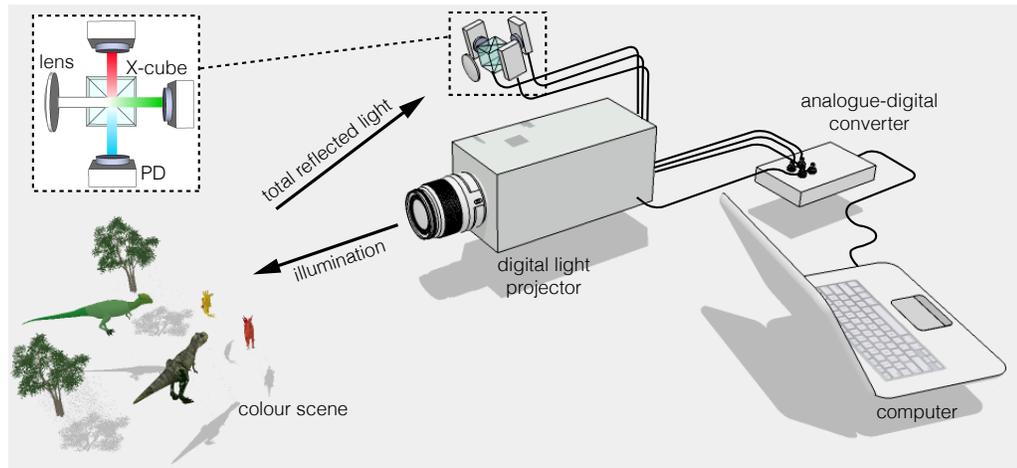


Figure 3.8: *Experimental setup of light commander based computational imaging system. A collection of small scale dinosaurs are placed approximately 1m from the projector system. This system is capable of sequential multi-wavelength illumination allowing full colour scenes to be reconstructed.*

suites for matrix inverse or compressed sensing based reconstructions due to the high computational load and ram requirements for storing such large matrices. Thus, in this case the standard iterative algorithm was used. As we require only binary illumination patterns, we operate the projector with all three LEDs permanently on and can take advantage of the colour planes to display 24 different binary images per frame, working again in the structured light mode of the DLP.

A standard digital light projector produces 24-bit colour depth images at a typical frame rate of 60Hz, meaning that the pattern display rate is equivalent to 1440Hz. Therefore, to reconstruct images by raster scanning or basis scanning with Hadamard matrices, the acquisition time and hence maximum achievable frame rate are summarised by Fig. 3.9.

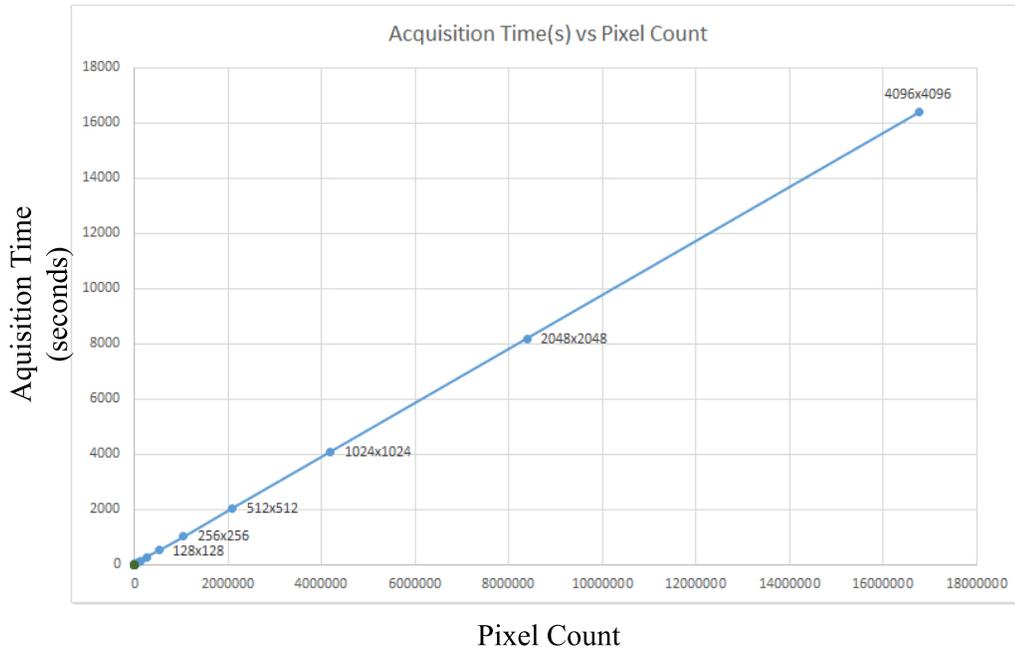


Figure 3.9: Graph demonstrating the linear relationship between the acquisition time in seconds versus the total pixel count of the final reconstructed image. This assumes 100 percent of Nyquist sampling and a constant pattern projection rate, 1440Hz in this case to reflect the rate of a standard 24 bit-plane commercial DLP. Common square resolution values have been highlighted on the scale.

However, DMD technology is now capable of achieving display rates more than 20kHz, significantly reducing the acquisition time for imaging with a single-pixel detector. This opens the possibility of reconstructing images faster than one second, allowing video or near-video real time reconstructions to potentially be made.



Figure 3.10: (a) Full-colour iterative reconstruction of a 3D scene, obtained by combining the separate reconstructions derived from three photo-detectors sensitive to (b) red, (c) green and (d) blue colour channels. The illumination patterns and final reconstruction utilized a 256×128 pixel resolution and required $> 20 \times$ Nyquist number of measurements. The 3D scene had dimensions of approximately $20\text{cm} \times 10\text{cm}$ and was located approximately 1m from the digital light projector.

2D scene reconstruction

Figure 3.10 shows a high quality full-colour reconstruction along with the individual colour channels which form the combination image. This real-world 3D scene here is represented in 2D in a similar manner consistent with a standard Bayer filtered CCD camera albeit at the lower resolution of 256×128 with the main difference being that in a Bayer filtered system, pixel interpolation is

always incorporated to give final full colour results the full stated resolution of the sensor. In contrast, here each colour channel has an equal resolution to the final reconstructed image using the single-pixel, full-colour combination reconstruction, a subtle but important difference in the image processing requirements between the two imaging methods. The 256×128 resolution balances expected levels of detail within the final image with acquisition time set by the maximum display rate of $1440Hz$. At the time of publication of this image, this was the first simultaneous real world multi-wavelength image of a scene by a single pixel/ghost imaging camera. This was a shift in reconstruction fidelity, approaching what could be described as single pixel ‘photography’ away from the binary transmissive objects of past experiments.

Compressed results

Compressive sensing applies well to the obtained data used to produce Figure 3.10 as acquisition time can become a major detrimental factor to the hybrid method employed for this investigation (See Chapter 2.5 for CS algorithm details). If one wishes to readily extend the same technique either into a non-visible region or have additional channels for higher colour discrimination, additional time for each detector added will have to be accounted for. Compressive sensing can therefore allow for more than twice the number of detectors for the same time investment with far greater apparent image quality over the same number of iterations.

Figure 3.11 Demonstrates compressive and non-compressive results for an increasing number of sampled encoded patterns. A final, massively over sampled iterative result is displayed to indicate approximately the number of samples required to match the compressive result of 8000 samples. From the sequence of images shown for this resolution, images corresponding to greater than 50%

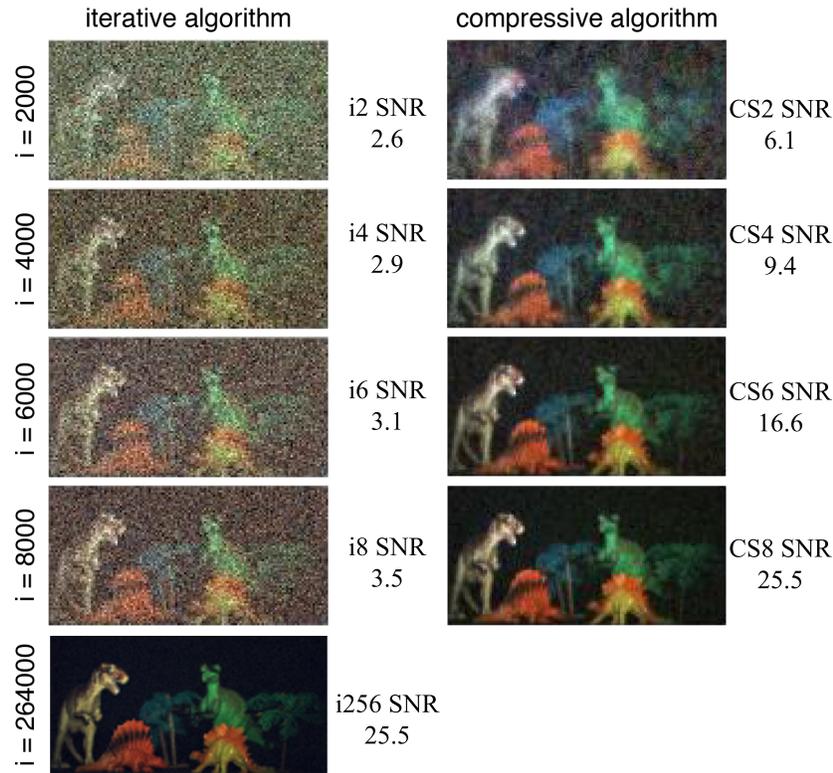


Figure 3.11: *Comparison of coloured image reconstructions with 128×64 pixel resolution for increasing sample size using an iterative and a compressive algorithm. For 8000 (approximately Nyquist) measurements the acquisition time is less than 12 seconds. SNR Estimates show CS vastly out performing iterative algorithms for all measurement levels.*

of the Nyquist limit shows faithful image reconstruction. This is contrast to the iterative approach in which even at Nyquist, image quality is still poor.

An advantage of compressive sensing as applied to single-pixel image reconstruction is that there is no pre-set ‘special’ requirements involved with applying the compressive sampling algorithm to recorded data. Any previous single-pixel dataset is applicable to be reconstructed via a compressed sensing method as opposed to basic iterative reconstruction. While old datasets

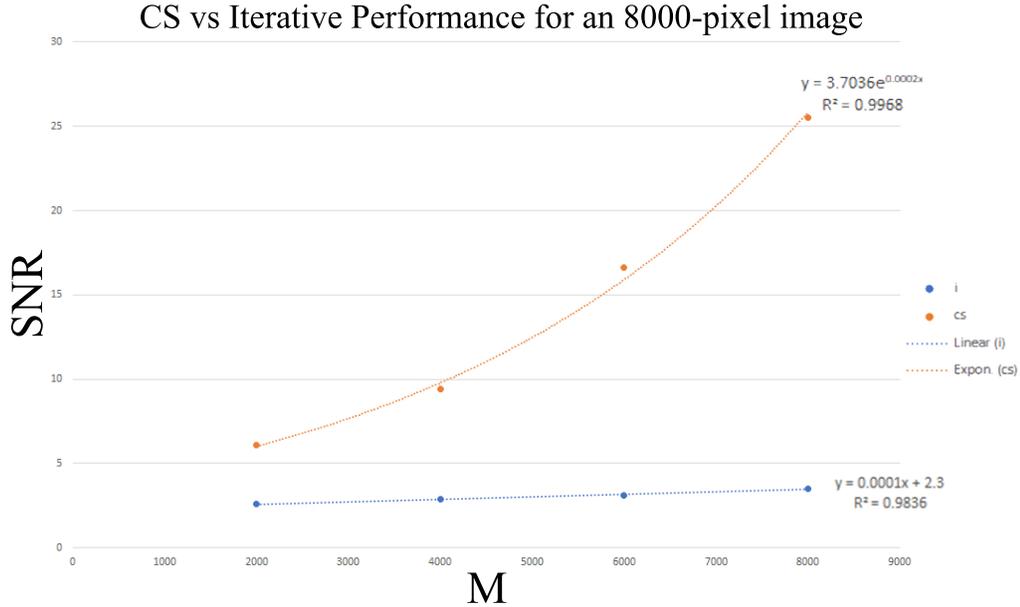


Figure 3.12: *Compressive vs Iterative algorithm performance comparison for an 8000 pixel final reconstruction. Iterative reconstructions have an extremely linear response with increasing measurement number while there is an exponential response from CS in like for like comparisons to iterative reconstructions.*

will most likely have taken the time to fully Nyquist sample the object, the required information needed to produce a result can always be reduced. Thus, new datasets can completely omit a fraction of the needed samples, being an example of sub-Nyquist sampling, saving on both time taken during data acquisition and on the total information required to be retained in producing reconstructions.

Figure 3.12 shows the performance comparison of CS and iterative reconstruction for the 8000-pixel dino image shown in Fig. 3.11. This plots the total number of measurements used to acquire the reconstruction, M , vs the SNR estimate of that reconstruction. While these are SNR estimates made using

the ROI threshold method described in Chapter 2, when appropriate trend lines are added it is possible to observe a very good linear fit with the iterative algorithm. Doing the same for CS, an exponential increase in performance with increasing M is observed. Note that it is not possible to supply the CS algorithm with 256k weighted patterns due to the nature of the matrix inverse process required to perform CS. This is a limitation of current computer memory capacity. However, we obtain near identical SNR estimates for the 8000 pattern CS reconstruction and the 256k iterative reconstruction, demonstrating the power of the CS algorithm to effectively suppress the background ‘floor’ of the reconstructions.

In addition to estimating the SNR, the presence of a reference image also allows the calculation of the Structural Similarity Index (SSIM)[45] of each pixel. SSIM is an algorithm applied to images or video, that attempts to quantify a value for ‘quality’ as is normally only determined by a human observer, it uses three comparison metrics: luminance, contrast, and structure.

Figure 3.13 is the visual representation of the SSIM calculation for the selection of iterative and compressive results chosen in Fig A.3. The brightness of each pixel in an SSIM image, taking a positive value between 0 and 255 represents the similarity in that pixel between the supplied and reference images. It is an interesting result as this metric demonstrates clearly where the real power of compressed sensing lies; The ability to drastically reduce the noise floor, compared to the relative peak intensity. Looking at the histogram result in Fig A.3 for the 8k compressive and 8k iterative results, we see a striking similarity in the signal distribution, however, the noise floor in the 8k iterative result is still such that all background pixels in the 8k iterative SSIM image score almost 0. However, the background level is much closer in even the 2k compressed sensing SSIM image to the background of the reference 264k iterative result. By

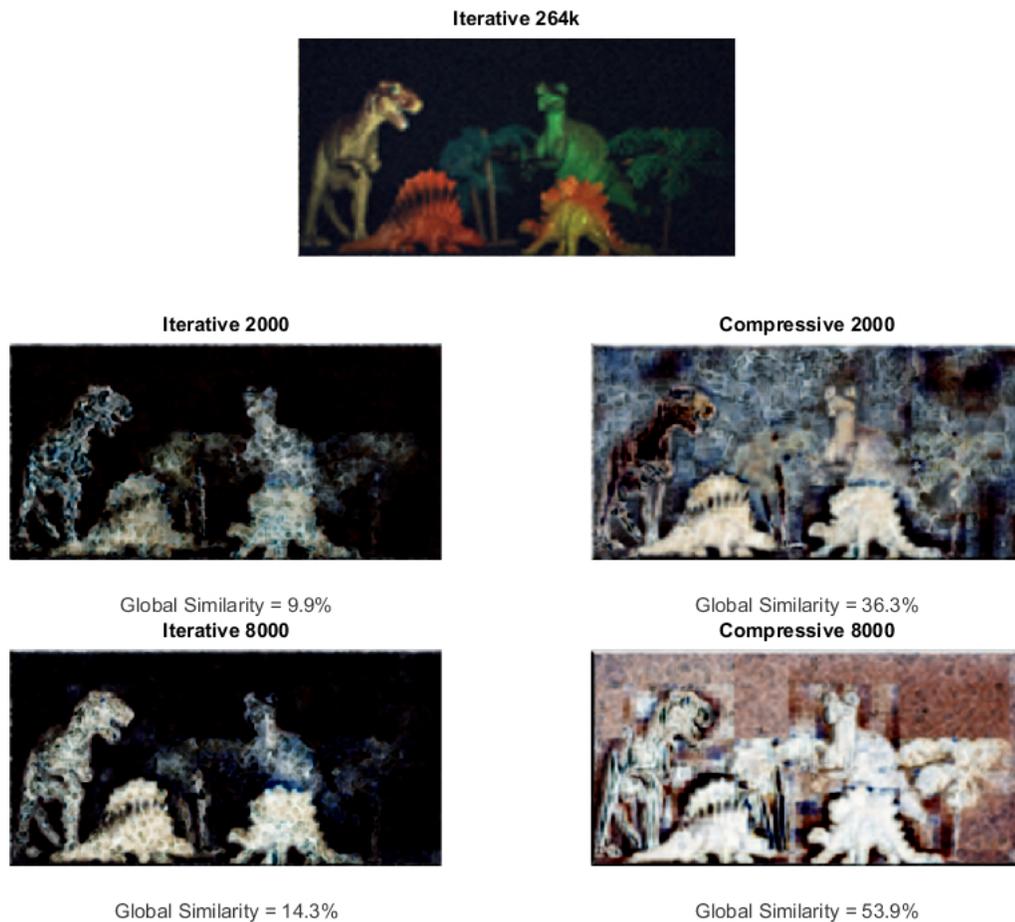


Figure 3.13: *SSI calculations for 2 iterative and 2 compressive results. The Global Index value is an averaged value which states, as a percentage the relative similarity between image pixels and reference image pixels overall.*

8k measurements in the CS result, we have a very good overall SSI score. This visual metric demonstrates the convergence in iterative results, and the rapid convergence in CS, which is due to the rapid lowering of background noise instead of higher signal pixels.

3.4 Discussion

It has been shown that computational imaging with three single-pixel detectors can be used to produce full-colour images of large scenes in just a few seconds of data acquisition. The use of a digital light projector and a suitable computer algorithm allows rapid structured illumination and hence short acquisition times, showing promise for a range of alternative imaging applications.

Additionally, by employing compressive sensing techniques, significant improvement to the image quality is observed when the number of measurements is far below the Nyquist limit. An advantage of compressed sensing is that no ‘active’ process is mandatory for any single pixel imaging system to take data compressively, indeed any non-compressive previously recorded set of encoded patterns and signals can be processed compressively to reconstruct the sub-Nyquist reconstruction. The advantage however of knowing that one is performing compressed sensing while acquiring data is a reduction in the required time to obtain all data needed for a reconstruction.

The low cost of additional photo-detectors and the large operating bandwidth afforded by DMD technology opens a range of alternative imaging solutions, such as hyper-spectral imaging and particularly for imaging at wavelengths where CCD or CMOS imaging technology is limited. Non-visible commercial cameras focus on small bandwidths which have some commercial interest, but there also exists some regions of the infrared for which readily available commercial cameras are not suited or difficult to operate in. This single pixel approach can operate over all visible and non-visible bandwidths with only the cost of a single detector per total number of desired colour channels compared to the cost of replacing a sensor with high spatial resolution. This holds true for any region of the EM spectrum in which it is readily possible to spatially

shape light fields.

