Compressive Spectral and Coherence Imaging

by

Ashwin Ashok Wagadarikar

Department of Electrical and Computer Engineering
Duke University

Date: _____

Approved:

_____
Dr. David Brady, Advisor

_____
Dr. Jungsang Kim

_____
Dr. David Smith

_____
Dr. Adam Wax

_____
Dr. Rebecca Willett

Dissertation submitted in partial fulfillment of the
requirements for the degree of Doctor of Philosophy
in the Department of Electrical and Computer Engineering
in the Graduate School of
Duke University

2010

ProQuest LLC
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106-1346

Abstract

Compressive Spectral and Coherence Imaging

by

Ashwin Ashok Wagadarikar

Department of Electrical and Computer Engineering
Duke University

Date: _____

Approved:

_____
Dr. David Brady, Advisor

_____
Dr. Jungsang Kim

_____
Dr. David Smith

_____
Dr. Adam Wax

_____
Dr. Rebecca Willett

An abstract of a dissertation submitted in partial fulfillment of the
requirements for the degree of Doctor of Philosophy
in the Department of Electrical and Computer Engineering
in the Graduate School of
Duke University

2010

Abstract

This dissertation describes two computational sensors that were used to demonstrate applications of generalized sampling of the optical field. The first sensor was an incoherent imaging system designed for compressive measurement of the power spectral density in the scene (spectral imaging). The other sensor was an interferometer used to compressively measure the mutual intensity of the optical field (coherence imaging) for imaging through turbulence. Each sensor made anisomorphic measurements of the optical signal of interest and digital post-processing of these measurements was required to recover the signal. The optical hardware and post-processing software were co-designed to permit acquisition of the signal of interest with sub-Nyquist rate sampling, given the prior information that the signal was sparse or compressible in some basis.

Compressive spectral imaging was achieved by a coded aperture snapshot spectral imager (CASSI), which used a coded aperture and a dispersive element to modulate the optical field and capture a 2D projection of the 3D spectral image of the scene in a snapshot. Prior information, such as piecewise smoothness of objects in the scene, could be enforced by numerical estimation algorithms to recover an estimate of the spectral image from the snapshot measurement.

Hypothesizing that turbulence between the scene and CASSI would introduce spectral diversity of the point spread function, CASSI's snapshot spectral imaging capability could be used to image objects in the scene through the turbulence. How-

ever, no turbulence-induced spectral diversity of the point spread function was observed experimentally. Thus, coherence functions, which are multi-dimensional functions that completely determine optical fields observed by intensity detectors, were considered. These functions have previously been used to image through turbulence after extensive and time-consuming sampling of such functions. Thus, compressive coherence imaging was attempted as an alternative means of imaging through turbulence.

Compressive coherence imaging was demonstrated by using a rotational shear interferometer to measure just a 2D subset of the 4D mutual intensity, a coherence function that captures the optical field correlation between all the pairs of points in the aperture. By imposing a sparsity constraint on the possible distribution of objects in the scene, both the object distribution and the isoplanatic phase distortion induced by the turbulence could be estimated with the small number of measurements made by the interferometer.

This dissertation is dedicated to my beloved grandparents.

Contents

List of Tables

List of Figures

Acknowledgements

I owe my deepest gratitude to Dr. David Brady for welcoming me into the Duke Imaging and Spectroscopy Program (DISP) in the fall of 2005 and for his guidance, judgement and support throughout my graduate career. With his patience, he has offered me the privilege of completing this work while allowing me to explore a variety of extracurricular passions. His immense capacity to generate new ideas has been truly inspirational.

Dr. Kerkil Choi, Dr. Daniel Marks, Dr. Joonku Hahn, Dr. Nathan Nagen and Dr. Renu John have all provided advice and insight through frequent conversations about my work. Kerkil, thanks for your infectious laugh and for being patient with me so many times. Joonku, thanks for helping me figure out how to align the RSI. I could not have done it without you. Dan, thanks for introducing me to coherence and for helping me with my preparation prior to my October 2009 trip to JPL. Nathan, thanks for helping me to realize the amount of effort that is required to calibrate an optical instrument. Renu, thanks for your help in putting the first CASSI prototype together.

I would like to thank my committee members, Dr. Jungsang Kim, Dr. David Smith, Dr. Adam Wax and Dr. Rebecca Willett for their time and feedback throughout this process. I would also like to thank my collaborators, with whom I have had fruitful discussions and published papers including Dr. Rebecca Willett, Dr. Nikos Pitsianis, Dr. Xiaobai Sun and Ryoichi Horisaki. Thanks also to Dr. Sudhakar

Prasad for useful discussions about the problem of imaging through turbulence.

In my daily work, I have been blessed with a friendly and cheerful group of fellow students and fellow members of DISP. I want to thank the students of DISP, both past and present, who have been a great source of motivation, encouragement, friendship and laughter. Thanks to Evan Cull, Scott McCain, Mohan Shankar, Andrew Portnoy, Christy Fernandez Cull, Sehoon Lim, David Kittle, Alex Mrozack, Nan Zheng and Ryoichi Horisaki for making my time with DISP so memorable. Special thanks to Paul Vosburgh who designed, machined and rapid-prototyped countless parts for me, mostly because I was too lazy to do it myself. Thanks to Steve Feller, someone I could always rely on to make me laugh, to help keep life in perspective and for all things related to the administration of DISP computers. Thanks also to Leah Goldsmith, the unofficial lab Mom.

For their support in administrative issues, I thank Autumn Wenner, Samantha Morton, Leah Goldsmith, Ellen Currin and August Burns. I gratefully acknowledge funding sources including Duke University, DARPA, AFOSR and Canada's Natural Sciences and Engineering Research Council (NSERC) for supporting my work with a postgraduate doctoral fellowship. Thanks also to SPIE for awarding me an educational scholarship in 2009.

Beyond research, a number of key friendships have helped me to maintain my sanity during this process. Vince Mao, Randy Evans, Leeanne Cox, Vladi Ivanov and Andrew Fontanella made CIEMAS 3515 the best office environment in the building there ever was. I have always been able to count on Vladi Ivanov, Andrew Fontanella and Sara Salahi for support. Thank you guys for the countless number of walks we have taken together around Duke's beautiful campus. Thanks to Michael Giacomelli for getting me out of the office and into the gym three times a week without fail starting in Spring 2008.

Over the past 14 months, my wonderful girlfriend Victoria Vanlandingham has

been a great source of motivation, strength, support and love. I am convinced that while I might have still completed this work, it would have taken many more months and I certainly would not have been as happy doing it. Thanks for being so awesome, sweetheart!

I would not be where I am today without the endless love and support of my family. My parents, Ashok and Mukul, and my brother Atul, have always believed in me and given me encouragement when I've needed it. I also want to acknowledge my late grandmother who was never short on advice for eating well to stay healthy and loved me in her own special way.

Finally, I would like to thank Duke University for making this place so conducive for gaining knowledge and doing wonderful things. This experience has been an amazing one and I am proud to say that I will soon be a Duke graduate.

1

Introduction

1.1 Computational optical imaging

Electronic imaging has revolutionized optical sensors. It has mostly replaced photochemical imaging because of the ease with which digital data can be stored and recalled. Optical sensors using electronic imaging devices as detectors, i.e. digital optical sensors generate a mapping between discrete samples of the particular signal of interest in an optical field (such as images or spectra) and discrete measurements on a detector array. This mapping may be written as

$$g_i = \int \phi_i(x) f(x) dx, \tag{1.1}$$

where $f(x)$ represents the signal, $\phi_i(x)$ represents the impulse response of the optical system that maps points in the signal to the detector measurement g_i. If the continuous signal can be represented with discrete samples f_n on a basis $\psi_n(x)$ such that

$$f(x) = \sum_n f_n \psi_n(x), \tag{1.2}$$

then equation (1.1) may be written in matrix-vector form as

$$g = \Phi f, \tag{1.3}$$

where g represents a vector of detector measurements, $\phi_{ij} = \int \phi_i(x)\psi_j(x)dx$, and f is a vector of discrete samples of the signal of interest.

Conventional optical sensors using electronic imaging include digital cameras and spectrometers. Such instruments have revolutionized several fields including machine vision, medical diagnostics, chemometric analysis and astronomy. They create an *isomorphic* mapping between the particular signal of interest in an optical field and the detector array by relying on optical elements such as lenses and gratings. The optical elements are a means of implementing analog processing on the optical field prior to the digitization of the image on the detector. Once recorded, the detector measurements can be easily displayed and analyzed.

While conventional optical sensors may directly acquire signals of interest such as the spectrum of a light source or an eye-pleasing image, improvements to their architectures and performance rely solely on advances in hardware and may be limited by the physical laws of nature such as diffraction. For example, consider a conventional dispersive spectrometer consisting of a slit aperture, a grating and a line detector array that directly records the spectrum of light incident on the aperture. While one may want to increase the width of the slit to increase the throughput of light onto the detector for increased signal-to-noise ratio (SNR), maintaining high spectral resolution necessitates that the width of the slit is as narrow as possible. Increasing the light throughput while maintaining spectral resolution would thus require much longer exposure times. A more ideal spectrometer would be one that would use a large entrance aperture while maintaining the spectral resolution of a slit spectrometer.

In some cases, hardware advances add only marginal value to conventional optical

sensors. For example, current digital cameras capture images with high resolution (more megapixels), only to have most of the information discarded due to storage requirements and because compression algorithms can maintain an accurate representation of an image with significantly fewer bits. If it's possible to compress an image, it's easy to argue that capturing a many megapixel image is rather redundant. A more ideal camera would only require enough pixels to capture the salient information in an image, i.e. the information determined by compression algorithms as necessary for generating a compressed image that appears is free of artifacts when viewed on a computer display.

In contrast to conventional optical sensors, computational optical sensors, or simply computational sensors, make *anisomorphic* mappings between the signal of interest and the detector measurements [1]. The design of such sensors involves jointly optimizing optical hardware with digital post-processing of the detector measurements [2]. The measurements are coded by the optical elements in front of the detector and thus cannot be as easily deciphered as the isomorphic measurements made by conventional optical sensors. By design, computational sensors rely on digital processing of these measurements to extract the desired signal of interest. By leveraging the power of software, a variety of hardware limitations experienced by conventional optical sensors can be overcome. For example, a coded aperture spectrometer replaces the entrance slit of a conventional dispersive spectrometer with a much wider (hence greater throughput) coded aperture and measures a convolution of the incident source spectrum and the coded aperture pattern on a two dimensional (2D) detector array. Post-processing of the detector measurements can produce an estimate of the source spectrum with equivalent spectral resolution to a slit spectrometer, but with considerably higher SNR due to the greater throughput of the instrument.

Just as software can be used for overcoming hardware limitations, the optical

hardware in a computational sensor may employ diverse optical coding elements to implement a sensing strategy that is suited towards a particular digital processing inversion algorithm used to recover the signal of interest from the detector measurements.

1.2 Generalized sampling

The sampling process used to obtain discrete, isomorphic samples of a continuous signal in conventional optical sensors relies on the Nyquist-Shannon sampling theorem. This theorem assumes that the signal of interest is bandlimited and specifies that to avoid losing information in the signal through aliasing, one must sample at a rate that is at least two times faster than the signal bandwidth. As described earlier, in cases such as focal imaging using a digital camera, the Nyquist rate may be so high that far too many samples result for transmission and storage.