There is a balance to be struck between acquisition time, and desired resolution. Depending on the specific object, temporal information may not be relevant, so many colour channels can be employed at high resolution without concern for the result. However, the resolution and number of channels is an easily scalable feature of the single-pixel imaging technique. For any given application, sequentially higher resolution runs of the system can be undertaken to discover the appropriate level of detail required for a given inherent noise value. The final resolution can be set accordingly, up to the maximum resolution of the DMD itself, which in this case was 1 Megapixel.

Chapter 4

Polarisation sensitive single-pixel imaging

4.1 Introduction to polarisation

A polarisation is the orientation in which an oscillating wave will ‘trace’ a locus of points for a fixed position in space, over some period of time [46, 47, 48]. Any oscillating wave will have a polarisation state and this state can be affected by interacting with a medium, or by scattering on a surface. How an object interacts with polarised light is a direct method of quantifying certain object properties related to the surface, type of material and even internal refractive index. However, polarisation is an often-underexploited property of light due to the moderate difficulty in polarisation applications which require spatial resolution of the states.

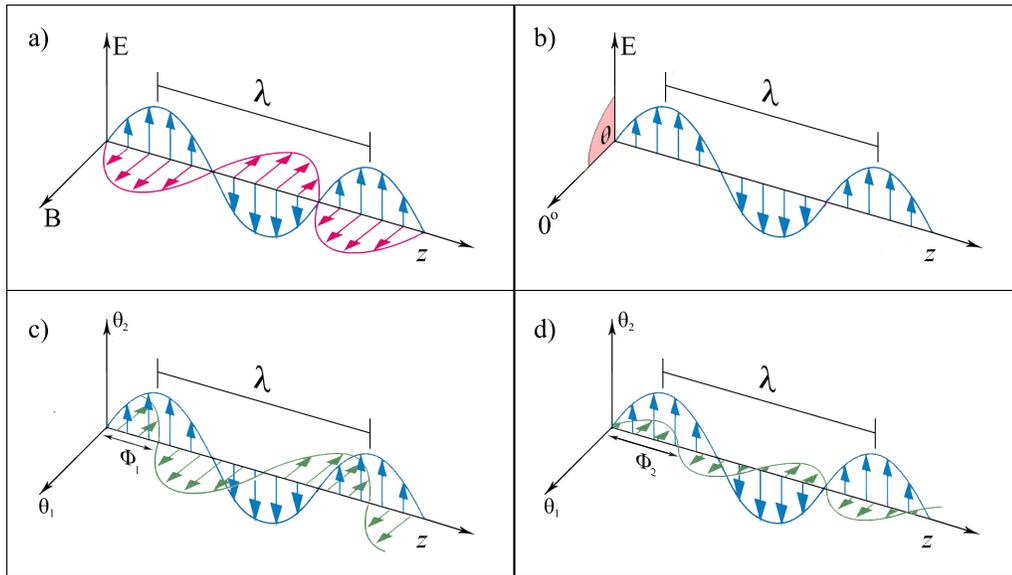


Figure 4.1: a) An electromagnetic field propagating in the direction z , with orthogonal electric and magnetic field components, denoted E and B respectively. b) The linear polarisation state of a light wave, with arbitrarily defined 0° relative angle. c) Two orthogonal waves in a circularly polarised state, with phase difference $\phi_1 = 90^\circ$ and the amplitudes of the waves are equal. d) Two orthogonal waves with a phase difference $\phi_2 \neq 90^\circ$.

4.1.1 Polarisation states

Figure 4.1 demonstrates a series of planewaves propagating along the z -axis in a range of different polarisation states. Figure 4.1 a) describes the electric field, E and magnetic field, B , components of an electromagnetic (EM) wave of light. E and B are always orthogonal to each other and the oscillations are intrinsically in phase. This orthogonality arises from the classical result of the planewave solution to Maxwell's equations where $E \bullet B = 0$, and the E and B fields can be written as

$$\vec{E}_x(z, t)\hat{j} = E_0 \sin(kz - \omega t + \varphi)\hat{j} \quad (4.1)$$

$$i\vec{B}_y(z, t)\hat{k} = B_0 \sin(kz - \omega t + \varphi)\hat{k} \quad (4.2)$$

Where $\vec{E}_x(z, t)$ and $\vec{B}_y(z, t)$ are the magnitudes of the planewaves at any given point and time, $k = \frac{2\pi n}{\lambda}$ is the wavenumber, \hat{k} and \hat{j} are standard basis vectors which satisfy a right hand coordinate system (the type of system applicable to the examples in Fig. 4.1). $\omega = 2\pi f$ is the angular frequency, t a given point in time and φ is the phase shift of the wave.

By convention, we will omit the polarisation state of the B field component and focus on the orientation of the E field component only when describing the ‘polarisation’ state of any EM planewave. Fig 4.1 b) demonstrates the linear polarisation state of an individual EM planewave propagating in the z-axis. Here an orthogonal 0° angle is defined. It is often useful, and a simplification of the relationship, to define such an angle as always perpendicular to the propagation of one of the EM planewaves, however, an arbitrary angle is also valid. This section in effect describes the propagation of a single EM planewave, however in nature, a light source with generally emit many such waves. These waves may be emitted in a single orientation or many. It is therefore the net combined E field of interacting waves which determines the resultant net polarisation state. A two wave system can be summarised in a column vector format as follows, and taking the form $e^{i\theta}$

$$\vec{E}(z, t) = \begin{pmatrix} e_x \\ e_y \\ 0 \end{pmatrix} e^{i(kz - \omega t + \phi)}$$

Where $\phi = \varphi_1 - \varphi_2$ is the phase difference between the two waves. While

a single individual wave will have some linear polarisation state, when there are two interfering waves, linear polarisation occurs when $\phi = 0$, and the orientation is dependent solely on the combined $\vec{E}(z, t)$ field component. Light from a source is said to be fully linearly polarised when all EM planewaves have no phase difference between them, and take a single net orientation. See Fig. 4.2 a) for the linear polarisation ‘trace example. When $\phi \neq 0$ the resultant trace for a fixed position in the z-axis over a period of time will take the form of an ellipse with semi-major and minor axes, A and B , as shown in Fig. 4.2 b).

Figure 4.1 c) describes the situation in which two orthogonal planewaves have a phase difference of $\phi = \pm\frac{\pi}{2}$. This ‘special case’ results in circularly polarised light. If $\phi = +\frac{\pi}{2}$, the orientation of the $\vec{E}(z, t)$ field component traces a circle on any fixed position in the z-axis in a counter clockwise direction. This is referred to as Left-Handed Circular Polarisation (LHCP). Conversely if $\phi = -\frac{\pi}{2}$, the circle will be traced in a clockwise fashion and is referred to as Right-Handed Circular Polarisation (RHCP), as seen in Fig 4.2 d). Figure 4.1 d) demonstrates two orthogonal planewaves for which $\phi \neq 0$ and $\phi \neq \pm\frac{\pi}{2}$. This is the general form of polarisation as any expression concerning the ‘polarisation ellipse’ can be used in the linear or circular polarisation case. The ellipsoid trace can be observed in Fig. 4.2 d).

The resultant ellipse drawn out by a propagating pair of waves can be used to fully describe the geometric properties of linear or non-linear polarisation states of a system. Fig 4.2 a)-d) describe all possible net polarisation states light can take. Therefore, parameterising polarisation by the magnitudes of the $|E|$ field components, the orientation relative to the y-axis, and the ‘ellipticity’ in the case of elliptical polarisation, fully describes the polarisation state. To summarize this parameterisation, the equation of an ellipse in its

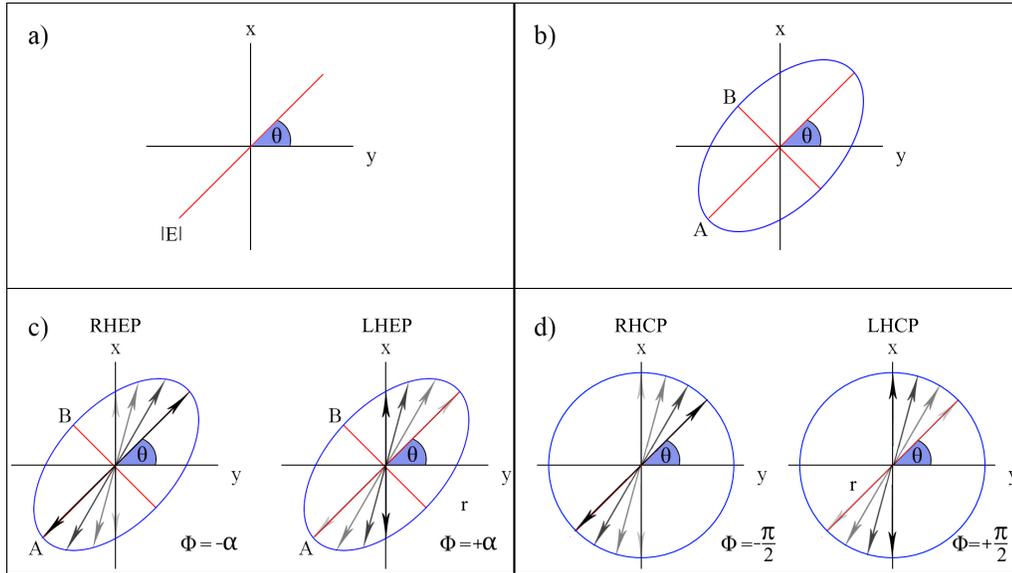


Figure 4.2: a) At a fixed point on the z -axis from Fig. 4.1 of a linear polarisation state, with angle θ to the y -axis. The size of oscillation is the magnitude of the electric field component $|E|$. b) Two out of phase waves will trace out an ellipsoid on the fixed point of the z -axis, with semi-major and minor axes, A and B . c) Negative values for ϕ cause Right-Handed Elliptical Polarisation (RHEP), conversely positive values give rise to Left-Handed Elliptical Polarisation (LHEP) with the oscillation tracing the ellipsoid in a counter clockwise and clockwise direction, respectively.

standard form (aligned to the xy -axes) is

$$\frac{E_{0x}^2}{A^2} + \frac{E_{0y}^2}{B^2} = 1 \quad (4.3)$$

Where E_{0x} and E_{0y} are the electric field components wrt to the xy coordinate system. However, for any given elliptical polarisation state, an ellipse with any orientation is possible. Therefore, from Eqn. 4.2 it is possible to derive [49] the equation of an ellipse in its non-standard form given by

$$\frac{E_x(z, t)^2}{E_{0x}^2} + \frac{E_y(z, t)^2}{E_{0y}^2} - \frac{2E_x(z, t)E_y(z, t)}{E_{0x}E_{0y}} \cos\phi = \sin^2\phi \quad (4.4)$$

Furthermore, an orientation angle, θ which is the angle the semi-major axis makes with the x-axis can be further derived from Eqn. 4.4 using both algebraic calculations, and trigonometric identities to give

$$\tan(2\theta) = \tan(2\alpha)\cos(\phi) \quad (4.5)$$

Where $\tan(\alpha) = \frac{E_{0x}}{E_{0y}}$ is known as the ‘auxiliary’ angle. The final parameter to define is the ellipticity angle, ϵ , which determines how far from a circle or line the ellipsoid trace will take. Further work on Eqn. 4.4 yields an expression for the equation of an ellipse in terms of ϵ as follows

$$\tan(2\epsilon) = \sin(2\theta)\tan(\phi) \quad (4.6)$$

4.1.2 Jones calculus

Jones calculus [50] is a matrix representation of polarisation which is useful for describing the individual states in a system, but also how those states change when interacting with optical components. Jones Vector and Matrix polarisation only deals with fully polarised or unpolarised states. From Eqn. 4.2 in a column vector format, and using $\tau = (kz - \omega t)$ as the wave phase, then the Jones column vector for the E field is

$$\vec{E} = \begin{pmatrix} E_x \\ E_y \end{pmatrix} = \begin{pmatrix} E_{0x}e^{i\tau_x} \\ E_{0y}e^{i\tau_y} \end{pmatrix} \quad (4.7)$$

Where E_{0x} and E_{0y} are the amplitude components and $e^{i\tau}$ is the phase component. Generally, we want to normalize the Jones vector. For a complex

number, this means that the dot product between the number and its complex conjugate is one. If we set the orientation of the Jones vector to be horizontal and vertical we can represent linear horizontal, vertical, -45° and $+45^\circ$ polarisation Jones vectors as follows

$$\vec{E}_H = \begin{pmatrix} E_x(t) \\ 0 \end{pmatrix} = \begin{pmatrix} E_{0x}e^{i\tau_x} \\ 0 \end{pmatrix} = \begin{pmatrix} N \\ 0 \end{pmatrix} = N \begin{pmatrix} 1 \\ 0 \end{pmatrix} \quad (4.8)$$

$$\vec{E}_V = \begin{pmatrix} E_y(t) \\ 0 \end{pmatrix} = \begin{pmatrix} 0 \\ E_{0y}e^{i\tau_y} \end{pmatrix} = \begin{pmatrix} 0 \\ N \end{pmatrix} = N \begin{pmatrix} 0 \\ 1 \end{pmatrix} \quad (4.9)$$

$$\vec{E}_{-45} = \begin{pmatrix} E_x(t) \\ 0 \end{pmatrix} = \begin{pmatrix} E_{0x}e^{i\tau_x} \\ E_{0y}e^{i\tau_y} \end{pmatrix} = \begin{pmatrix} N \\ N \end{pmatrix} = N \begin{pmatrix} 1 \\ 1 \end{pmatrix} \quad (4.10)$$

$$\vec{E}_{+45} = \begin{pmatrix} E_y(t) \\ 0 \end{pmatrix} = \begin{pmatrix} E_{0x}e^{i\tau_x} \\ E_{0y}e^{i\tau_y} \end{pmatrix} = \begin{pmatrix} N \\ N \end{pmatrix} = N \begin{pmatrix} 1 \\ 1 \end{pmatrix} \quad (4.11)$$

Where $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$ and $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$ are the normalized Jones vectors and N is a normalization factor.

In the case of circular polarisation, there is a $\phi = \pm\frac{\pi}{2}$ phase difference, giving RHCP for negative and LHCP for positive. These two polarisation states can also be represented by Jones Vectors

$$\vec{E}_{LHCP} = \begin{pmatrix} E_{0x}e^{i\tau_x} \\ E_{0y}e^{i\tau_y} \end{pmatrix} = \begin{pmatrix} N \\ Ne^{i\frac{\pi}{2}} \end{pmatrix} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ i \end{pmatrix} \quad (4.12)$$

$$\vec{E}_{RHCP} = \begin{pmatrix} E_{0x}e^{i\tau_x} \\ E_{0y}e^{i\tau_y} \end{pmatrix} = \begin{pmatrix} N \\ Ne^{-i\frac{\pi}{2}} \end{pmatrix} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ -i \end{pmatrix} \quad (4.13)$$

For a general ellipse with semi-major and minor axes, A and B respectively. The Jones vector is

$$\vec{E}_o = \begin{pmatrix} E_{ox} \\ E_{oy}e^{i\phi} \end{pmatrix} = \begin{pmatrix} A \\ B(\cos\phi + i\sin\phi) \end{pmatrix} \quad (4.14)$$

Using the identity $e^{i\phi} = \cos\phi + i\sin\phi$.

Optical components can also be represented by a Jones Matrix, and matrix operations between these optical components and incident polarised light can determine the resultant polarisation state after propagation through the optical component.

To derive the Jones matrix for a general optical device we can consider the Jones vector for a beam of light

$$\vec{E}_i = \begin{pmatrix} E_{ix} \\ E_{iy} \end{pmatrix} \quad (4.15)$$

And then consider the same beam of light after propagation through an optical device

$$\vec{E}_t = \begin{pmatrix} E_{tx} \\ E_{ty} \end{pmatrix} \quad (4.16)$$

These two vectors can be fully described by a system of linear equations with coefficients a,b,c,d.

$$E_{tx} = aE_{ix} + bE_{iy} \quad (4.17)$$

$$E_{ty} = cE_{ix} + dE_{iy} \quad (4.18)$$

Satisfying $\vec{E}_t = \vec{J}\vec{E}_i$ and rewriting in matrix form gives

$$\vec{J} = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$$

Table 4.1 demonstrates the most common Jones Matrices for physical optical devices. This is useful when summarising any system dealing with polarisation components, as it allows potentially complex systems to be reduced to a simple

Table 4.1: Jones Matrices

VLP		$\begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}$			RHCP	$\frac{1}{2}$	$\begin{bmatrix} 1 & i \\ -i & 1 \end{bmatrix}$
HLP		$\begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}$			LHCP	$\frac{1}{2}$	$\begin{bmatrix} 1 & -i \\ i & 1 \end{bmatrix}$
+45	$\frac{1}{2}$	$\begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}$			QWP Fast V	$e^{\frac{i\pi}{4}}$	$\begin{bmatrix} 1 & 0 \\ 0 & -i \end{bmatrix}$
					QWP Fast H	$e^{\frac{i\pi}{4}}$	$\begin{bmatrix} 1 & 0 \\ 0 & i \end{bmatrix}$

matrix per component. The effect of the light through the system can then be readily determined in simulation. The Jones matrix method is unique in its dealing of the instantaneous electric field, making it the preferred choice when using coherent light sources such as lasers. Other methods described in the proceeding section are chosen then in cases of randomly and quickly altering polarization states, which can be found in natural light sources.

4.1.3 Stokes' Parameters

Another useful characterization method for polarisation is the four Stokes' parameters [51]. These can be defined by a number of notation methods, namely S_0, S_1, S_2, S_3 or I, Q, U, V and even S_1, S_2, S_3, S_4 which can lead to some confusion. Therefore, for the purposes of this discussion we will constrain to use the I, Q, U, V convention in this section, but will use the S_0, S_1, S_2, S_3 in future references as this is the standard accepted convention. These parameters are extremely useful in practical polarisation applications because they can be easily measured and unlike the Jone's vector, can deal with partial polarised states, which is common in many sources.