In contrast, the co-design of hardware and software processing algorithms for computational sensors can enable more general sampling strategies. For instance, the implementation of a strategy known as compressive sampling allows the acquisition of a signal at sub-Nyquist rate *given* the prior information that the signal is sparse or compressible in some basis other than the sampling basis. This suggests that the signal is dependent on a smaller number of degrees of freedom than the number dictated by its bandwidth. Application of compressive sampling in optical imaging may exploit the fact that images have structure and tend to be sparse in some basis other than the Dirac sampling basis.

This dissertation provides examples of computational sensors which sample the signal of interest through an integrated approach to optical and algorithm design. This enables the powerful ability to potentially eliminate dimensional tradeoffs such as the need to temporally scan an instrument to capture a multi-dimensional (more than two) signal of interest. In essence, such sensors measure a compressed and

efficient representation of the signal that can be decoded digitally through post-processing.

1.3 Motivation

This dissertation focuses on two computational sensors that demonstrate applications of generalized sampling of the optical field. The first sensor, described in chapters 2, 3 and 4, is a spectral imaging system that is used to compressively measure the power spectral density of a scene. While there are several approaches to spectral imaging, most require some temporal scanning to capture the desired information. In contrast, the sensor described here demonstrates how the introduction of coding into optical hardware and the use of non-linear numerical estimation methods for processing the detector measurements can potentially permit a snapshot approach to capturing a spectral image.

The second sensor, described in chapter 5, is used to image sparse objects through turbulence by sampling a small subset of the mutual intensity of the optical field in its aperture. While there are several approaches to imaging through turbulence, most require expensive equipment or post-processing of several captured images in batch mode. In contrast, the proposed sensor only consists of a beam splitter, two fold mirrors and a detector array, and it uses a numerical estimation method that exploits the sparsity of objects in the scene. It has the potential to image sparse objects in space through the turbulence of Earth's atmosphere in real time.

1.4 Organization

This section outlines the structure of this dissertation and summarizes the work completed towards this dissertation, which began in March 2007. It also identifies people, besides myself and my advisor Dr. David Brady, who have contributed significantly to the work described in each chapter.

Chapter 2 begins with a review of spectral imaging techniques. The focus of this chapter is the implementation of a coded aperture snapshot spectral imager, or CASSI for short. I was first introduced to CASSI as a general class of spectral imaging systems by Dr. Brady in January 2007. After inventing a snapshot spectral imaging system with a dual disperser architecture [3], Dr. Brady talked to me about the possibility of achieving snapshot spectral imaging with a system containing a coded aperture and only a single disperser. Such a system would be an extension of the coded aperture spectrometer, a system I studied while completing my MS thesis work, to spectral imaging. Following through on his ideas and suggestions, I worked on building a proof-of-concept CASSI with a single disperser architecture. Experimental results obtained by using this prototype are described in this chapter. The results were first presented to the scientific community at SPIE's Optics and Photonics conference in San Diego, CA in August 2007 [4], and later published in *Applied Optics*' special issue on Computational Optical Sensing and Imaging (COSI) in May 2008 [5]. A figure from the paper also appeared on the cover of that issue. Dr. Renu John helped me with building CASSI in the lab, while Dr. Rebecca Willett contributed by proposing the GPSR algorithm and assisting me with datacube recovery from CASSI measurements.

The experience of building, characterizing, and processing measurements made by the first CASSI prototype was valuable. It led to the conception of a second generation CASSI prototype after recognizing the need for a custom-design system with good image quality. After construction, this prototype was used to capture snapshot spectral images with better image quality than the previous prototype, in part due to more rigorous system calibration. The system design, calibration process and experimental results are described in chapter 3. The system was used to capture data for a comparative analysis of reconstruction algorithms to process CASSI measurements. This work was presented as an invited talk at SPIE's Optics and Photonics

conference in San Diego, CA in August 2008 [6]. The snapshot ability of this CASSI was then used to capture spectral images of a dynamically changing scene at video rate. This work was published in the journal *Optics Express* in May 2009 [7]. Dr. Nathan Hagen helped me to significantly improve the CASSI calibration process. Dr. Jose Bioucas Dias, Dr. Mario Figueiredo and Joao Oliveira helped me to adapt the TwIST algorithm for spectral imaging with CASSI. The NeAREst algorithm for datacube recovery from CASSI measurements was conceived and implemented by Dr. Xiaobai Sun and Dr. Nikos Pitsianis.

While CASSI was conceived from recent ideas on compressive sampling, its performance and limitations had not been analyzed based on ideas in the compressive sampling literature. Chapter 4 begins with a background on compressive sampling and noiseless Matlab simulations are used to analyze the performance of CASSI as an example of a compressive sampling instrument. The simulations also provide insight for an alternative design for CASSI, that was first proposed by Ryoichi Horisaki, a visiting student to the group. Finally, the performance of CASSI and the alternative design is studied in the presence of noise. The work described in this chapter was conducted under the guidance of Dr. Kerkil Choi and was submitted for publication to the journal *IEEE Transactions on Image Processing* in July 2009.

To demonstrate an application of CASSI, the system was briefly considered as a means to image remote objects through atmospheric turbulence. Hypothesizing that turbulence between remote objects and CASSI would introduce spectral diversity of the point spread function of the imaging system, CASSI's snapshot spectral imaging capability could be used to image objects in the scene through the turbulence. However, no turbulence-induced spectral diversity of the point spread function was observed experimentally. Thus, a new approach to imaging through turbulence was investigated that relied on compressive measurement of the mutual intensity of the optical field, rather than the power spectral density. In this approach, prior knowl-

edge about the sparsity of the scene is assumed that parallels the assumption made for CASSI. Chapter 5 describes the use of a rotational shear interferometer (RSI) to image objects through turbulence by measuring just a 2D subset of the 4D mutual intensity function. A alternating minimization algorithm is used to digitally process these measurements to jointly estimate the image of sparse point sources in the scene along with the turbulence induced phase screen in the aperture. This work was presented at the Optical Society of America's topical meeting on Computational Optical Sensing and Imaging (COSI) in San Jose, CA in October 2009 [8]. It was later accepted for publication in *Optics Letters* in February 2010 [9]. Dr. Daniel Marks was a key resource as he first introduced me to the intricacies of the RSI, which is not an easy instrument to use. Dr. Joonku Hahn helped me greatly with aligning the RSI in the lab. Dr. Kerkil Choi proposed the alternating minimization algorithm used to image objects through turbulence from RSI measurements. Ryoichi Horisaki assisted me with implementing the algorithm.