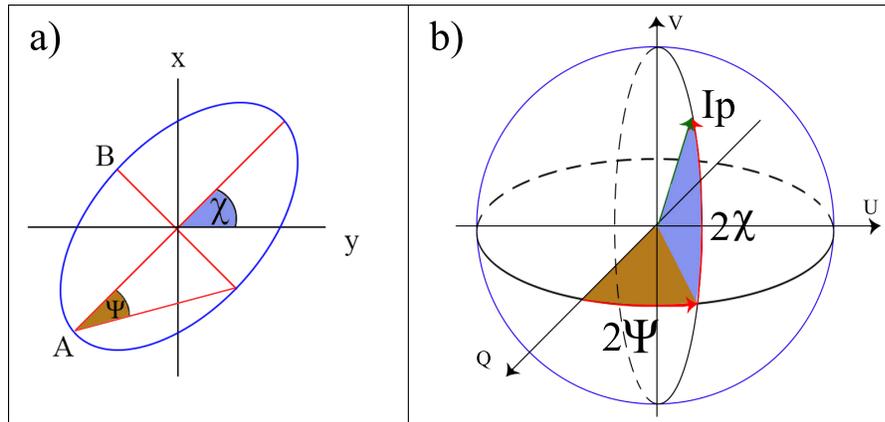


Figure 4.3: a) *The Polarisation Ellipse with defined orientation angle, χ and ellipticity angle Ψ , these angles are chosen by convention when one wishes to project angles into spherical coordinates.* b) *Poincaré sphere for the Stokes' parameters. Along the horizontal we have various orientations of linear polarisation, as one selects a vector which moves towards the 'poles' of the sphere, the light becomes non-linear with an increasing ellipticity in the resulting polarisation ellipse. At the pole, the light is circularly polarised with LHCP or RHCP if the vector is pointing to the north or south pole, respectively.*

Figure 4.3 a) Is a diagram of the well-established polarisation ellipse, the general case for a state of polarisation. Here we have defined the orientation angle as χ , which is still the angle with which the semi-major axis makes with the y axis, and the ellipticity angle, ψ which determines the shape of the ellipse, allowing it to take the form of a line in the case of linear polarisation, a circle for circular polarisation and any other shape for all other polarisation states.

In a practical setting, we may wish to characterise some unknown light beam in terms of its polarisation ellipse. The first stage of that characterisation is to measure the total intensity of light we are receiving, containing no polarisation information. This forms the first Stokes' parameter, I and allows us to normalise all subsequent measurements to this total intensity value, which results

in all Stokes' parameters taking values between $-1 \leq S_{IQUV} \leq 1$. From Fig 4.1 c) we can see that any beam of light has a polarisation orientation which is the sum of the two Electric Field vectors where this orientation angle can take any value $0 \leq \chi \leq 2\pi$ and any ellipticity angle $-\frac{\pi}{2} \leq \Psi \leq \frac{\pi}{2}$. We can measure individual linear components of the polarisation states using a linear polariser which can be rotated through 2π with an arbitrarily defined but constant 0° value. This optical sheet only allows transmission of light rays which match the current orientation of the polariser. Measuring orthogonal linear polarisation states allows us to characterize both the orientation and ellipticity of the polarisation ellipse. The second and third Stokes' parameter are therefore defined as the difference between orthogonal states, with Q and U rotated by 45° relative to each other

$$Q = I_{0^\circ} - I_{90^\circ} \quad (4.19)$$

$$U = I_{45^\circ} - I_{135^\circ} \quad (4.20)$$

Where the I_{θ° are the various orientations of the linear polariser for a given measurement. Working from the equation of an ellipse in its standard form as shown in Eqn. 4.3, and the eccentricity of the ellipse as expressed in terms of the semi-major and minor axes, namely

$$e = \sqrt{1 - \frac{b^2}{a^2}} \quad (4.21)$$

it is possible to derive [52] the first two Stokes' parameters in terms of the ellipse orientation θ_e and its eccentricity, e .

$$Q = \frac{e^2 \cos 2\theta_e}{2 - e^2} \quad (4.22)$$

$$U = \frac{e^2 \sin 2\theta_e}{2 - e^2} \quad (4.23)$$

It follows that it is possible to rearrange Equation 4.23 to give e and θ_e in terms of Q and U

$$e^2 = \frac{2\sqrt{Q^2 + U^2}}{1 + \sqrt{Q^2 + U^2}} \quad (4.24)$$

$$\tan 2\theta_e = \frac{U}{Q} \quad (4.25)$$

Equation 4.25 are two powerful expressions which allow the near full characterisation of the polarisation ellipse from just four measurements with a varying linear polariser. To complete the ‘picture’ of the polarisation ellipse we must address the chirality of the polarisation state, which is only non-zero if the light is non-linearly polarised. Therefore, for known linearly polarised light sources, the polarisation state can be fully characterised by only the first three Stokes’ parameters, with the fourth assumed to be zero.

If we take our linear polariser, and introduce an optical device known as a ‘quarter wave plate’, which offsets the phase of transmitted light by $\frac{\pi}{4}$, and thus can be selected to only allow a non-linear polarised source rotating in either a clockwise or counter clockwise direction passage through the optical device. Thus, we can define the final Stokes’ parameter as follows

$$V = 2I_c - I \quad (4.26)$$

We can use derivation to acquire an expression of V in terms of e , but not θ_e as before

$$V^2 = \frac{4(1 - e^2)}{(2 - e^2)^2} \quad (4.27)$$

Thus, a value of $V = -1$ indicates counter clockwise polarisation and $V = +1$ indicates clockwise orientation. $V = 0$ indicates linear polarisation. As before we can express this in terms of the eccentricity

$$e^2 = \frac{2(-1 + v^2 + \sqrt{1 - V^2})}{V^2} \quad (4.28)$$

So we are able to obtain the eccentricity, and chirality of the polarisation state from a measurement with the QWP optical device, but not the orientation of the ellipse from this parameter alone.

Figure 4.3 b) demonstrates the Stokes' parameters as projected onto a 3D coordinate system, with the resulting surface for possible values of I known as a Poincare Sphere. Here the physical position or coordinates of the I vector determine the ellipticity and orientation of the polarisation ellipse. If the polarisation state never leaves the horizontal axis, for example, the polarisation state will be purely linear, with a changing orientation as the state 'moves' around the circle of the horizontal of the sphere. When one leaves the horizontal axis in either a positive or negative direction, the state becomes non-linear, with a larger ellipticity as it approaches one of the poles of the sphere, at this point the state is circular. This also illustrates the special case nature of linear and circular polarisation quite concisely, where linear polarisation is the special case where the polarisation state never leaves the 2D horizontal plane, and circular polarisation is the special case where the polarisation state is either

directly north, or south, with each representing LHCP or RHCP respectively. Here χ and psi are represented as angles mapped onto this spherical coordinate system. Additionally, we can define an expression for the total intensity, I_t as a fraction of the degree of polarisation present, p which can allow for a partial state of polarisation in our calculations, such that $I_t = Ip$. The Stokes' parameters can readily be expressed in terms of the spherical coordinates of the Poincare sphere as

$$Q = Ipcos2\Psi cos2\chi \quad (4.29)$$

$$U = Ipsin2\Psi cos2\chi \quad (4.30)$$

$$V = Ipsin2\chi \quad (4.31)$$

The factor of two for both the orientation angle and ellipticity angle are a result of symmetry effects of the ellipse wrt each angle.

4.2 Polarisation imaging with single-pixel cameras

Traditional visible light cameras can be readily made to be polarisation sensitive[37]. These imaging devices can probe both internal [53] and external [54] properties of various media and surfaces. Normally multiple cameras [55] are employed where each single CCD can produce an image for one polarisation state, so pixel correspondence is required between all cameras to a high degree of accuracy if an application calls for more than one polarisation state to be observed simultaneously. Such applications include Stokes' parameter imaging. In the case of linear Stokes' imaging, 4 cameras need to be spatially and temporally

synchronized with each other to provide the necessary data for recovery of the desired image, where each camera measures perpendicular polarisation states from a lab defined arbitrary angle of linearly polarisation illumination.

It was hypothesized that the fast, high fidelity system described in Chapter 3 of this work could be adapted to be fully polarisation sensitive [56]. This would offer some distinct advantages over a camera with spatial resolution for the first time, rather than simply being a demonstration of matching current capabilities of standard CCD imaging systems. When reconstructing single pixel images simultaneously with multiple detectors, the illumination source sets the angle and perspective of the image, in this case the DMD projector, since this is the only component in the device with has spatial resolution. Indeed, the only difference between spatially separated detectors in such a setup is the apparent illumination angle corresponding to the position of the detector, a rather unique property of computational single pixel imaging. This apparent illumination angle can be set from a single direction with the use of a BS so the reflection ‘observed’ by the detectors is locked to a single position. It is this property which lends the greatest advantage to investigating single-pixel imaging polarisation applications, as no pixel correspondence is needed regardless of how many simultaneous images are required for any given application. In addition, the potential ability for single pixel imaging to readily reconstruct from outside the visible portion of the spectrum allows the investigation of non-visible wavelength dependent polarisation properties such as birefringence.

A single-pixel camera system was successfully made polarisation sensitive and could reconstruct the linear Stokes’ parameters simultaneously. Our main goal for this system was to validate that it could indeed reconstruct images containing polarisation information for transmissive and reflective objects in

‘real-time’, and that simultaneous Stokes’ parameter image reconstruction can be achieved using differential signals, as opposed to differential images as is the current method for CCD based polarisation imaging systems. Furthermore the system can be adapted to operate at a ‘slow’ pattern projection rate of $\sim 1440\text{Hz}$, which is effective when temporal information is not required in the scene. This ‘slow rate’ creates a sufficient level of light intensity per projection to allow standard low-cost photodetectors to be implemented in the detection process, thus producing better results when given many samples to average over. Conversely when temporal information is desired, the system can be operated at a ‘fast’ pattern projection rate of up to $\sim 22\text{kHz}$. This allows near video-rate acquisition of the scene, however in this mode of operation photomultiplier tubes (PMT’s) have been employed to compensate for low light levels resulting from reduced sampling time. The choice of pattern projection rate in this system therefore offers flexibility depending on the desired application.

4.2.1 Stokes’ parameter image reconstruction

The first stage of the image reconstruction process is correlating measured signals from each detector with its associated illumination pattern. Using $D_{i\theta}$ as the measured detector signal we can define $D'_{i\theta} = D_{i\theta} - \widehat{D}_\theta$ where $D'_{i\theta}$ is the detector signal after subtraction by the current average detector signal, and $P_i(x, y) = P_i(x, y) - \widehat{P}_i(x, y)$, where $P_i(x, y)$ is the associated illumination pattern and $\widehat{\cdot}$ denotes the average value. Since the Hadamard testing regime natively provides 50/50 Black and White patterns, the average pattern is always 0 or 0.5 depending on how black pixels are defined as -1 or 0 respectively.

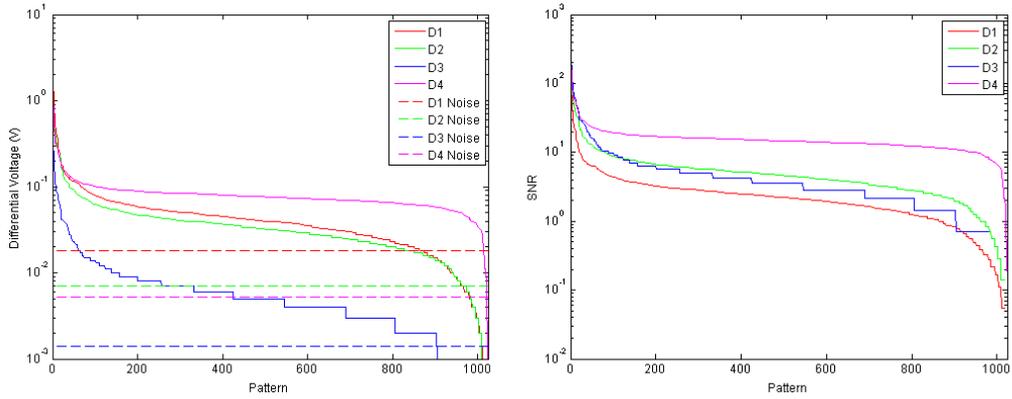


Figure 4.4: The left figure shows measured signals and associated standard deviation for each detector $D1$ - $D4$, plotted in descending order. The right figure shows the associated signal to noise ratio calculated for every pattern. This analysis was performed on a smaller data set containing 1024 patterns. The mean SNR for each detector in a ‘typical’ run is $D1 = 2.88$, $D2 = 5.89$, $D3 = 5.93$ and $D4 = 15.89$.

The iterative reconstruction is then obtained by

$$O_{\theta}(x, y) = \frac{1}{M} \sum_{i=1}^M D'_{i\theta} P_i(x, y), \quad (4.32)$$

for M iterations, and $O_{\theta}(x, y)$ is the resulting 2D image reconstruction for polarisation θ . After the complete basis set of M Hadamard patterns has been displayed, a fully-sampled image is reconstructed for each photodetector.

Figure 4.4 shows the analysis of a ‘typical’ single run of the system to quantify the strength of the signals versus the calculated noise on each of the detectors. The noise is calculated as the mean difference between two sets of signals whilst the scene remains static. In this experiment, it is shown that only a fraction of the total patterns provides significant contribution (above the noise floor) to the final reconstruction. Utilising only the significant fraction of patterns would enable a form of compressed sensing for the following frame, provided

the scene remains static, however this is not applicable in real-time to dynamic scenes since the significant pattern set would change.

The Stokes' parameter image corresponding to the total intensity, is obtained from summing the images corresponding to the orthogonal polarisation states, defined as

$$S_0(x, y) = O_{0^\circ} + O_{90^\circ} \equiv O_{+45^\circ} + O_{-45^\circ}. \quad (4.33)$$

One method of calculating images representing the linear Stokes parameters is by subtracting two images from each other, corresponding to orthogonal polarisation states, defined as

$$S_1(x, y) = O_{0^\circ} - O_{90^\circ} \quad (4.34)$$

$$S_2(x, y) = O_{+45^\circ} - O_{-45^\circ}. \quad (4.35)$$

Alternatively, the Stokes' images can be obtained by utilising the difference in the signals measured by detectors sensitive to orthogonal polarisation states, such that

$$S_1(x, y) = \frac{1}{M} \sum_{i=1}^M (D'_{i0^\circ} - D'_{i90^\circ}) P_i(x, y), \quad (4.36)$$

$$S_2(x, y) = \frac{1}{M} \sum_{i=1}^M (D'_{i-45^\circ} - D'_{i45^\circ}) P_i(x, y). \quad (4.37)$$

Once the first three linear Stokes' parameter images have been reconstructed, we can calculate the degree of linear polarization for each pixel if we assume the chirality stokes parameter, $S_3(x, y) = 0$, thus DoP is obtained by

$$DoP = \frac{\sqrt{S_1(x, y)^2 + S_2(x, y)^2 + S_3(x, y)^2}}{S_0(x, y)} \quad (4.38)$$

Furthermore, when Stokes' images are reconstructed for a specular multi-faced object, it can be appropriate to express the Angle of Polarization(AoP) which is produced from the following expression

$$AoP = \frac{1}{2} \arctan\left(\frac{S_1}{S_2}\right) \quad (4.39)$$

In this experiment reconstruction of Stokes' parameter images is performed utilising Eqs. 4.36 and 4.37, which reduces the computational load.

4.3 Experimental Method

The experimental configuration used to produce polarisation sensitive reconstructions is shown in Fig.4.5. The digital light projector consists of a DMD, and four coloured LEDs, (red, green, blue and near-IR). The DMD spatially modulates the light of the illumination LEDs to project binary patterns onto an object. A linear polarising filter positioned after the projection lens restricts the projected light to one polarisation state. The object shown in Fig. 4.5(a) is a stressed clear plastic sheet with a triangular portion removed, where the encoded light-field has transmitted through the plastic and is subsequently reflected back from a white diffuse screen placed behind the Perspex sheet, thus enabling the birefringence effects due to internal stress to be observed. The PD cannot be placed directly in the propagation path of the output of the DMD yet must observe light which has been transmitted through the Perspex sheet to observe internal birefringence effects. Therefore, the diffuse card placed directly behind the Perspex sheet enables the PD to be positioned at some off-axis angle in reflection to collect the maximum amount of light from the surface of the diffuse sheet. The four photo-detectors result in four sets of

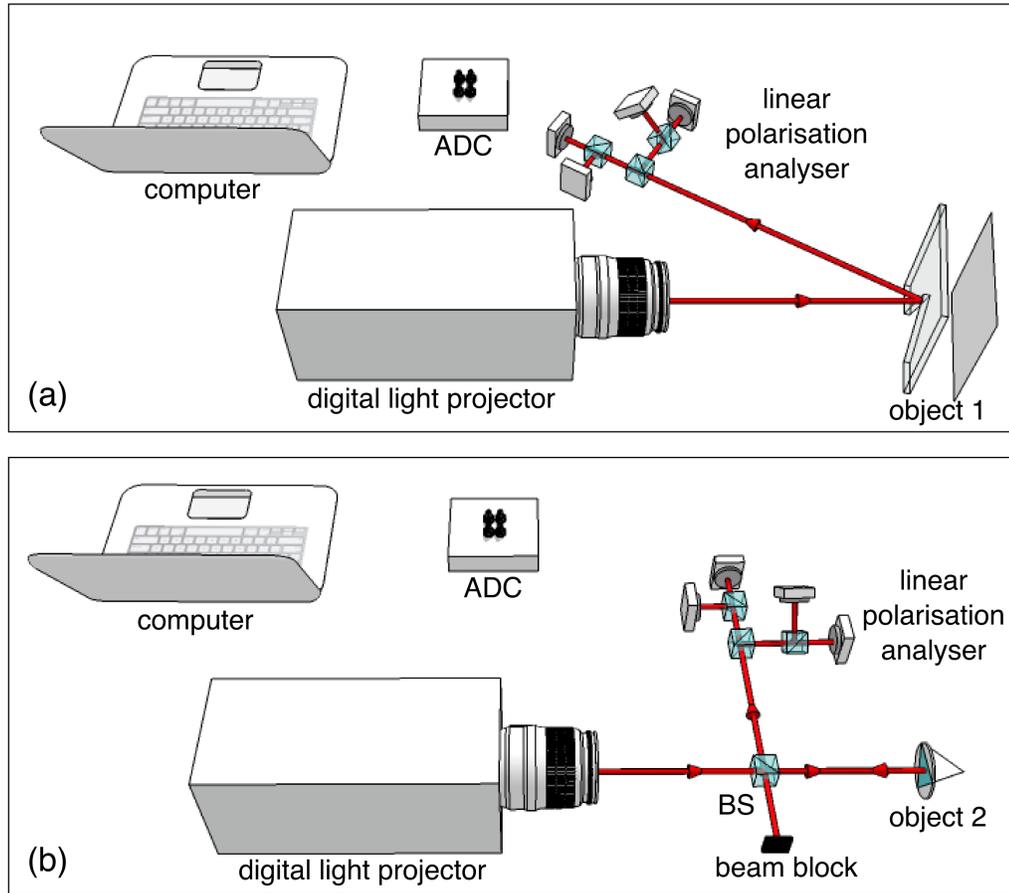


Figure 4.5: (a) An experimental system used to characterize linear polarization properties of diffusely reflected light from object 1, showing the digital light projector, the analogue to digital converter (ADC), linear polarization analyser and a computer running the control and reconstruction software. (b) A modified experimental system used to characterize linear polarization properties of specularly reflected light from object 2, in this instance a corner cube. Light reflected from the object is propagated through the additional BS into the linear polarization analyser.

signals for each differential pattern and is akin to a four CCD Stokes imaging system [55]. The patterns applied to the DMD and corresponding measured

signals are used for image reconstruction utilising an efficient computer algorithm (See Chapter 2 Eqn 2.6).