Finally, Chapter 6 summarizes the work described in this dissertation and identifies some open questions for some future avenues of research.

2

Spectral imaging using a coded aperture snapshot spectral imager

The focus of this chapter is a new instrument for spectral imaging called a Coded Aperture Snapshot Spectral Imager (CASSI). The chapter begins with an introduction to the problem of spectral imaging in section 2.1 and a review of diverse spectral imaging techniques in section 2.2. Section 2.3 provides an overview of CASSI systems. Section 2.4 introduces the new CASSI design and briefly compare it to a previously developed CASSI design. Section 2.5 details a mathematical model for the new design. Section 2.6 describes the design of the first experimental prototype of the new CASSI design. Section 2.7 describes how the Gradient Projection for Sparse Reconstruction method can be used to reconstruct a datacube from a CASSI measurement and section 2.8 demonstrates its use in recovering a spectral image of a static scene.

2.1 Introduction to spectral imaging

In many imaging applications, color increases the information content of the objects in the scene being imaged. In biomedicine, color signals the biochemical makeup of different regions of tissue. When coupled to more targeted indicators such as antibodies, this can allow the detection of multiple, specific analytes [10].

However, the native ability of our eyes to perceive and evaluate color information is limited to the visible range, roughly 380 to 750 nm, of the electromagnetic spectrum. This limitation extends to conventional digital cameras that typically consist of a color filter array of red, green and blue filters on top of a monochromatic imaging sensor array [11]. Thus, all spectral information within the visible spectral range is binned into three broad spectral ranges roughly corresponding to the three primary colors, red, green and blue. While such information assists in perceiving shape, surface texture and mutual spatial relation in the depth of 3D space, it is insufficient in applications where it is necessary to observe and study the accurate spectral signature of objects in the scene. Such signatures allow discrimination of specific object properties. In addition, the spectral information of interest often extends to the invisible range of the electromagnetic spectrum.

Spectral imaging combines 2D imaging and spectroscopy. In contrast to the conventional three-color imagers (i.e. digital cameras), a spectral imager attempts to precisely measure an optical spectrum at every pixel in an image. The spectral range of the instrument can be designed to span a particular region of the electromagnetic spectrum with some desired spectral resolution. The output of the instrument is a 3D dataset, known as the datacube, containing the two spatial dimensions and one spectral dimension. It can be visualized as a data set consisting of a number of images representing the brightness at each pixel as a function of wavelength [10]. Figure 2.1 (reproduced from [12]) demonstrates the spectral imaging concept by

showing different spectra at different locations in an image.

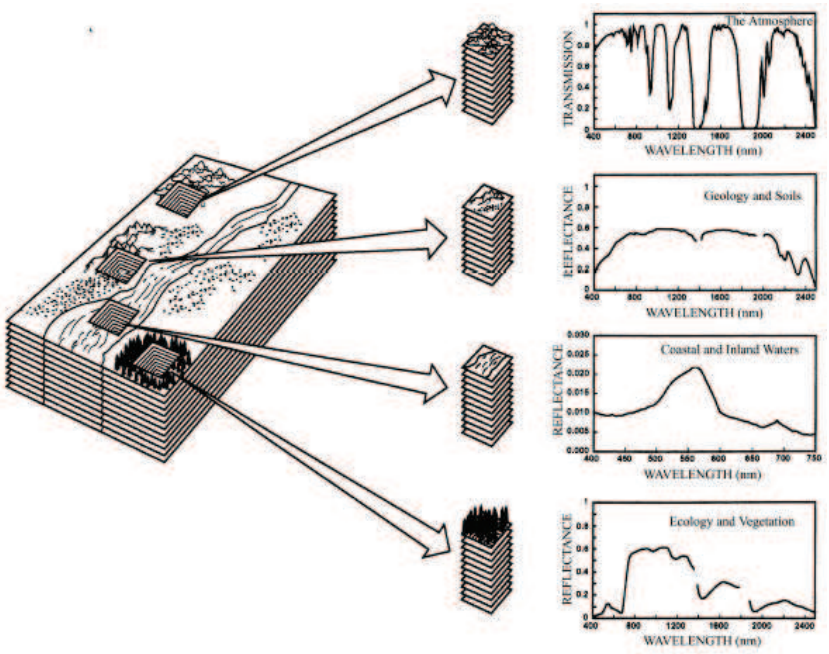

FIGURE 2.1: The concept of spectral imaging is illustrated with spectra gathered at four different spatial locations in an image.

2.2 Approaches to spectral imaging

There are many different types of spectral imagers. They may roughly be divided up into three types including (i) dispersive imagers, (ii) interferometric imagers and (iii) tunable filtered imagers. Figure 2.2 (reproduced from [13]) demonstrates the different subsets of the datacube that conventional imagers are able to measure without any scanning. Due to the inherent dimensionality difference between a 3D datacube and a 2D detector array, these imagers rely on temporal scanning to measure all the voxels, small unit volume elements in the datacube.