Figure 4.5 (a) was the initial configuration of experimental setup to primarily test the Stokes' parameters for a given object. Object 1 in this case was a stressed transparent piece of Perspex with a triangular portion removed at the bottom. This was done to create a natural focus for stress in the object. Thus, if our system was truly polarization sensitive, it would be easy to observe the rapid polarisation state changes one would expect within the stressed object. This is known as birefringence, where a transparent medium has a varying refractive index dependent on the direction and polarisation state of the transmitted light beam. These refractive index changes can occur in some material which have varying degrees of internal pressure, or stress placed on them. Birefringence is therefore wavelength dependent since the refractive index at any given point in the internal structure of the object will cause a direction change, in which the magnitude of that change is dependent on the specific wavelength of the light.

Since the exact propagation path of light in Fig. 4.5 a) involves non-intuitive free space propagation, it can be summarised as follows;

- The projector fitted with a linear polariser provides the initial fully linearly polarised illumination with an encoded pattern.
- There is an 8% reflection at the front surface of the Perspex sheet, this is effectively wasted intensity.
- Light then propagates through the Perspex and falls incident on the diffuse reflector sheet placed behind.
- An additional 8% reflection occurs on the back surface of the object from

the diffuse reflection, further attenuating total intensity

- The reflected light then propagates once more through the Perspex following paths determined by refractive index variations within the structure of the object.
- Some fraction of diffuse light which has propagated twice through the object enters the first BS within the polarisation analyser.
- This analyser ‘directs’ and splits different linear polarisation states to the correct PD, which measures the total intensity transmitted through the analyser for that specific polarisation state.

In a further effort to obtain the highest quality results, attempts were made to simplify all other aspects of the experiment. This included limiting the projection and detection wavelength to be monochromatic per ‘run’ of the system, where a single run constituted the projection of a complete set of Hadamard patterns. This is a form of sequential multi-wavelength illumination and from a technical point of view, is the simplest to implement in practice, as the number of detectors required to successfully perform image reconstruction is minimised and no further synchronisation is needed for the inclusion of multiple wavelength’s, only repeating a new run with a new wavelength. Furthermore, the linear polariser was left as a manually rotatable device, so each arbitrary polarisation angle was set by hand rather than automating this process. While automation has obvious advantages, it does add an additional program interface between the computer and physical experiment, and for this investigation, only a few arbitrary angles were required to demonstrate the polarisation sensitivity of the system. In any future investigation, where many angles might be required, automation of the rotation of the linear polariser is the first improvement to this system which could be relatively easily undertaken. The system could then find use in a wide variety of polarisation imaging

applications without the need for a human operator beyond the operation of a connected computer, such as industrial testing of transparent media where multiple orientation angles are required in quick succession.

Figure 4.5 (b) details the second configuration of the system used to probe additional fundamental properties of linear polarisation, including Angle of Polarisation(AoP). An additional BS is placed between the polarisation analyser and the object. This allows the specular reflected light from the object to be measured by selecting the appropriate output of the BS. Object 2 used in this second demonstration is a ‘corner cube’, an internally mirrored cube where surfaces were at 45° angles to the input illumination. These well-defined angles and highly reflective non-transparent surfaces allowed the AoP to be readily represented in a reconstructed image after the use of an appropriate algorithm on the Stokes’ parameter images, produced from each ‘twinned’ detector within the polarisation analyser.

4.4 Results

The Stokes’ parameter images, S_0 , S_1 and S_2 , are acquired simultaneously for a setting of the incident polarisation. The object, shown in Fig. 4.6 is a stressed clear plastic sheet having a triangular section removed. The centre of each image indicates a point of stress where we observe colour dependent intensity variations in S_1 and S_2 that are not seen in S_0 , indicating birefringent effects and hence validating the polarisation sensitivity of the system. The results in Fig. 4.6 have been obtained whilst operating the system in the ‘slow’ mode, with a pattern projection rate of $\sim 1440\text{Hz}$. The lower display rate, and thus increased sampling time, allows the use of low-cost, wide-band photodiodes. The choice of resolution (96×96 pixels) was chosen to provide enough detail

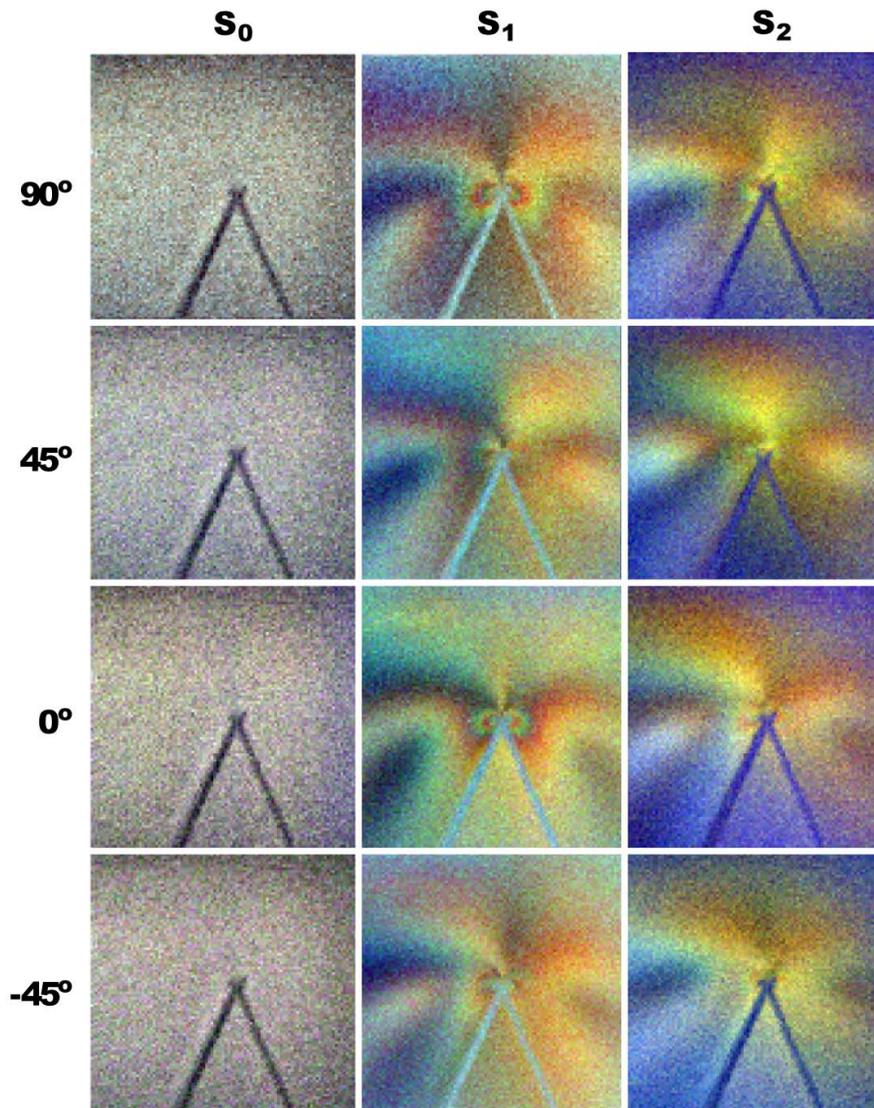


Figure 4.6: Stokes' parameter images S_0 , S_1 and S_2 reconstructed in colour (RGB), for which the object is a stressed clear plastic sheet which has had a triangular portion removed. The birefringence effects which are wavelength dependent occur within the plastic. The angle (in degrees) of the projected linear polarisation state is indicated on the left-hand side and relative to an arbitrarily defined 0° state.

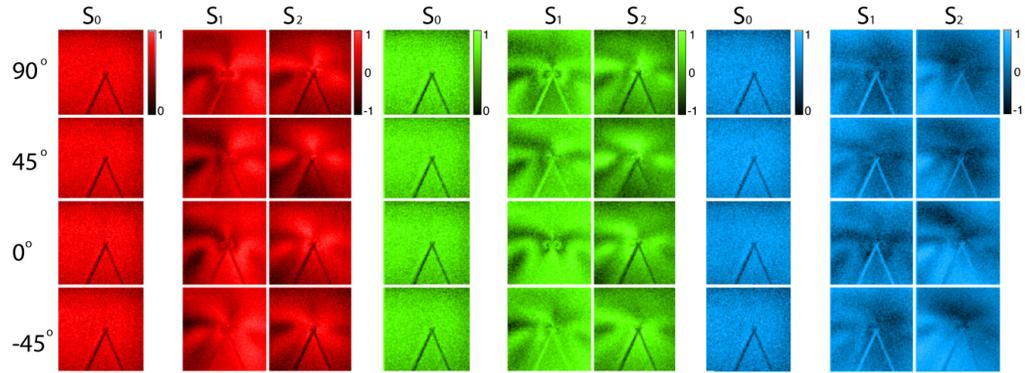


Figure 4.7: *Iterative reconstructions of the first three Stokes parameter images S_0 , S_1 and S_2 used in the RGB combination images. The object is a clear plastic sheet under stress. The three coloured LED sources have an approximate bandwidth of 40 nm, which leads to broadening of the birefringence features. Each image resolution is 96×96 pixels.*

in the images to reveal birefringence effect whilst also minimizing the fully-sampled acquisition time of approximately 15 seconds per image.

Figure 4.7 is the deconstructed colour planes which form the RGB full colour images shown in Fig. 4.6. The specific width and other features present in objects which show birefringent properties are highly dependent on wavelength, thus for proper analysis and interpretation of the results shown in Fig. 4.6, the individual colour planes are given which were recovered from the sequential illumination process.

Table 4.2 displays the corresponding SNR estimates for the 0° linear polarization illumination reconstructions in Fig. 4.6 and Fig.4.7 for the RGB and colour channels respectively. Since the SNR was nearly invariant with linear polarisation angle, only the 0° results were included in the table.

Table 4.2: SNR Estimates for Fig. 4.6 and Fig.4.7 at 0°

	S_0	S_1	S_2
RGB	3.9	4.3	3.9
RED	4.1	4.8	4.1
GREEN	4	4.5	3.9
BLUE	4.1	4.3	4.5

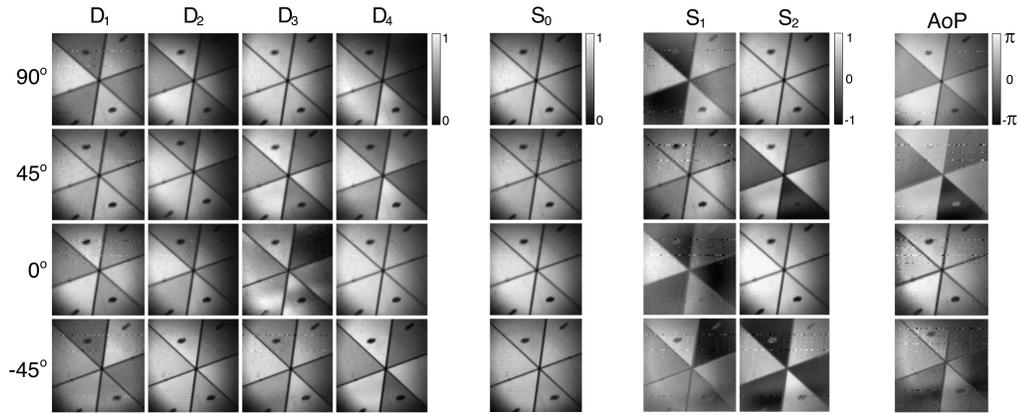


Figure 4.8: *Reconstructed images of a corner cube obtained by the polarization sensitive single pixel camera prototype for various incident linear polarization illumination. D_1 to D_4 columns represent the reconstruction from each individual detector. S_0 , S_1 and S_2 columns represent the first three Stokes parameter images obtained through combinations of the signals as described in the main text. The angle of polarization column, AoP , shows the effect of a multi-faced surface where each mirror segment was at a different angle than its neighbouring segment, with respect to the linear polarization illumination angle.*

4.4.1 Angle of Polarization

Figure 4.8 shows a complete set of image reconstructions obtained from the system when observing a ‘corner cube’ optical component, utilising a modified

experimental setup as shown in Fig. 4.5(b). A corner cube was chosen because it is essentially a set of mirrors which have different angles with respect to the incident linear illumination angle. The results allow a qualitative characterisation of the systems response to a multi-faced object, and may allow direct comparison to previously reported Stokes parameter imaging utilising detector arrays. We observe differences in relative intensity for the different mirror faces as expected. Single-pixel cameras operating in a projection or masking mode, tend to suffer when a portion of the scene exhibits high specular reflectivity, since this can result in overwhelming signals measured for certain patterns and not others, degrading the final reconstruction. In this instance, the entire object exhibits specular reflection and indeed this resulted in high-signals measured independent of pattern choice, and led to high-quality images to be reconstructed. This is a further validation that the system is truly polarisation sensitive as only with polarisation information can the exact angle of a surface relative to the input illumination be determined from a 2D image where only a single angle of observation has been recorded. This is only true where the surface is highly polarisation preserving, thus a highly reflective smooth mirror surface fits this assumption well.

4.4.2 Degree of Linear Polarization

Orthogonal polarisation imaging is a commonly used technique for investigating the non-reflective components in diffuse objects. Therefore, in a subsequent experiment the inert object was replaced with a living object, in this case a small cactus plant, which was chosen for the distinct differences between the sharp ‘needles’ and soft ‘flesh’. The results are shown in Fig. 4.9. As with Fig. 4.6 the first column, representing S_0 (reflected intensity) does not change as a consequence of different incident polarisation. Whereas analysis of S_1 and

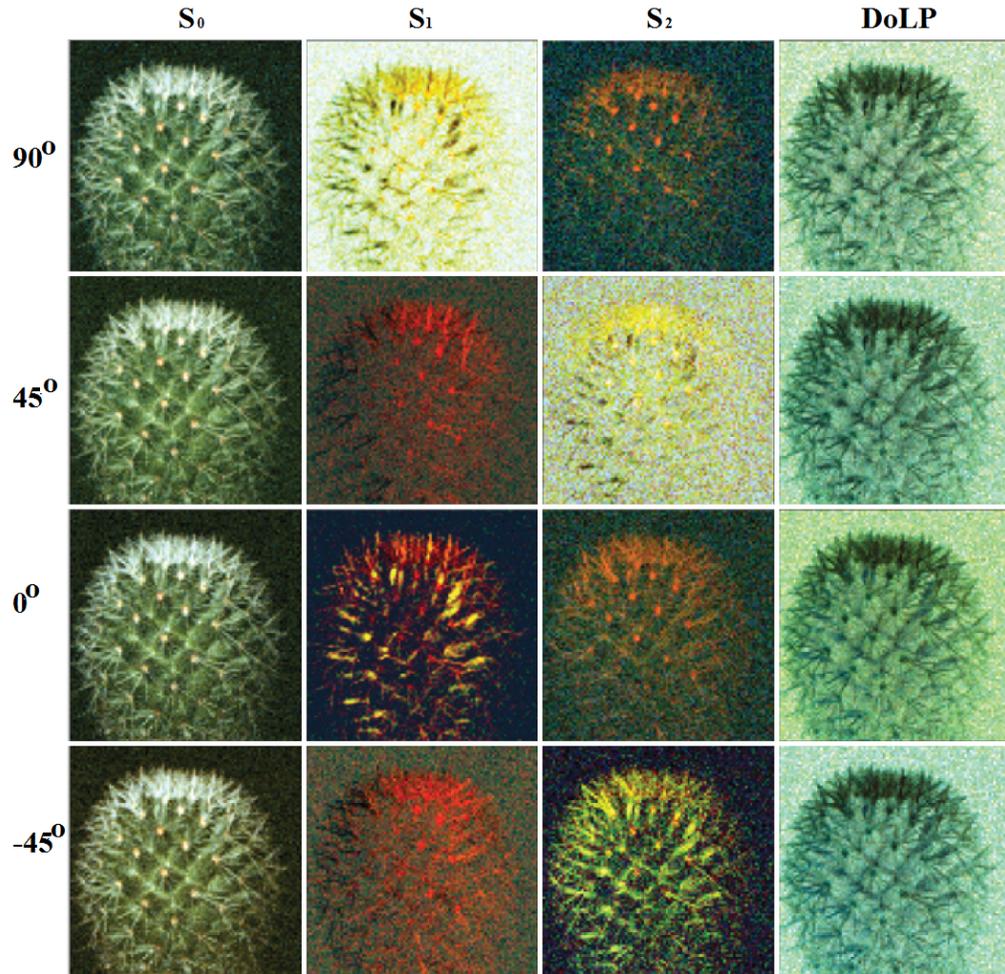


Figure 4.9: *Stokes parameter image reconstructions for a small cactus plant. The sharp ‘needles’ on the cactus are shown to exhibit greater polarisation maintaining reflectivity compared to the soft ‘flesh’ of the plant as indicated by stark contrasting with DoLP being invariant with angle of linear illumination as expected.*

S_2 reveal differences as a result of the incident illumination, in particular the ‘needles’ appear to exhibit enhanced contrast compared to the ‘flesh’ for different incident polarisation. Note that 90° and 0° images of S_1 are the inverse of each other, similarly with 45° and -45° images of S_2 . When the incident

Table 4.3: SNR Estimates for Fig. 4.9

	S_0	S_1	S_2	$DOLP$
90°	11.5	4.4	3.3	4.7
45°	10	3.1	2.3	4.8
0°	10.2	8.8	4.8	4.9
-45°	10.6	3.4	8.2	4.8

illumination is at an angle of 45° with respect to the angle of the PBS, the corresponding photodetectors will measure a similar light intensity leading to an image exhibiting mostly uniform noise, due to the normalisation process used here. We also note that these noisy images appear red, which is believed to be the consequence of the inherent differences in spectral response of the photodetectors. The final column represents the recovered DoLP information, the invariance with illumination polarisation angle verifies the system can recover DoLP accurately.

Table 4.3 contains the SNR estimates for the corresponding image ‘table’ shown in Fig. 4.9. We see an approximately uniform SNR estimate for intensity and DOLP as expected with higher SNR images for those Stokes’ parameter images which matched the ‘straight’ through position within the polarisation analyser. It should be noted then that the S_1 , 90° result has lower than expected SNR. However, this may be due to the SNR estimation process as white background images are moderately more challenging to estimate than black background images.

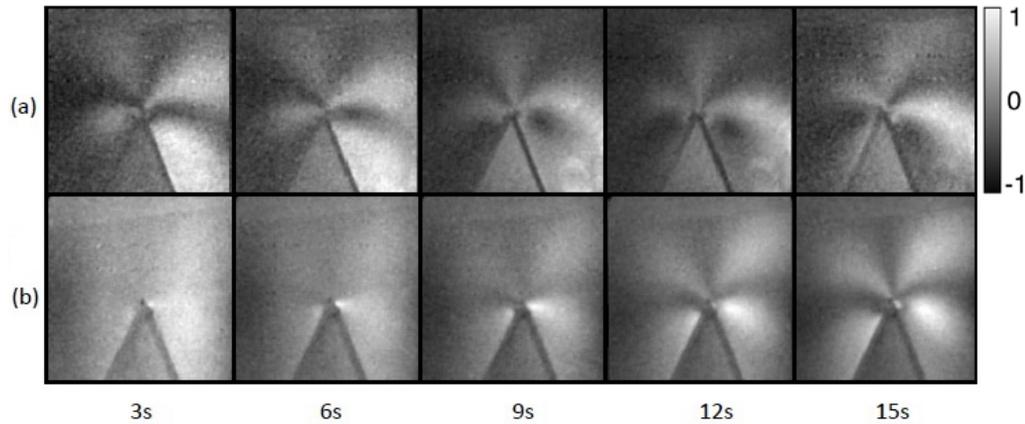


Figure 4.10: A selection of S_1 images acquired at different intervals for scenes exhibiting dynamic behaviour. The object is a clear plastic sheet with a triangular portion removed, and illuminated with red light. The S_1 images have a resolution of 128×128 pixels. (a) While an object was maintained at constant stress the linear polarising filter in the illumination path was rotated. (b) The polarising filter is fixed while the stress applied to the plastic sheet is gradually increased.

4.5 Video rate single-pixel polarization imaging

To acquire Stokes parameter images from dynamic scenes requires utilising faster pattern projection rates or utilising fewer patterns from the set via compressed sensing techniques. The latter presents a non-trivial computational reconstruction algorithm, this increases computational load which can make it difficult to achieve near video framerates and ‘real-time’ reconstruction, thus it was beyond the scope of this work, although this approach is currently under investigation as a possible method of improving the technique. Instead, the pattern projection rate was increased to $\sim 22\text{kHz}$ and the resolution of the pat-

terns limited to 128×128 pixels. Along with differential pattern projection this leads to a maximum video frame rate of $\sim 0.7\text{Hz}$. Importantly however, when operating the system in ‘fast’ mode results in less light received by detectors and hence lower signal-to-noise, leading to poor image quality. To compensate for reduced light levels the photodetectors were replaced with PMT’s.

Figure 4.10 shows the results acquired from this system in one colour channel while the scene exhibited dynamic behaviour. For clarity, only S_1 is displayed here. The object described in Fig.4.6 was used for this investigation. In Fig.4.6a the plastic was maintained under constant stress while the linear polarising filter in the illumination path was rotated by hand. We observe that the intensity variations in each image rotate in accordance with the linear polariser, verifying that the Stokes’ imaging system is capable of producing images at near video-rates. In Fig.4.10b the linear polarising filter was fixed while the stress applied to the plastic was gradually increased. We note that the contrast between the intensity variations in each image increases as the applied stress to the plastic also increases, further validating the near video-rate capabilities of the system.

4.6 Summary

Conventional polarisation imaging is notoriously challenging when simultaneous multiple polarisation states are to be imaged, as is essential in Stokes’ parameter reconstruction. This is due to the near pixel perfect correspondence that is required between individual cameras to provide accurate reconstruction. Taking the single pixel approach to Stokes parameter imaging provides an innate benefit in that in effect, there is always perfect pixel correspondence when multiple single element detectors are positioned at the output of an appropriate

polarising beam splitter.

Technological advances in the DMD allowing extremely fast mirror flip rates of over 20kHz, was the fundamental advantage which allowed near video rates to be achieved. This allowed direct imaging of temporally dynamic objects, such as the Perspex sheet under short scale varying degrees of stress or a rotating linear polariser.

Chapter 5

Current Investigations with single-pixel detectors and future work

5.1 Long Range and Astronomical Observa- tion

Astronomical observation and imaging is a wide-ranging field which encompasses both ground and space based imaging platforms. Observations across the entire electromagnetic spectrum can yield valuable and fascinating insights into astronomical phenomena [57, 58], where each region of the EM spectrum of interest to the astronomy community requires an appropriate detection system and where applicable, have high spatial resolution in order to obtain the highest quality data under demanding conditions.

Most ground based observations have been constrained to visible and radio

bandwidths due to atmospheric absorption blocking a large portion of the EM spectrum, however, there exists ‘windows’ of transmissivity where the atmosphere becomes opaque to relatively narrowband regions, primarily due to gaps in the absorption of water vapour and other greenhouse gases in the atmosphere. Some of the main observational windows lie within the near IR/short wave infrared (SWIR) region, thus having a high-resolution camera sensitive to these regions allows observations beyond visible bandwidths when atmospheric absorption is a problem.

The single pixel imaging techniques demonstrated thus far have been shown to operate successfully in SWIR regions using an appropriate spectral filter, can operate in predetermined bandwidths. This suggests that the single pixel imaging approach could be used as method to exploit these atmospheric windows as an alternate to currently available and potentially expensive SWIR cameras, or where there may be regions of interest where no current commercial camera has peak sensitivity. It should be noted that these regions of high transmissivity do experience fluctuations in absorption levels primarily due to location and local atmospheric composition. Additionally, due to refraction effects in the atmosphere, it is well known and understood that light at the longer end of the spectrum penetrates the atmosphere much better than shorter wavelength light, thus when observing over large distances where atmospheric turbulence is significant, simply observing in a longer wavelength can improve image quality. The highly tuneable nature of a single pixel camera imaging system with a range of spectral filtration and high operational bandwidth when utilising spatial modulation technologies such as the DMD, further gives motivation for implementing a system capable of astronomical observation based on single pixel imaging technologies.

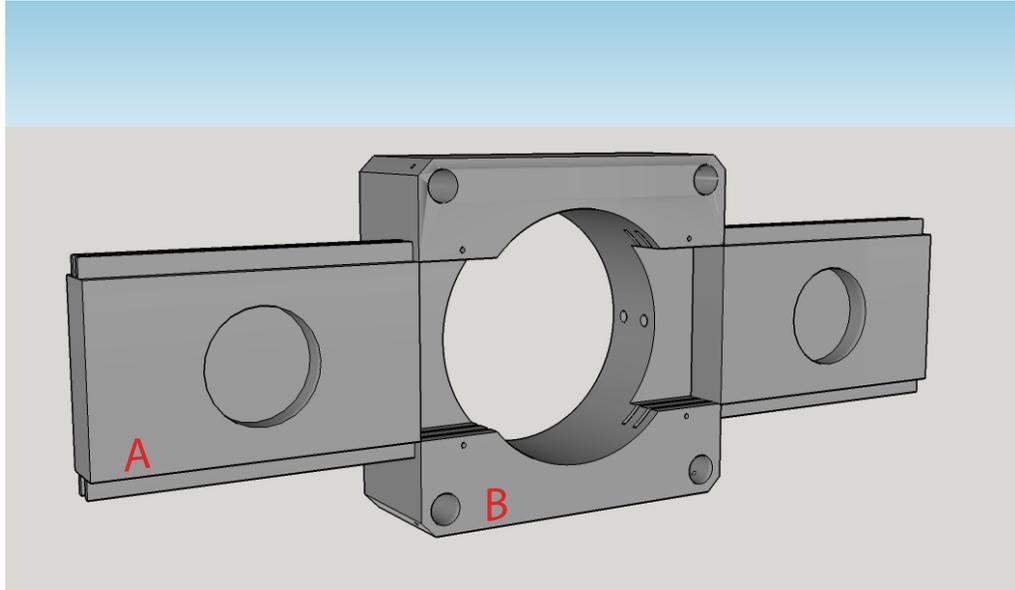


Figure 5.1: *Original design of the attachment component for mounting a single pixel imaging system onto a telescope. A) This section of the design is a removable slider which was made to house a standard 2cm concave mirror with focal length approx 50mm. B) The main body of the mount with a 2inch aperture to fit standard commercial telescopes. Grooves enabled mirror sliders to be inserted from both sides to an adjustable position to capture the dual reflection outputs of the DMD.*

5.2 Design and Experimental Setup

Figure 5.1 is the initial concept design for a portable telescope adaptor for a single pixel imaging system. This design features adjustable mirror ‘sliders’ which house standard 2cm concave mirrors for collecting light from the outputs of the DMD. The DMD is moved into the focal plane of the telescope and the mirrors are adjusted to the correct separation to collect the maximum amount of light from the DMD outputs.

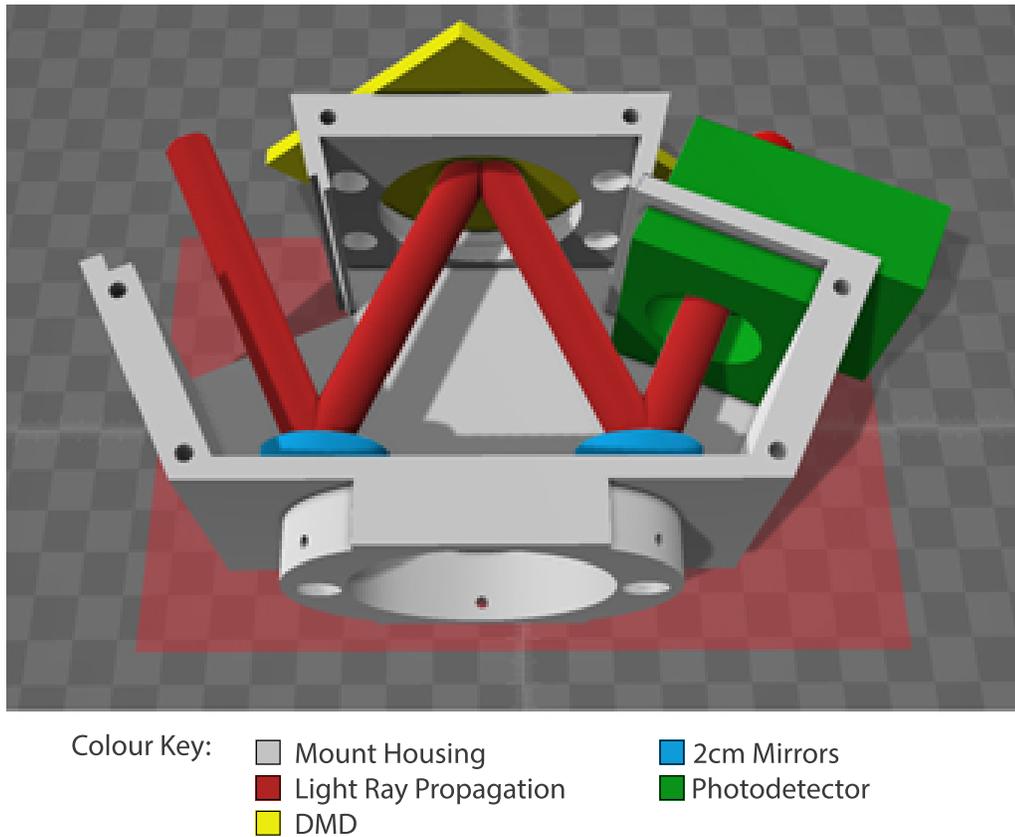


Figure 5.2: *Updated and final design model for the telescope single pixel camera mount. Rays have been added which show the propagation of light through the system. the mirror positions are now fixed and the entire unit can be sealed by adding a top casing to the unit for added background light reduction. A second detector is normally present at the empty output. This allows several applications such as differential or multi-wavelength imaging. With credit to Dr Graham Gibson for the final 3D design.*

Figure 5.2 displays the 3D model of the finalized single pixel portable housing system. This self-contained unit has fixed mirrors which are at the correct distance, separation and angle for a given DMD such that the total light from

the dual outputs of the DMD chip will be reflected onto photodiodes, which are attached via modular components of the casing. This entire module was mounted onto a 2 inch metal tube selected for a length which corresponded to the output of the telescope being approximately the same size as the DMD chip, to accommodate maximum light collection. In this specific instance, this ended up being 4 inches of tubing. This tube would then easily fit on a standard 2 inch Baadar focuser at the primary output of the telescope. Excluding the additional metal tubing, this is the standard way most imaging devices are attached to commercial telescopes and was done in this manner to demonstrate the versatility of the SPC system working in similar manner to conventional cameras. Figure 5.3 shows the 3D printed final prototype design for the single pixel telescope adaptor. This final design can also work well as a general housing for a DMD based single pixel imaging system, with pre-set mirrors already positioned to capture either incoming or outgoing light from the system. The unit is modular, with removable sliders for placing either light sources or detectors as required. For mounting this portable unit onto the telescope, two detectors are used and the system always functions in an incoming light mode, as it is hardly practical to attempt to project patterns onto objects at range. The detectors used in this specific example were Thorlabs DET36A for visible, which is a biased Si PD and Thorlabs DET20C for IR, which is an InGaAs standard PD.

5.2.1 Preliminary Results

The 3D printed prototype housing for the single pixel imaging system shown in Fig. 5.3 was mounted on a commercial 8 inch Maksutov-Newtonian telescope. After finding a suitable observation site, buildings between 1-2kms away were observed and the images recorded in both visible and SWIR.

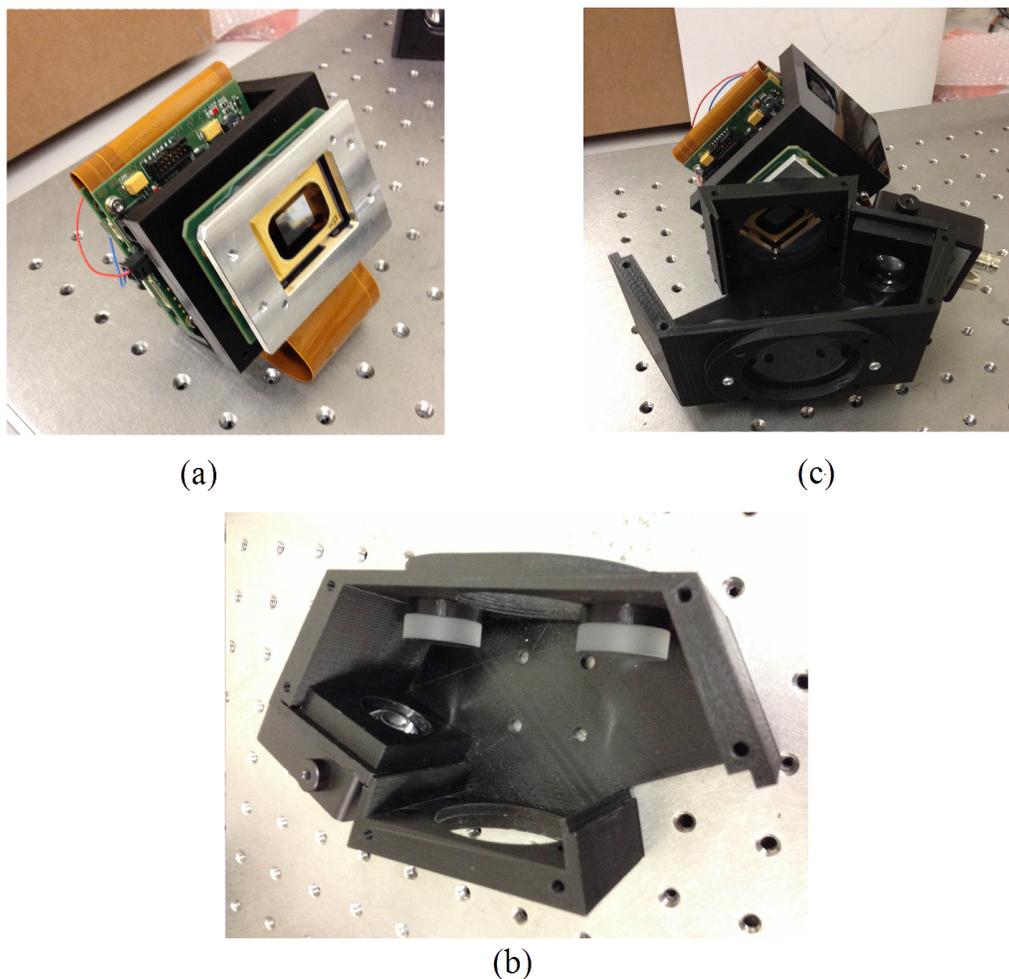


Figure 5.3: *a) Standard $0.2\mu - 2\mu$ coated DMD b) 3D Printed mount housing produced from the 3D model shown in Fig. 5.2 c) DMD attached to the 3D printed model ready to be mounted on a telescope or used in any number of lab based experiments. This standardised mount proved to be applicable in a number of recent single pixel imaging applications beyond the telescope imaging system.*

Figure 5.4 demonstrates the ability of the system to image in both visible and SWIR at much greater distances than had been previously shown. Single-pixel cameras had only been shown to produce successful reconstructions in

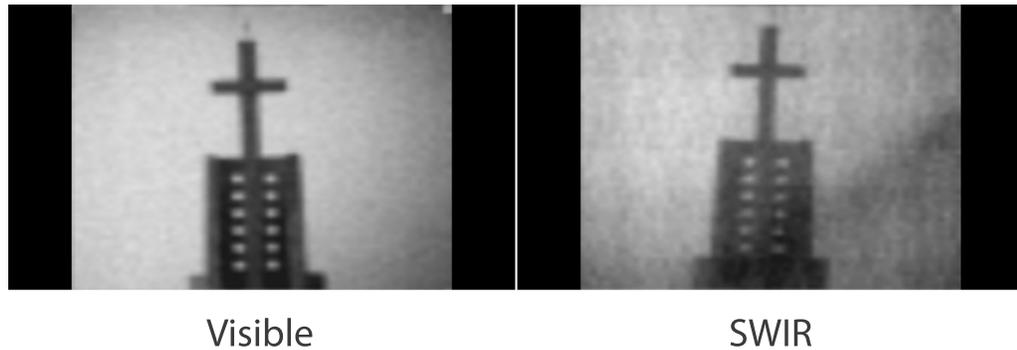


Figure 5.4: *Initial reconstruction result from a single pixel camera mounted as shown in Fig 5.3 which is being used with a 8 inch Maksutov-Newtonian commercial telescope. The target was the steeple of a church 2km from the single pixel camera observation site. The reconstruction was produced at a near video rate of approximately 10Hz where visible and SWIR signals were captured simultaneously.*

the range of a few metres at maximum. The motivation for obtaining this result was initially to see if atmospheric turbulence would be reduced in the SWIR reconstruction vs the Visible result. However, overall lower signals were observed in SWIR leading to a slight decrease in image quality as observed. The method for obtaining this result is in effect identical to that demonstrated in Chapter 4 for polarisation imaging, except we have a SWIR detector. This would specifically be a ‘fast’ acquisition regime, with the DMD operating at over 20kHz to produce near-video rate reconstructions. Hadamard patterns were again used, fully sampling at Nyquist using the basic iterative reconstruction algorithm, and not undergoing any compressed sensing.