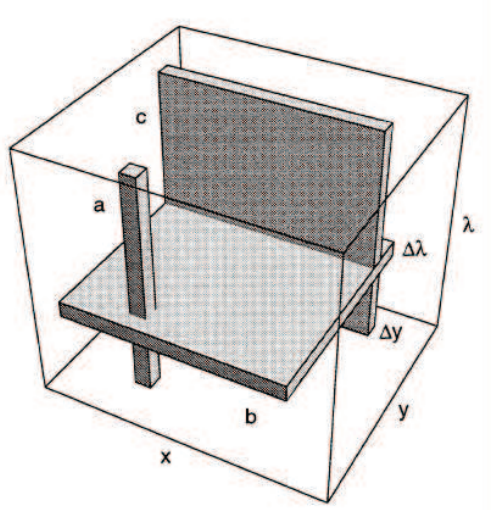

FIGURE 2.2: Lower dimensional volumes that simple spectral imagers are able to capture without any scanning. (a) represents a whiskbroom instrument, (b) represents a filtered camera or a Fourier transform instrument, and (c) represents a pushbroom instrument.

2.2.1 Dispersive spectral imaging

Whiskbroom spectral imager

The simplest method to collect a datacube is to use a spectrometer that uses a grating and a linear array of detectors and samples a very small spatial area, thus measuring voxels in the column **a** as shown in figure 2.2. Such a system is known as a whiskbroom spectral imager. The datacube is built up over time by scanning the instrument over a desired spatial area. The best known instrument of this kind is Airborne Visible/Infrared Imaging Spectrometer (AVIRIS), developed in 1983 at the Jet Propulsion Lab (JPL), Pasadena, California, USA. It was the first hyperspectral imager to measure the solar reflectance spectrum from $400 - 2500$ nm. The scanner had 224 channels at 10 nm intervals across the spectrum [12]. The main problem with such an instrument is the amount of time it takes to temporally scan the entire datacube and the poor light collection efficiency of the instrument.

Spatially filtered pushbroom spectral imager

Another type of spectral imager uses pushbroom scanning, capturing spectral data along one spatial dimension in each time step. Fore-optics are used to image the scene on to a field-limiting entrance slit. The entrance slit and a dispersive element allow the 2D detector to sample the spectral dimension and one spatial dimension of the datacube simultaneously. The second spatial dimension is captured by scanning the slit in the direction of motion. The output of each dataset collected from a pushbroom device is a one-dimensional (1D) spectral image and produces information which is described by a plate in a spatial dimension, **c** as shown in figure 2.2. The disadvantage of such a technique is that relative motion in the scene during the process of sampling the second spatial dimension can result in a distorted or warped image. In addition, the use of a slit means that out-of-slit photons are discarded, thus limiting the light throughput of the instrument.

Coded aperture pushbroom spectral imager

It is also possible to use a coded aperture spectrometer [14, 15] as a pushbroom spectral imager. In this instrument the entrance slit of a dispersive spectrometer is replaced with a coded aperture. The instrument measures a convolution of the source spectrum with the aperture code in a snapshot and post-processing of the recorded detector data is required to generate an estimate of the source spectrum. The light throughput and the corresponding signal strength measured by the modified instrument is significantly larger, while the spectral resolution is identical to that of a slit spectrometer. If the horizontal distribution of light incident on the aperture code is non-uniform, then the spectrometer behaves as a 1D spectral imager. This characteristic can be used to extend the utility of a coded aperture spectrometer to pushbroom 2D spectral imaging by scanning the instrument perpendicular to the direction of dispersion [16]. Note that this is in contrast to a slit-based pushbroom

spectral imager, which scans parallel to the direction of dispersion.

Tomographic spectral imager

More computationally intensive solutions to the light collection problem have also been developed. Mooney et al. [17] developed a direct-view design that maximizes the light gathering efficiency by not requiring any field stop (no SLM or slit etc.). With this design, the source is viewed through a rotating dispersive element. Measurements are taken at different rotation angles. These measurements are projective measurements through the datacube that can be tomographically reconstructed. While the light gathering efficiency of such an instrument is high, the geometry of the system limits the range of angles over which projections are made. Based on the Fourier slice theorem, this results in an unsampled conical region in Fourier space (the missing cone problem) [18].

2.2.2 Interferometric spectral imaging

A Fourier transform spectral imager uses a Michelson interferometer in the pupil plane of a conventional imaging system. In a Michelson interferometer, input light from an aperture is collimated using a lens and split into two paths. Both paths are retro-reflected on to a detector using flat mirrors. Scanning one of the two mirrors in the interferometer varies the path difference between the two arms of the interferometer and produces an interferogram at each pixel on the detector array as a function of the temporally scanned path difference. Performing a Fourier transform of interferograms produces a datacube from the recorded interferogram cube [19]. Such an instrument is useful in low light applications, as it provides higher signal to noise ratios than the spectrally filtered imagers [19]. The main disadvantage of such an instrument is that it requires moving parts.

2.2.3 Tunable filter spectral imaging

Examples of spectrally filtered imagers include cameras filtered by acousto-optic tunable filters (AOTFs) and liquid crystal tunable filters (LCTFs). AOTFs are birefringent crystals that provide an electronically tunable spectral notch passband in response to an applied acoustic field. In addition, AOTFs allow the simultaneous tuning of multiple spectral passbands, allowing multiplex measurements of the datacube if desired. LCTFs also provide a notch passband that can be controlled by incorporating liquid crystal waveplate retarders within the birefringent filter. Both types of filters provide wavelength selection and tunability. AOTFs provide higher transmittance and higher free spectral range than LCTFs at comparable spectral bandpass, but do not provide the imaging quality that LCTFs do [20]. Spectral imagers using either type of filter have the advantage of requiring no moving parts. However, they have the drawback that the spectra for a given pixel are not taken simultaneously. This can result in pixel-to-pixel spectral mixing if the scene is moving while the spectra are being scanned. The output of one image of a filtered camera is a plate, $\mathbf{b}$, as shown in figure 2.2, whose thickness is determined by the spectral bandpass of each filter in the filtered camera.

2.2.4 Snapshot spectral imaging

While all of the spectral imager designs discussed so far involve temporal scanning to measure all the voxels in the datacube, examples of snapshot instruments have also been developed. The elimination of any sort of scanning to capture a datacube is motivated by two major factors. First, the absence of moving parts ensures that the system layout of the spectral imager becomes fixed, which results in a rugged design. Second, instantaneous collection of both spectral and spatial information of the object space means that dynamic events within the field of view can be captured and is limited only by the exposure time of each frame [13]. Applications of a snap-

shot spectral imager include observation of fluorescent probes that track fast (0.01-1 s) cellular metabolic processes under a microscope [21], surveillance of fast moving space objects [22], and snapshot spectral imaging of the retina that is unaffected by patient motion [23].

One trend in snapshot spectral imaging is to rely solely on optical hardware to redirect voxels of the datacube to different pixels on a detector array and requires negligible post-processing for datacube recovery. An example of a system applying this approach is the 4D imaging spectrometer [24] that uses reformatting fiber optics to map a 2D image of the scene to a linear array that serves as an input slit to an imaging spectrometer. As another example, Harvey et al. [19] use a Lyot filter that simultaneously achieves spectral filtering and image replication using a sequence of polarizing beam splitters. Filtering is performed without any light rejection. The detector placed behind the Lyot filter captures multiple narrow-band non-overlapping images of the scene on a single detector array. This design is limited by the number of spectral channels that can be measured while preserving the property of non-overlapping images on a single detector array. As another example, an image slicing spectrometer [25] is a system that uses a custom-made mirror consisting of long strips with different tilts to slice and redirect the image of the scene onto non-overlapping zones on the detector array. A prism disperses the sliced image zones into their neighboring void regions. Since this approach establishes an isomorphic mapping between voxels in the datacube and the pixels on the detector array, simple image remapping is sufficient to recover a datacube. Potential limitations include the need for an expensive large format charge-coupled device (CCD), an array of reimaging lenses and custom fabrication of the image slicing mirror assembly.

Another trend is to rely on optical hardware to record a 2D coded projection of the 3D datacube on the detector and using post-processing of the measurements to recover the 3D datacube. A computed tomography imaging spectrometer (CTIS) [26]

is a static, snapshot instrument that captures multiple projections of the datacube at once. Nearly all the collected light is measured on the detector [23]. However, the instrument requires a large focal plane area and also suffers from the missing cone problem [27]. The missing cone problem for the CTIS implies that it is difficult to reconstruct datacubes that contain little spatial structure and sharp spectral features [28].

Another strategy that leverages both hardware and software while reducing the size of a snapshot spectral imager involves using pixelated spectral filters placed immediately in front of the detector array. Different linear combinations of the spectral channels are captured at each pixel. The simplest example of a pixel filtering system is a 3-channel red-green-blue (RGB) imager with RGB filters distributed in a Bayer pattern. Simple algorithms can be used to interpolate the RGB datacube from the Bayer sampled measurements. This strategy of integrating pixel-level filters with detectors may also be extended to more than three colors along with more sophisticated algorithms for datacube recovery [1].

2.3 Coded aperture snapshot spectral imaging

Coded apertures can provide an alternative, programmable means of pixel-level filtering. Coded aperture snapshot spectral imagers (CASSI) utilize combinations of coded apertures and one or more dispersive elements to modulate the optical field from a scene [1]. A 2D detector array captures a single multiplexed projection of the full 3D datacube. The nature of the multiplexing performed depends on the relative position of the coded aperture(s) and the dispersive element(s) within the instruments. The number of measurements made is significantly smaller than the number of voxels in the datacube that are eventually reconstructed. Diverse signal processing techniques, which will be discussed in this chapter and in chapter 3, may be used to reconstruct the datacube. Due to the modulation of the datacube using

a coded aperture, CASSI systems do not measure certain voxels. Thus, unlike the localized missing cone of information in the Fourier domain of a CTIS datacube, the missing information in the Fourier domain representation of the CASSI datacubes is distributed throughout the volume.

The dual disperser CASSI design [3] consists of two sequentially dispersive arms arranged in opposition so that the dispersion in the second arm cancels the dispersion introduced by the first arm. A coded aperture is placed between the two arms. Such a design applies spatially-varying, spectral filter functions with narrow features similar to Bayer pattern RGB sampling [1]. Through these filters, the detector measures a projective measurement of the datacube in the spectral domain. In essence, this design sacrifices spatial information to gain spectral information about the datacube. Spectral information from each spatial location in the scene is multiplexed over a localized region on the detector. A useful property of the design is that the measurement resembles the scene, making it easy to focus the camera on objects in the scene. This also makes it possible to perform local block processing of the detector data to generate smaller datacubes of subsets of the entire scene. Recovery of the datacube from the detector measurement may be performed using an expectation-maximization method combined with a wavelet denoising technique designed for spectral images.

In his thesis, Portnoy [29] describes a long-wave IR (LWIR) CASSI design using two coded aperture and two prisms that could selectively measure narrow spectral bands at periodic spatial locations. The features on each coded aperture and the dispersion introduced by each prism are designed so that each detector pixel measures exactly one spectral channel at each location in the scene. The spectral channels are not multiplexed on the detector array. Since the spectral channels from a given location in the scene are split across multiple pixels and measured individually, not every spatial location can be measured. Reconstruction of the full datacube requires

interpolation of the measured data in a process that parallels the reconstruction of RGB images from a sensor with a Bayer filter.

The examples of the dual disperser and LWIR CASSI demonstrate the flexibility provided by the use of coded apertures and other inexpensive optical components for designing novel snapshot spectral imaging architectures. The remainder of this chapter and chapters 3 and 4 focus on the design, characterization and analysis of another CASSI design.