Figure 5.5 shows 6 frames from a real-time video rate single-pixel camera reconstruction of a dynamic object at extreme horizontal distances of approximately 20km. These objects were wind turbines belonging to Whitelee windfarm on

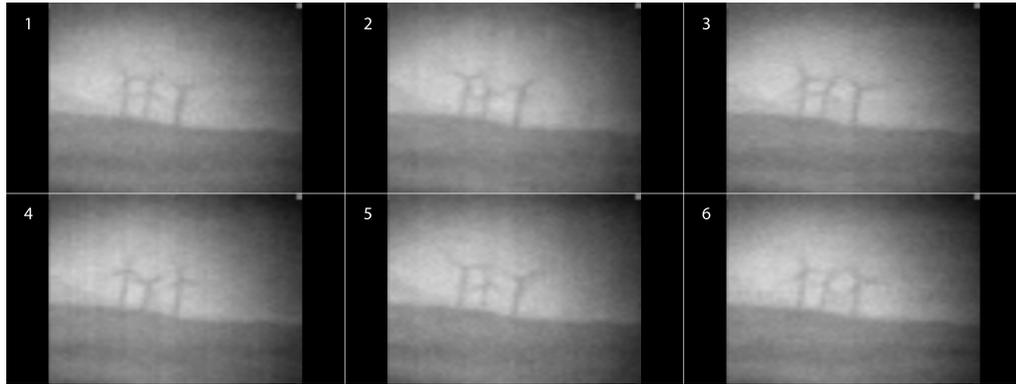


Figure 5.5: *A series of SWIR images were captured with the single pixel camera system at near video rate reconstruction rates of approx. 10Hz. The targets imaged here are of wind turbines belonging to a wind farm situated 20km from the observation location. The resolution of the reconstruction and fidelity of the images is of sufficient quality to observe the wind turbine's rotation.*

the outskirts of Glasgow City.

5.3 Additional Future Applications

5.3.1 METAMATERIALS for Optics

There have been many advances in recent years in the development of new materials which demonstrate properties or effects not found in nature. These materials generally are made from various substances where micro or nanoscale 'features' from together to have a noticeable effect on larger scales which is generally not present on the 'feature' scale of the material. A transparent metamaterial has been developed [59, 60] which has optical applications, demonstrating analogue view rotation [61] of 2D and 3D objects. This new material, referred

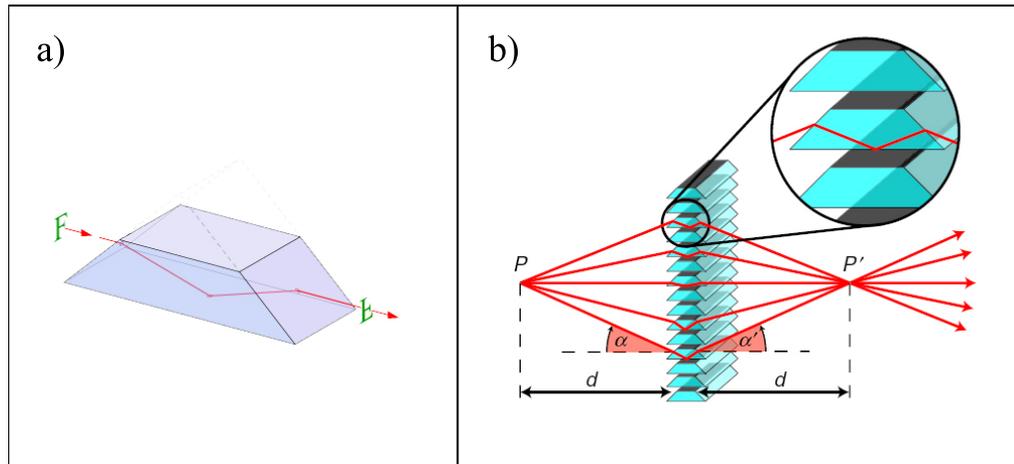


Figure 5.6: a) A single dove prism flipping the image of ‘F’ upside down as one would expect an image to appear when viewed through such a prism. b) Many dove prisms create an ‘array’ which has the net effect of flipping each component of object, P which is imaged upside down at the image plane, P' .
Figure Credit : Johannes Courtial

to as a ‘METATOY’ (metamaterials for optics) could potentially have an application within a single pixel camera system. These METATOY ‘windows’ are formed from an arrangement of two dove prism arrays which are rotated relative to one another. The ray-rotation[62] which takes place within these dove prism arrays is a form of generalised refraction [63].

Figure 5.6 Demonstrates the ray flipping process of a dove prism array. In section b) we see how many small dove prisms can individually flip the vertical component of individual light rays coming from discrete projection angles from the object P . A METATOY window is formed when two such dove prism arrays are aligned at some orientation angle θ wrt the ray-flipping axis of each window to reach other. The addition of this second dove prism array again flips the incoming light ray a second time, returning it to its original sign,

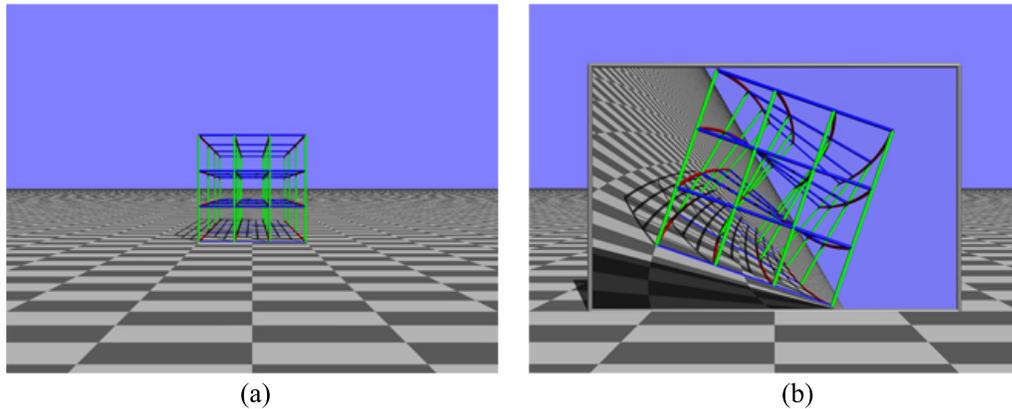


Figure 5.7: a) A 3D object placed at some distance $x = P + P'$ from the observer, in a simulated 3D environment b) A METATOY window placed at a distance P' from the observer. The extended object seems to have undergone a visual ‘twist’. including the visible background.

while having been rotated through the dove prism arrays orientation angle. The well documented effect[64] of these windows is that 2D objects placed behind a window will appear rotated. The angle of this rotation is dependent on the linear distance the object is placed behind the window. Thus, these METATOY windows can be used as a simple ranging device for 2D objects. For 3D objects the behaviour of the image observed through a METATOY window is more complex. Since each plane of the object is at a different distance from the window, the object will appear ‘twisted’ as each plane has underwent a different rotation. This optical metamaterial is effectively acting as an analogue image distorter and/or rotator. Spatial Modulation of a light-field has potential overlap for application in single-pixel camera systems.

Figure 5.7 a) and b) were recorded on the ray-tracing and interactive environment software TIM [65]. a) Shows a wireframe 3D cube object placed some distance x away from the observer. In section b) we see that a simulated

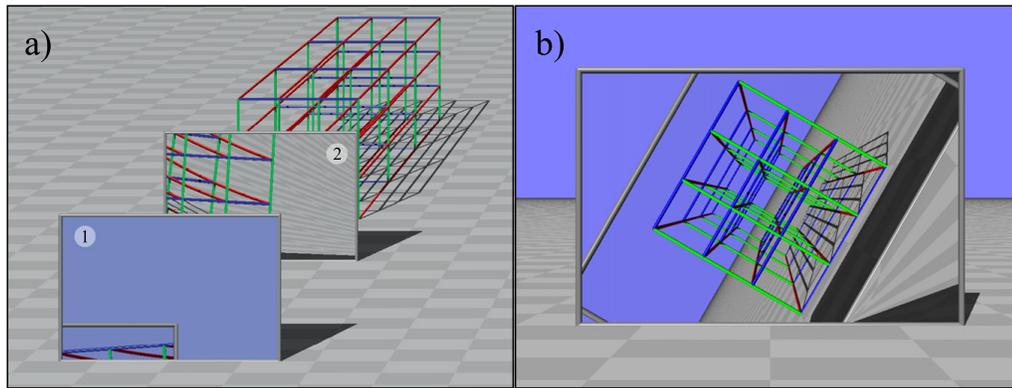


Figure 5.8: a) An isometric view of the two METATOY window setup for view rotation b) The view rotated 3D object is seen here with some single rotation angle for the entire image, but no warping of the object as seen in Fig. 5.7 b)

METATOY window has been placed between the observer and the object, so that the observer is in one imaging plane P' with the object front P distance behind the window. From the observation made through the window, we can see that the extended object has been ‘twisted’, along with the rendered background. The rotation angle that each plane has been rotated through is dependent solely on the distance that plane sits behind the window, thus each 2D cross-section of the wireframe cube appears rotated by a different angle, leading the warping or twisted effect observed.

Figure 5.8 a) Demonstrates the view rotation setup in a simulated METATOY window experiment, where the distance between the two windows, and the distance between the first window and the observer must be set precisely for the view rotation effect to occur. b) Demonstrates how an object would appear when the correct conditions for view rotation have been met, fully compensating for the twisting effect noted in extended 3D objects when a single METATOY window is present.

5.3.2 Image enhancement via spectral image masking

Demonstrations [18] have been made of the apparent correlation between the transmissivity of the object and the image quality in single-pixel imaging. The main conclusion of this finding was that the lower the transmissivity at constant signal strength, the higher the quality of the final image. This is because as the fraction of the object which is visible to the detector is reduced, there is a far greater fractional contribution from any pattern which scores highly, i.e any pattern which had many white squares in the location of the object. Whereas when you had a much more transmissive object, the fractional contribution of each pattern was lower overall. In effect, single pixel imaging is very good at finding small bright objects, but suffers to a degree when faced with real extended scenes with subtle variations in intensity over the entire field of view. Overall however, the fractional difference increase in the weighted patterns obtained by forcing the total object area to be $\frac{1}{3}$ or less would always lead to an improvement in observed image quality, even at a reduced overall total light intensity for the $\frac{1}{3}$ or less illumination. Therefore, the following is a proposed method in which single-pixel imaging device can use multiple detectors to image a single object, achieving the image quality of a low transmissive object even when the object is highly transmissive.

Figure 5.9 is the theoretical experimental setup of a system capable of improving image quality for highly transmissive greyscale or binary objects. The physical system setup is identical to that in Chapter 3 of this work focusing on simultaneous multi-wavelength image reconstruction. Here the greyscale object has been masked by 3 strips of coloured filters, RGB in this case. Since the spectral information of the object can be disregarded, each detector only reconstructs $\frac{1}{3}$ of the original object in total, respective to the matching spectral bandwidth. These three sectional components of the reconstruction are

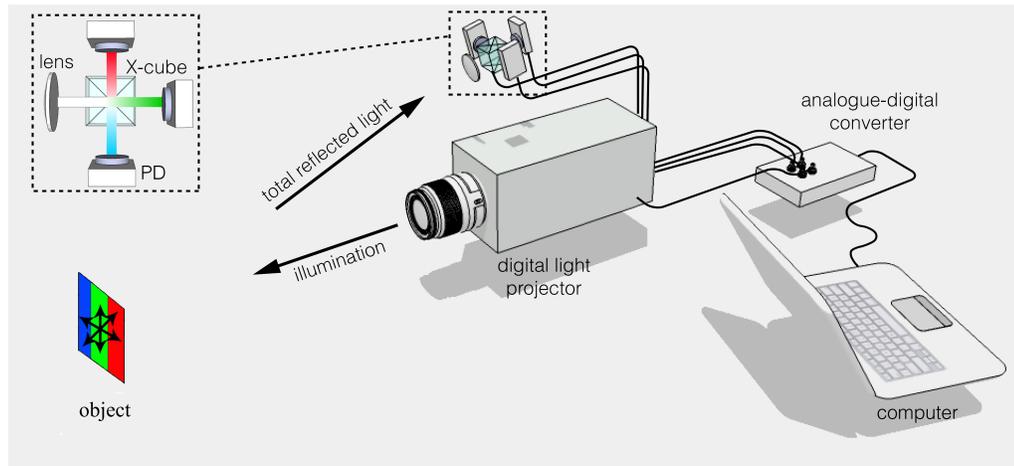


Figure 5.9: *The simultaneous multi-wavelength single pixel imaging system modified for image enhancement by spectral image masking. Here a highly transmissive yet greyscale object is placed in the path of the illumination. The projector illuminates the object with white light where the object has been masked by coloured filters into 3 distinct sections, RGB in this case.*

overlaid to form one final reconstructed image. In theory, the final image should have the relative ‘quality’ of a transmissive object $\frac{1}{3}$ of the objects actual transmissivity. This is case of using spectral information for a different purpose than spectral imaging. The object could be further subdivided into various narrowband spectral ‘segments’ further increasing the quality of the final image, this simply takes the same number of extra corresponding detectors with corresponding spectral filtration.

This proposed experiment could also prove useful in the case where single pixel detector arrays are used to image non-visible portions of the spectrum, where the application does not call for narrowband observations of the object in the non-visible region. Such an array of non-visible detectors accompanied by a DMD coated in an appropriate coating could produce single pixel images with

greatly enhanced quality than standard single pixel imaging, where in non-visible regions it is often the case that image qualities become compromised due to lower signals or higher noise profiles in the more expensive detectors commonly used in non-visible regions such as IR or UV.

5.4 The Future of Single Pixel imaging

The single largest obstacle to performing high resolution single pixel imaging on acceptable reconstruction timescales, is computational load. While physical image projection devices exist which have resolutions of 1MP and display rates of 24kHz, the storage capacity of these devices has been severely limited. In the work presented here, there is no image greater than 128x128, even for temporally static objects where the reconstruction timescale is not a factor. Since the new line of DMD's operate at such high speeds, it is impossible for two-way communication to exist between the controlling computer and the DMD device, therefore, patterns must be preloaded onto on-board memory chip each time the DMD is initialised or the pattern set altered. There is little support for third party implementation of larger solid-state storage devices however this will change in the future allowing even greater resolutions to be preloaded onto the DMD.

These higher resolutions are the perfect example of how new image compression techniques will drive the ability of this technology to be applied to real world problems in competition with more traditional imaging devices. There have been recent examples of compression algorithms operating on megapixel images[66, 67, 68] with a small fraction of the total data and still producing high fidelity images. A combination of image compression, and ever faster computational power is now what sets the limits on this technology's capability

as a serious imaging alternative.

Chapter 6

Conclusions

6.1 Summary

In final conclusions, the single-pixel camera (SPC) has been successfully applied and tested in multi-channel, polarisation and ranging experiments. These implementations demonstrate the potential to be a viable alternative as a camera system, given the improvements shown in image quality and the capabilities of the underlying technology. It has been demonstrated that these applications, which previously required the use of a camera based on a detector array with spatial resolution, can be successfully supplanted with a single-pixel (SP) imaging system using computational correlation methods. This work detailed the progression of prototype SPC systems as each imaging application was investigated and where appropriate, compared to a traditional imaging system based on a detector with spatial resolution. Every SPC iteration improves on some fundamental aspect of the system, from final image resolution, to image acquisition rate. These core improvements opened new potential imaging applications for investigation, comprising some of the later work presented.

In contrast, some of the early systems shown were not capable of producing the images demonstrated in later results. A limiting factor for this included attempting to image objects with temporal variations in their spatial components with shorter time-scales than the acquisition time, or where the object had features smaller than the systems maximum resolution. As the final system described in this work is capable of fast acquisition, moderate resolution, and simultaneous multi-wavelength reconstruction, the SPC as a viable alternate imaging system was realised.

6.2 Multichannel Imaging

One of the primary goals and results of development in the system, is the ability for a SPC camera to image in discrete and wide-ranging portions of the electromagnetic (EM) spectrum. Constructing a camera based on a detector with spatial resolution requires a detection medium which is sensitive to the expected illumination bandwidth. It requires the ability to produce low-noise micro-scale components based on the chosen medium. These components need to be arranged into some spatial layout so that they receive the appropriate portion of a light-field, which has been refracted by an imaging optic placed in front of the system. This limits traditional cameras to specific wavebands, and causes the affordability and the ease of implementation of the camera system to vary drastically depending on factors such as commercial interests of a wavelength, and at extreme non-visible regions. In contrast, imaging based on a single-element detector only requires a single component which is sensitive to the expected illumination. This keeps the scale in affordability of any type of SPC camera system much more constrained. The underlying technology behind a SPC based on correlation methods is a spatial light modulator and the single-element detector. A SLM can be capable of a much wider band-

width operation than a traditional camera, and thus can be retained through multiple waveband ‘barriers’ such as visible and near infrared (NIR). Thus, only a new single element detector is required when moving between different portions of the EM spectrum. Different wavelengths within a single region of the EM spectrum can also be imaged simply by filtering a detector, adding additional detectors for multiple simultaneous wavelengths and changing the detection medium of the single element when the desired wavelength moves beyond the capabilities of the current set of detectors. This creates a very dynamic and widely applicable system, where the exact construction of the camera is dependent on the expected properties of the target object. It also gives rise to a disparaging comparison between detector array based cameras and a SPC in that the addition of detectors to the first, increase resolution, but addition of detectors to the second, increase the number of simultaneous wavelengths it is capable of measuring.

6.3 Polarisation SPC imaging

A primary example of an application which could only be undertaken when the system was robust enough to enable fast image acquisition was described in the polarisation sensitive SPC. In this iteration of the prototype. Polarising beam-splitters returned specific polarisation states to multiple detectors which form the basis for producing 2D images. The differential image created from the output of two cross polarised detectors is the basis for Stokes parameter imaging. This parameterisation of linear and non-linear polarisation states allows the calculation of many important polarisation parameters and is also relatively easy to directly measure from a physical system. This fast system capable of displaying a framerate of between 1Hz-10Hz allowed the probing of transparent media under stress as a dynamic stress factor was applied. Bire-

fringe effects could be observed altering as the stress factor was increased or decreased. In addition, changing polarisation states could be observed for a spinning polarising filter at various spin rates. This demonstration was to achieve similar results, for the moderate resolution of the SPC, that a four-camera polarising imaging system could obtain. This camera system generated the same set of images as the SPC, but was difficult to implement as it required near pixel-perfect alignment between the four CCD cameras. Since the direction of the observation angle was set by the SLM, and there was only one such angle, the SPC intrinsically did not need to be aligned in this manner. The images obtained from the SPC were always perfectly aligned. If the resolution of the images obtained by the SPC was deemed acceptable for polarisation imaging application, the natural perfect pixel alignment of the SPC was an advantage over traditional imaging systems.

6.3.1 Long Range SPC

The final prototype demonstrated in this work was the initial attempt at long range SPC imaging. The system was successfully adapted to function on a commercial telescope and take images of stationary and moving objects at distances between $5km - 20km$. The prototype system at this late stage in the work was now capable of (near) real-time video, displaying at framerates between 1Hz-10Hz and providing simultaneous multi-wavelength imaging. In this case, broadband visible and NIR. The initial investigation aimed to produce simultaneous visible and NIR images of a target dynamic object $> 5km$ through turbulent atmosphere. Since longer wavelengths of light are diffracted less, the NIR was expected to show less atmospheric turbulence than the visible image, however at this late stage, the moderate resolution of the camera was not high enough to show this expected result. A higher resolution SPC

camera would be able to confirm this expectation with a similar setup, and is a future goal of the technology. This is another case of an application of SPCs that is technology dependent rather than the need for better understanding or algorithm reconstruction.