2.4 CASSI - a single disperser design

The coded aperture spectrometer described earlier can be modified with the introduction of a front objective lens and by replacing the grating with a prism to produce a CASSI using just a single dispersive element. A basic schematic of the instrument is shown in figure 2.3. A standard imaging lens is used to form an image of a remote scene in the plane of the coded aperture. The coded aperture modulates the spatial information over all wavelengths in the datacube with the coded pattern. Imaging the datacube from this plane through the dispersive element results in multiple images of the code-modulated scene at wavelength-dependent locations on the detector array. The intensity pattern measured on the detector contains a coded mixture of spatial and spectral information about the scene. Figure 2.4 demonstrates the three step sensing process on the datacube.

Like the dual disperser, the single disperser does not directly measure each voxel in the desired 3D datacube. The instrument disperses spectral information from each spatial location in the scene over a certain area across the detector. Thus, spatial *and* spectral information from the scene is multiplexed on the detector pixels. Also, raw measurement of a scene on the detector rarely reveals spatial structure of the scene and makes block processing more challenging. The single disperser requires fewer optical elements than the dual disperser, which makes optical alignment of

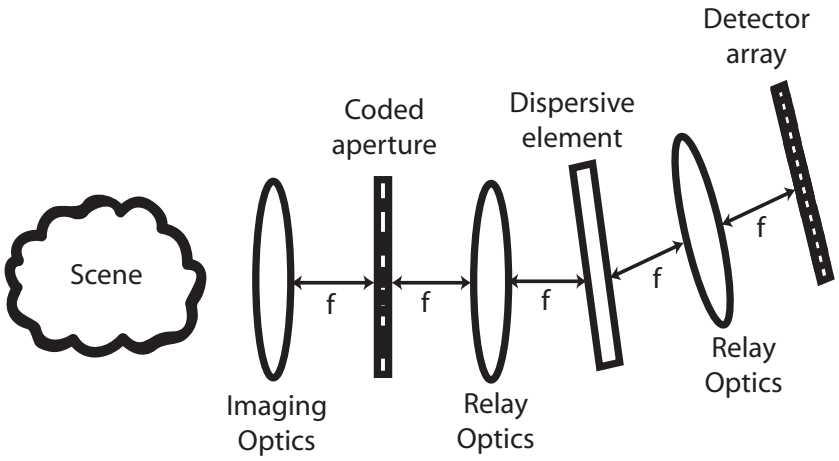

FIGURE 2.3: Schematic of a Coded Aperture Snapshot Spectral Imager (CASSI) using a single dispersive element. The imaging optics image the scene on to a 2D coded aperture. The relay optics relay the image from the plane of the coded aperture to a 2D detector through the dispersive element.

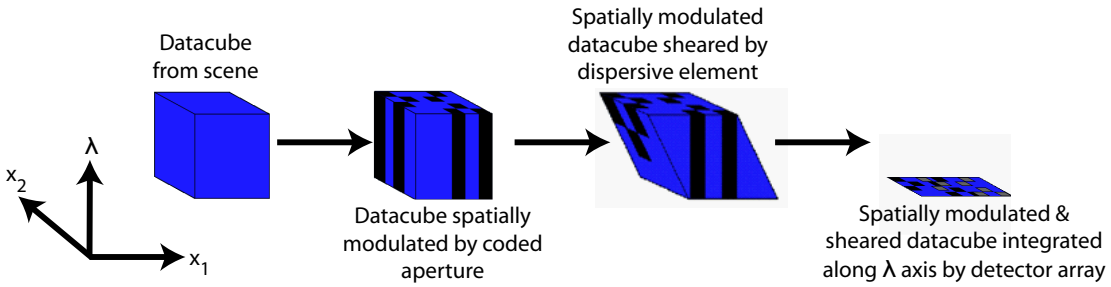

FIGURE 2.4: The CASSI three step sensing process on the datacube representing the spatio-spectral information in the scene. First, the coded aperture spatially modulates each spectral slice of the datacube. Then the dispersive element shears this spatially modulated datacube along one spatial axis. Finally, the spatially modulated and sheared datacube is integrated along the wavelength axis by a 2D detector array.

this design simpler.

The single disperser design also avoids the CTIS constraint of measuring multiple projections of the datacube and using a large focal plane array. Essentially, measuring just one spectrally-dispersed projection of the datacube that is spatially modulated by the aperture code over all wavelengths can be used to estimate the entire datacube.

The remainder of this chapter and chapter 3 will describe the design, charac-

terization and experimental results from single disperser CASSI prototypes. These instruments shall henceforth be referred to simply as CASSI.

2.5 CASSI system model

This section mathematically describes the CASSI sensing process. The reader may refer back to figures 2.3 and 2.4 to observe how the model describes the propagation of light through the instrument. The power spectral density of the image of the scene formed by the objective lens at the plane of the aperture code is denoted as $f_0(x'', y'', \lambda)$. The power spectral density immediately after the coded aperture is

$$f_1(x', y', \lambda) \;=\; T(x', y')\, f_0(x', y', \lambda), \tag{2.1}$$

where $T(x, y)$ is the transmission pattern printed on the coded aperture. The pattern is designed as an array of square features, with each feature being composed of 2×2 elements. Each element has the same size as that of a detector pixel, Δ. Let $t_{i,j}$ represent the binary value at the $(i, j)^{th}$ element, with a 1 representing a transmissive code element and a 0 representing an opaque code element. Then, $T(x, y)$ can be described as

$$T(x, y) = \sum_{i,j} t_{i,j}\, \tau(i, j; x, y), \tag{2.2}$$

with

$$\tau(i, j; x, y) = \text{rect}\left(\frac{x}{\Delta} - i, \frac{y}{\Delta} - j\right).$$