6.4 SPC Future And Final Statement

Future SPC applications will most likely be tied to the upcoming technology improvements in DMD and SLMs. Even faster acquisition rates over short timescales are now possible, offering up to 36kHz burst projection rates. This allows the probing of objects which have shorter scale dynamic events, such as gas diffusion in a contained space. Compressed sensing algorithm improvement has enabled higher resolution images to be acquired with the same data overhead as images a tenth or less in total number of pixels. This represents a factor of 2-3 improvement over the CS results demonstrated in this work. These two avenues could potentially enable SPC to be applied to areas in which only a few years ago it would have been deemed a ‘no go’ area for SPC. Such applications could include testing of atmospheric turbulence at multi-wavelengths. You would need high resolution around the full native 1MP of the DMD chip and 50Hz reconstruction rate to achieve this. While we are not quite there yet, the fact that this lies in the realm of possibility for SPC is an exciting development.

This work has successfully demonstrated multiple, distinct practical applications of a SPC system. Improvements in the underlying technology and implementation led to the ability to investigate new imaging concepts. Some improvements such as acquisition time allowed temporally dynamic media to be investigated for the first time and highlights how far this technique has come

from the days of binary transmissive 2D objects fixed to an optical bench and in complete darkness. The final work on a long-range implementation further confirms this, as a truly robust imaging system can be used in multiple environments and in varying conditions.

Appendix A

Image Analysis

This appendix will include supplementary figures which demonstrate the thresholding process undertaken to obtain the SNR estimates provided in the main text, for those images where a quantitative analysis was done.

A.1 SNR Estimates

Since many reconstructions were processed in bulk, some additional images were analysed but not included here, as the outcome was extremely similar in each instance. Thus providing no additional benefit or insight in their inclusion. Therefore, some figures are instead a sample of images taken from the output of the image analysis program written in MATLAB for this body of work.

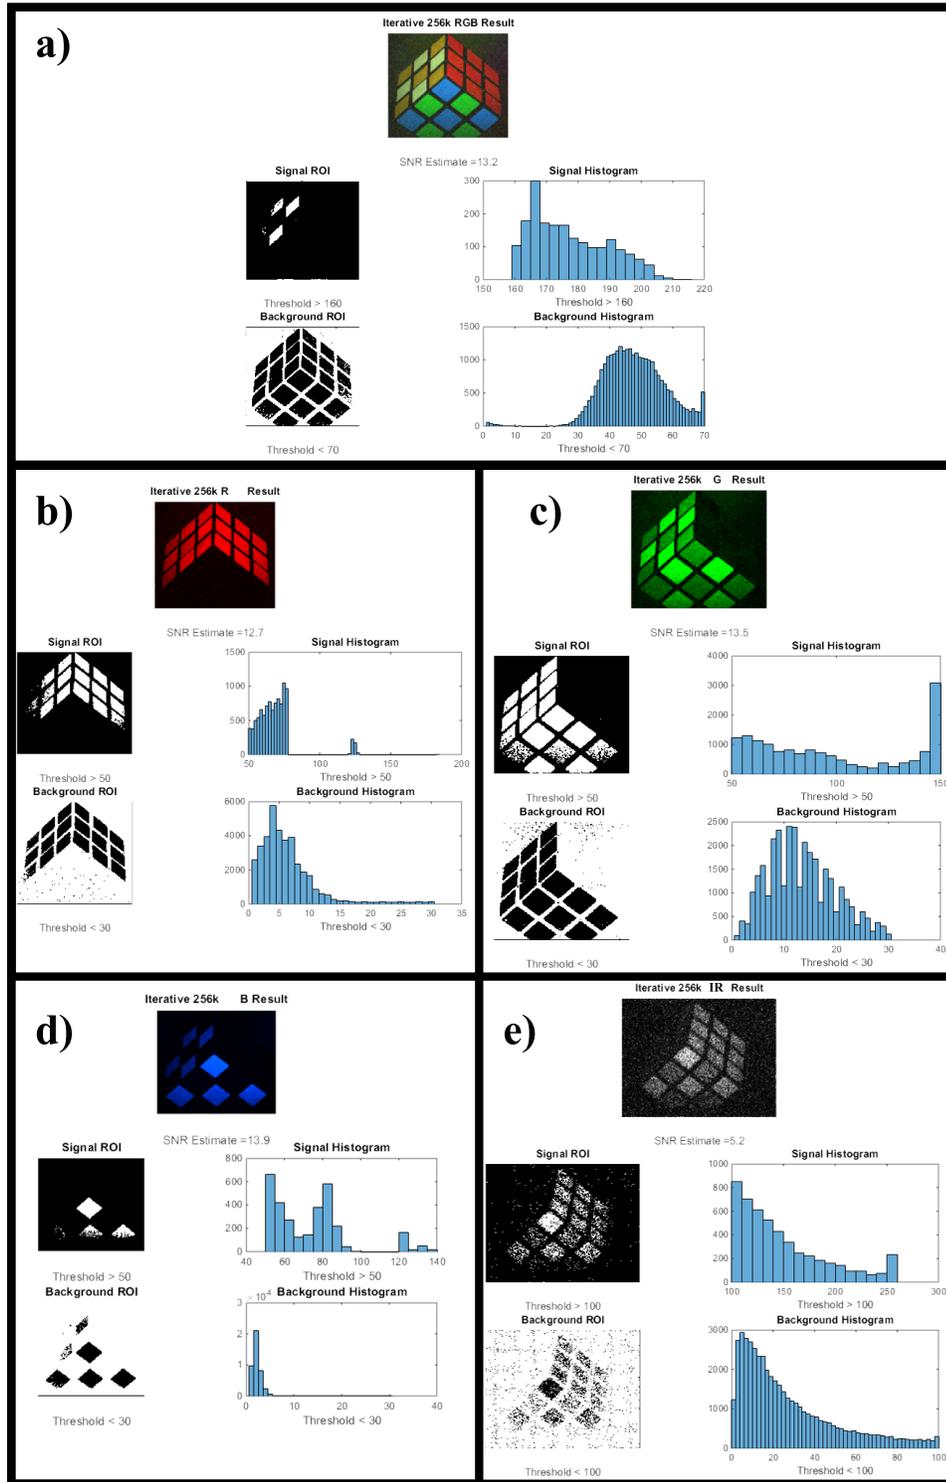


Figure A.1: *Signal to noise ratio estimation made by thresholding signal and background ROI in the RGB image data. This figure includes the analysis of the NIR result shown in Fig.3.7*

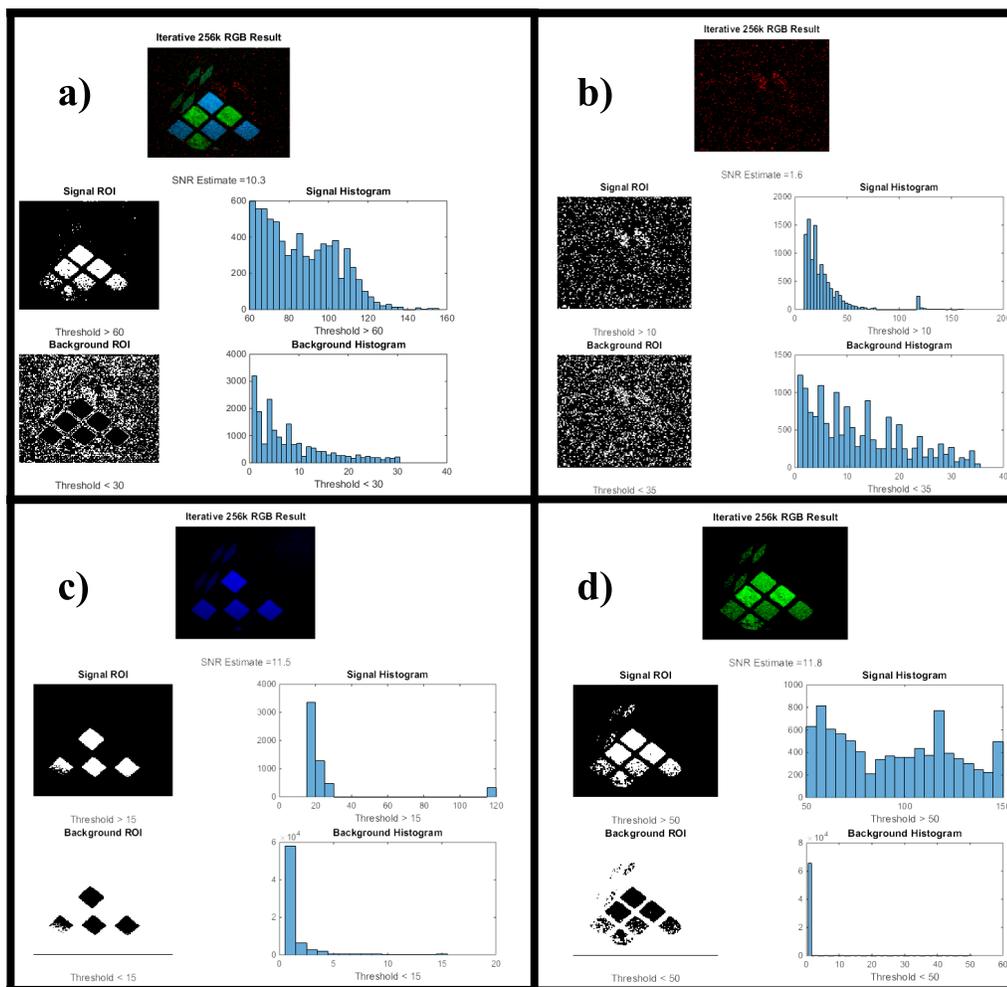


Figure A.2: *Signal to noise ratio estimates for the fluorescent SPC results from Fig.3.6. Red channel was too low to give a robust estimate on SNR and should be considered not fully representative of the SNR*

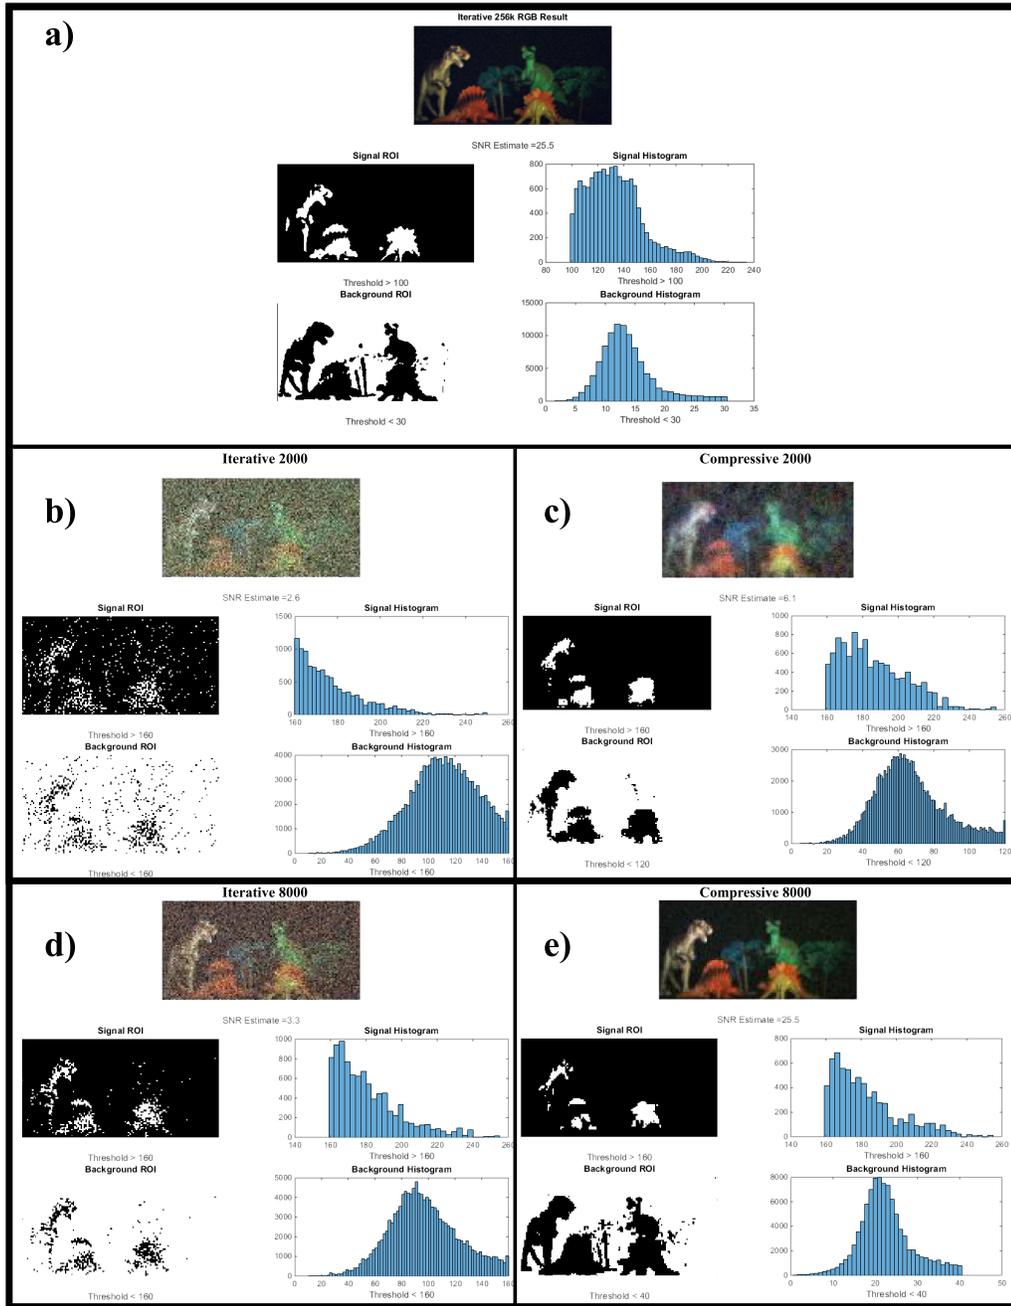


Figure A.3: Signal to noise ratio estimates for a selection of varying M value reconstructions featuring both iterative and compressive results from Fig.3.11. All other images for the interim M number results were calculated in a similar manner.

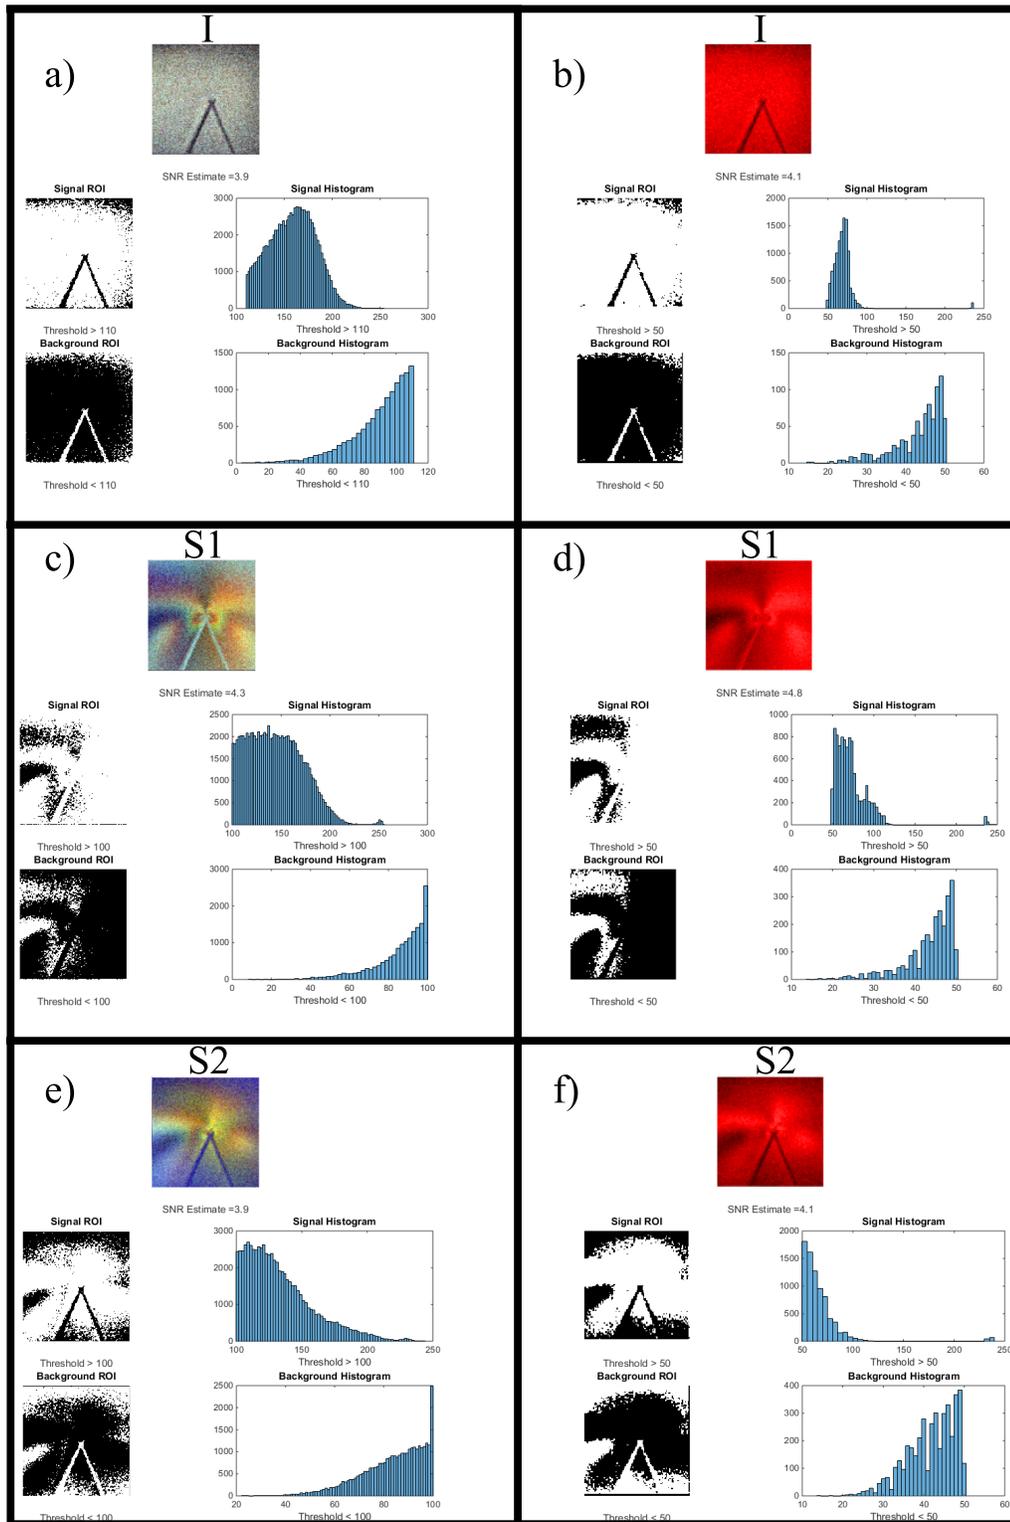


Figure A.4: Signal to noise ratio estimates for the RGB and RED colour channels of the Perspex sheet. Other Colour channels and polarisation illumination angles were calculated in a similar manner.

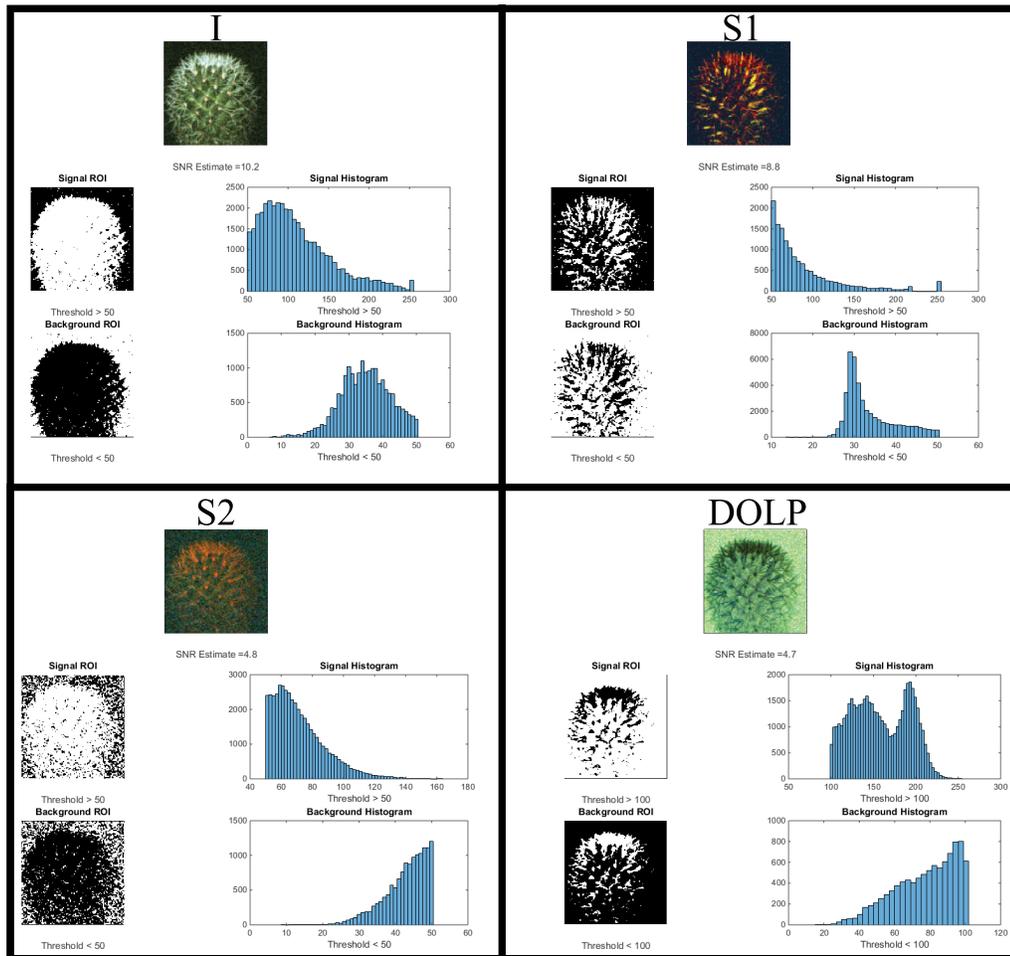


Figure A.5: *Signal to noise ratio estimates for the 0 degree linear polarisation illumination of the cactus with intensity, I , the first and second Stokes' parameters, $S1$ and $S2$ respectively. The Degree of Linear Polarisation (DOLP).*

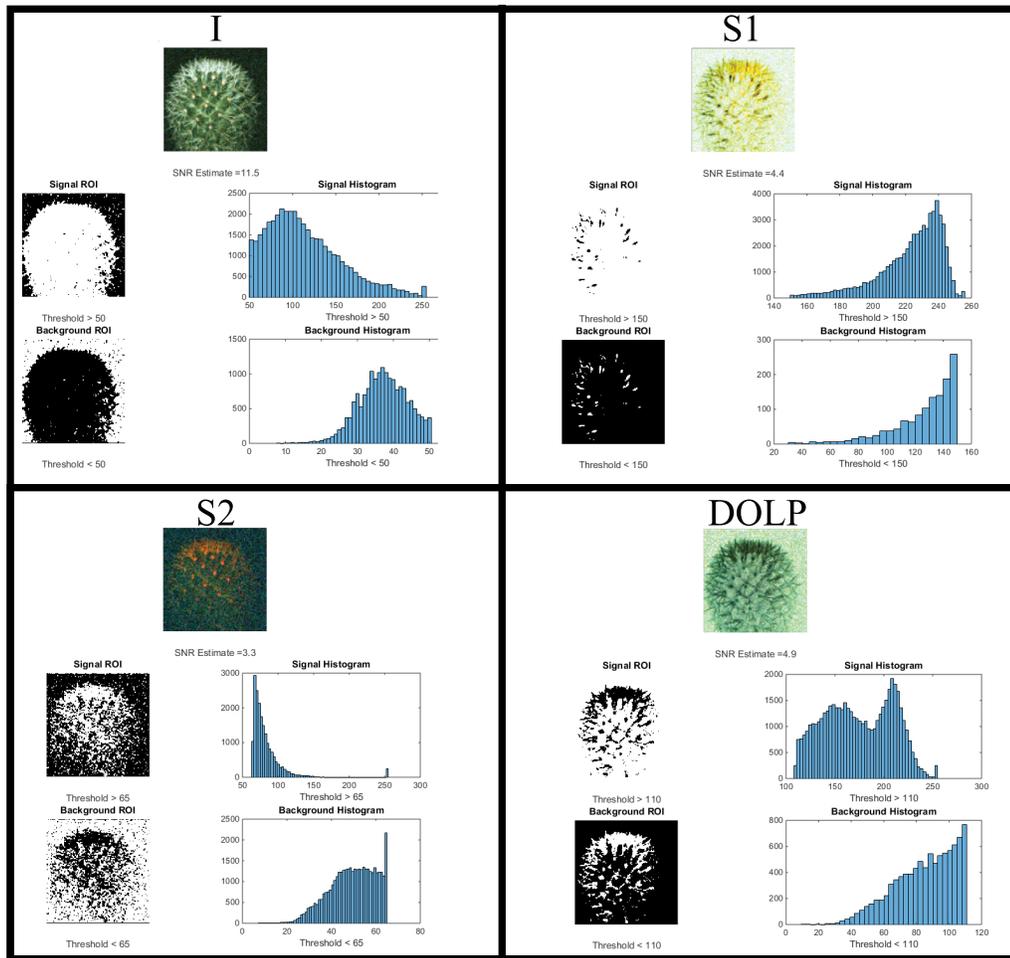


Figure A.6: *Signal to noise ratio estimates for the 90 degree linear polarisation illumination of the cactus with intensity, I , the first and second Stokes' parameters, $S1$ and $S2$ respectively. The Degree of Linear Polarisation ($DOLP$). The SNR estimates for the 45 and -45 linear illumination angle results were completed in a similar manner.*

Bibliography

- [1] Pierre Mertz and Frank Gray. A theory of scanning and its relation to the characteristics of the transmitted signal in telephotography and television. *Bell System Technical Journal*, 13(3):464–515, 1934.
- [2] Nicholas J Wade and Stanley Finger. The eye as an optical instrument: From camera obscura to helmholtz’s perspective. *Perception*, 30(10):1157–1177, 2001.
- [3] M. Minsky. Microscopy apparatus, December 19 1961. US Patent 3,013,467.
- [4] Robert H Webb. Confocal optical microscopy. *Reports on Progress in Physics*, 59(3):427, 1996.
- [5] Peter E. Powers, Thomas J. Kulp, and Randall Kennedy. Demonstration of differential backscatter absorption gas imaging. *Appl. Opt.*, 39(9):1440–1448, Mar 2000.
- [6] Baris I Erkmen and Jeffrey H Shapiro. Unified theory of ghost imaging with gaussian-state light. *Physical Review A*, 77(4):043809, 2008.
- [7] T. B. Pittman, Y. H. Shih, D. V. Strekalov, and A. V. Sergienko. Optical imaging by means of two-photon quantum entanglement. *Phys. Rev. A*, 52:R3429–R3432, Nov 1995.

-
- [8] D. V. Strekalov, A. V. Sergienko, D. N. Klyshko, and Y. H. Shih. Observation of two-photon “ghost” interference and diffraction. *Phys. Rev. Lett.*, 74:3600–3603, May 1995.
- [9] Giuliano Scarcelli, Vincenzo Berardi, and Yanhua Shih. Phase-conjugate mirror via two-photon thermal light imaging. *Applied Physics Letters*, 88(6):061106, 2006.
- [10] Lorenzo Basano and Pasquale Ottonello. Experiment in lensless ghost imaging with thermal light. *Applied Physics Letters*, 89(9):091109, 2006.
- [11] Ryan S. Bennink, Sean J. Bentley, and Robert W. Boyd. “two-photon” coincidence imaging with a classical source. *Phys. Rev. Lett.*, 89:113601, Aug 2002.
- [12] A. Gatti, E. Brambilla, M. Bache, and L. A. Lugiato. Ghost imaging with thermal light: Comparing entanglement and classical correlation. *Phys. Rev. Lett.*, 93:093602, Aug 2004.
- [13] A. Gatti, E. Brambilla, M. Bache, and L. A. Lugiato. Correlated imaging, quantum and classical. *Phys. Rev. A*, 70:013802, Jul 2004.
- [14] Alejandra Valencia, Giuliano Scarcelli, Milena D’Angelo, and Yanhua Shih. Two-photon imaging with thermal light. *Phys. Rev. Lett.*, 94:063601, Feb 2005.
- [15] F. Ferri, D. Magatti, A. Gatti, M. Bache, E. Brambilla, and L. A. Lugiato. High-resolution ghost image and ghost diffraction experiments with thermal light. *Phys. Rev. Lett.*, 94:183602, May 2005.
- [16] Jeffrey H. Shapiro and Robert W. Boyd. The physics of ghost imaging. *Quantum Information Processing*, 11(4):949–993, 2012.

-
- [17] Jeffrey H. Shapiro. Computational ghost imaging. *Phys. Rev. A*, 78:061802, Dec 2008.
- [18] Baoqing Sun, Stephen S. Welsh, Matthew P. Edgar, Jeffrey H. Shapiro, and Miles J. Padgett. Normalized ghost imaging. *Opt. Express*, 20(15):16892–16901, Jul 2012.
- [19] M.F. Duarte, M.A. Davenport, D. Takhar, J.N. Laska, Ting Sun, K.F. Kelly, and R.G. Baraniuk. Single-pixel imaging via compressive sampling. *Signal Processing Magazine, IEEE*, 25(2):83–91, 2008.
- [20] Vicente Durán, Pere Clemente, Mercedes Fernández-Alonso, Enrique Tajahuerce, and Jesús Lancis. Single-pixel polarimetric imaging. *Opt. Lett.*, 37(5):824–826, Mar 2012.
- [21] F Soldevila, E Irlles, V Durán, P Clemente, Mercedes Fernández-Alonso, Enrique Tajahuerce, and Jesús Lancis. Single-pixel polarimetric imaging spectrometer by compressive sensing. *Applied Physics B*, pages 1–8, 2013.
- [22] B. Sun, M. P. Edgar, R. Bowman, L. E. Vittert, S. Welsh, A. Bowman, and M. J. Padgett. 3d computational imaging with single-pixel detectors. *Science*, 340(6134):844–847, 2013.
- [23] J.B. Sampsell. Spatial light modulator, September 4 1990. US Patent 4,954,789.
- [24] L.J. Hornbeck. Spatial light modulator, September 11 1990. US Patent 4,956,619.
- [25] C.O. Bozler and S. Rabe. Spatial light modulator, July 21 1998. US Patent 5,784,189.
- [26] F. Ferri, D. Magatti, L. A. Lugiato, and A. Gatti. Differential ghost imaging. *Phys. Rev. Lett.*, 104:253603, Jun 2010.

- [27] Unser Michael Sage, Daniel. Teaching image processing programming in java. *IEEE Signal processing magazine*, 20(6):43–52, 2003.
- [28] Dennis Wageneck Mazzetta J, Caudle. Digital camera imaging evaluation. *EOI*, 8, 2005.
- [29] Herman M. Imaging and applied optics. *Optical Society of America*.
- [30] Matthew Herman. Compressive sensing with partial-complete, multiscale hadamard waveforms. In *Imaging and Applied Optics*, page CM4C.3. Optical Society of America, 2013.
- [31] Ling-An Wu Ming-Jie Sun, Zi-Hao Xu. Collective noise model for focal plane modulated single-pixel imaging. *OLE*, 100(4):18–22, 2018.
- [32] Ori Katz, Yaron Bromberg, and Yaron Silberberg. Compressive ghost imaging. *Applied Physics Letters*, 95(13):131110, 2009.
- [33] D.L. Donoho. Compressed sensing. *Information Theory, IEEE Transactions on*, 52(4):1289–1306, april 2006.
- [34] M Narasimha and A Peterson. On the computation of the discrete cosine transform. *IEEE Transactions on Communications*, 26(6):934–936, 1978.
- [35] Emmanuel J Candès, Justin Romberg, and Terence Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on information theory*, 52(2):489–509, 2006.
- [36] Stephen Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- [37] Michael Bass. *Handbook of Optics: Volume II*. McGraw-Hill Professional, 2010.

- [38] Pradeep Sen, Billy Chen, Gaurav Garg, Stephen R Marschner, Mark Horowitz, Marc Levoy, and Hendrik Lensch. Dual photography. In *ACM Transactions on Graphics (TOG)*, volume 24, pages 745–755. ACM, 2005.
- [39] Daryl Preece, Richard Bowman, Anna Linnenberger, Graham Gibson, Steven Serati, and Miles Padgett. Increasing trap stiffness with position clamping in holographic optical tweezers. *Optics express*, 17(25):22718–22725, 2009.
- [40] John V Frangioni. In vivo near-infrared fluorescence imaging. *Current opinion in chemical biology*, 7(5):626–634, 2003.
- [41] Daniel R Larson, Warren R Zipfel, Rebecca M Williams, Stephen W Clark, Marcel P Bruchez, Frank W Wise, and Watt W Webb. Water-soluble quantum dots for multiphoton fluorescence imaging in vivo. *Science*, 300(5624):1434–1436, 2003.
- [42] Mats GL Gustafsson. Nonlinear structured-illumination microscopy: wide-field fluorescence imaging with theoretically unlimited resolution. *Proceedings of the National Academy of Sciences of the United States of America*, 102(37):13081–13086, 2005.
- [43] Nian Tian, Qingchun Guo, Anle Wang, Dongli Xu, and Ling Fu. Fluorescence ghost imaging with pseudothermal light. *Opt. Lett.*, 36(16):3302–3304, Aug 2011.
- [44] Stephen S. Welsh, Matthew P. Edgar, Phillip Jonathan, Baoqing Sun, and Miles J. Padgett. Multi-wavelength compressive computational ghost imaging, 2013.
- [45] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.

-
- [46] Jean G Van Bladel. *Electromagnetic fields*, volume 19. John Wiley & Sons, 2007.
- [47] Mateusz Plewicki, Fabian Weise, Stefan M Weber, and Albrecht Lindinger. Phase, amplitude, and polarization shaping with a pulse shaper in a mach-zehnder interferometer. *Applied optics*, 45(32):8354–8359, 2006.
- [48] Mateusz Plewicki. *Phase, amplitude, and polarization pulse shaping in order to influence molecular processes*. PhD thesis, Freie Universität Berlin, 2007.
- [49] Edward Collett. *Field guide to polarization*, 2005.
- [50] Anthony Gerrard and James M Burch. *Introduction to matrix methods in optics*. Courier Corporation, 2012.
- [51] William H McMaster. Polarization and the stokes parameters. *American Journal of Physics*, 22(6):351–362, 1954.
- [52] C.A. Bennett. *Principles of Physical Optics*. Wiley, 2008.
- [53] Megan Paciaroni and Mark Linne. Single-shot, two-dimensional ballistic imaging through scattering media. *Applied optics*, 43(26):5100–5109, 2004.
- [54] Stavros G Demos and RR Alfano. Optical polarization imaging. *Applied Optics*, 36(1):150–155, 1997.
- [55] Roche M-Reinhardt J Pezzaniti J P Pezzaniti J L, Chenault D and Schultz H. Testing and results of an infrared polarized scene generator concept demonstrator. *SPIE Defense and Security Symposium (Bellingham, WA: International Society for Optics and Photonics)*.

- [56] Stephen S Welsh, Matthew P Edgar, Richard Bowman, Baoqing Sun, and Miles J Padgett. Near video-rate linear stokes imaging with single-pixel detectors. *Journal of Optics*, 17(2):025705, 2015.
- [57] John B Hearnshaw. *The analysis of starlight: one hundred and fifty years of astronomical spectroscopy*. CUP Archive, 1990.
- [58] David F Gray. *The observation and analysis of stellar photospheres*. Cambridge University Press, 2005.
- [59] Alasdair C Hamilton, Bhuvanesh Sundar, John Nelson, and Johannes Courtial. Local light-ray rotation. *Journal of Optics A: Pure and Applied Optics*, 11(8):085705, 2009.
- [60] J Courtial, N Chen, S Ogilvie, BC Kirkpatrick, AC Hamilton, G Gibson, T Tyc, E Logean, and T Scharf. Experimental demonstration of windows representing complex ray-optical refractive-index interfaces.
- [61] Stephen S Welsh and Johannes Courtial. View rotation with parallel ray-rotating windows. *Optics Communications*, 285(24):4802–4806, 2012.
- [62] Alasdair C Hamilton, Bhuvanesh Sundar, and Johannes Courtial. Local light-ray rotation around arbitrary axes. *Journal of Optics*, 12(9):095101, 2010.
- [63] Alasdair C Hamilton and Johannes Courtial. Generalized refraction using lenslet arrays. *Journal of Optics A: Pure and Applied Optics*, 11(6):065502, 2009.
- [64] Alasdair C Hamilton and Johannes Courtial. Optical properties of a dove-prism sheet. *Journal of Optics A: Pure and Applied Optics*, 10(12):125302, 2008.

-
- [65] Dean Lambert, Alasdair C Hamilton, George Constable, Harsh Snehan-shu, Sharvil Talati, and Johannes Courtial. Tim, a ray-tracing program for metatoy research and its dissemination. *Computer Physics Commu-nications*, 183(3):711–732, 2012.
- [66] Wai Lam Chan, Kriti Charan, Dharmpal Takhar, Kevin F Kelly, Richard G Baraniuk, and Daniel M Mittleman. A single-pixel terahertz imaging system based on compressed sensing. *Applied Physics Letters*, 93(12):121105, 2008.
- [67] Wen-Kai Yu, Ming-Fei Li, Xu-Ri Yao, Xue-Feng Liu, Ling-An Wu, and Guang-Jie Zhai. Adaptive compressive ghost imaging based on wavelet trees and sparse representation. *Optics express*, 22(6):7133–7144, 2014.
- [68] Marc Aβmann and Manfred Bayer. Compressive adaptive computational ghost imaging. *Scientific reports*, 3, 2013.