After further propagation through relay optics and the dispersive element, the power spectral density at the detector plane is

$$\begin{aligned} f_2(x, y, \lambda) \;&=\; \iint f_1(x', y', \lambda)\, h(x' - \gamma(\lambda) - x, y' - y, \lambda)\, dx'\, dy' \\ &=\; \iint T(x', y')\, f_0(x', y', \lambda)\, h(x' - \gamma(\lambda) - x, y' - y, \lambda)\, dx'\, dy', \end{aligned} \tag{2.3}$$

where $h(x'-x, y'-y, \lambda)$ assumes a shift-invariant optical impulse response of the relay optics and dispersive element. $\gamma(\lambda)$ describes the dispersion induced by the dispersive element. The detector array measures the intensity of incident light rather than the spectral density, which implies that the image on the detector array is the result of integrating the datacube along the wavelength axis over the spectral range Λ. In addition, the detector array is spatially pixelated by the pixel function

$$\mathrm{p}(m, n; x, y) = \mathrm{rect}\left(\frac{x}{\Delta} - m, \frac{y}{\Delta} - n\right),$$

with Δ denoting the side length of the pixel in each spatial dimension. As a result, the measurement at the $(m, n)^{th}$ pixel is

$$\begin{aligned}
g_{n,m} &= \int_\Lambda \iint f_2(x, y, \lambda)\, \mathrm{p}(m, n; x, y)\, dx\, dy\, d\lambda \\
&= \int_\Lambda \iiiint T(x', y')\, f_0(x', y', \lambda)\, h(x' - \gamma(\lambda) - x, y' - y, \lambda) \\
&\quad \times \mathrm{p}(m, n; x, y)\, dx'\, dy'\, dx\, dy\, d\lambda \\
&= \int_\Lambda \iiiint \sum_{i,j} t_{i,j}\, \mathrm{rect}\left(\frac{x'}{\Delta} - i, \frac{y'}{\Delta} - j\right) f_0(x', y', \lambda)\, h(x' - \gamma(\lambda) - x, y' - y, \lambda) \\
&\quad \times \mathrm{rect}\left(\frac{x}{\Delta} - m, \frac{y}{\Delta} - n\right) dx'\, dy'\, dx\, dy\, d\lambda \\
&= \sum_{i,j} t_{i,j}\, \Omega_{i,j,m,n}. \qquad (2.4)
\end{aligned}$$

Since CASSI images the aperture code onto the CCD along the y axis, $\Omega_{i,j,m,n} = \Omega_{i,n,m,n}\delta_{j,n}$, so that

$$\begin{aligned}
g_{n,m} &= \sum_{i,j} t_{i,j}\, \Omega_{i,j,m,n} \\
&= \sum_{i} t_{i,n}\, \Omega_{i,n,m}. \qquad (2.5)
\end{aligned}$$

Here

$$\begin{aligned}\Omega_{i,n,m} &= \int_\Lambda \iiiint \operatorname{rect}\left(\frac{x'}{\Delta} - i, \frac{y'}{\Delta} - n\right) \operatorname{rect}\left(\frac{x}{\Delta} - m, \frac{y}{\Delta} - n\right) \\ &\quad \times f_0(x', y', \lambda)\, h(x' - \gamma(\lambda) - x, y' - y, \lambda)\; dx'\, dy'\, dx\, dy\, d\lambda. \end{aligned} \tag{2.6}$$

Letting $x'' = x' - i\Delta$ and $y'' = y' - n\Delta$,

$$\begin{aligned}\Omega_{i,n,m} &= \int_\Lambda \iiiint \operatorname{rect}\left(\frac{x''}{\Delta}, \frac{y''}{\Delta}\right) \operatorname{rect}\left(\frac{x}{\Delta} - m, \frac{y}{\Delta} - n\right) \\ &\quad \times f_0(x'' + i\Delta, y'' + n\Delta, \lambda)\, h(x'' + i\Delta - \gamma(\lambda) - x, y'' + n\Delta - y, \lambda) \\ &\quad \times dx''\, dy''\, dx\, dy\, d\lambda. \end{aligned} \tag{2.7}$$

Letting $x''' = x - m\Delta$ and $y''' = y - n\Delta$,

$$\begin{aligned}\Omega_{i,n,m} &= \int_\Lambda \iiiint \operatorname{rect}\left(\frac{x''}{\Delta}, \frac{y''}{\Delta}\right) \operatorname{rect}\left(\frac{x'''}{\Delta}, \frac{y'''}{\Delta}\right) \\ &\quad \times f_0(x'' + i\Delta, y'' + n\Delta, \lambda)\, h(x'' + i\Delta - \gamma(\lambda) - x''' - m\Delta, y'' - y''', \lambda) \\ &\quad \times dx''\, dy''\, dx'''\, dy'''\, d\lambda. \end{aligned} \tag{2.8}$$

Assuming that the dispersion by the dispersive element is linear over the spectral range of the system so that $\gamma(\lambda) = \alpha\lambda$, and letting $\lambda' = \lambda + \left(\frac{m\Delta - i\Delta}{\alpha}\right)$,

$$\begin{aligned}\Omega_{i,n,m} &= \int_\Lambda \iiiint dx''dy''\, dx'''dy'''\, d\lambda' \operatorname{rect}\left(\frac{x''}{\Delta}, \frac{y''}{\Delta}\right) \operatorname{rect}\left(\frac{x'''}{\Delta}, \frac{y'''}{\Delta}\right) \\ &\quad \times f_0\left(x'' + i\Delta, y'' + n\Delta, \lambda' - \left(\frac{m\Delta - i\Delta}{\alpha}\right)\right) \\ &\quad \times h\left(x'' - x''' - \alpha\lambda', y'' - y''', \lambda' - \left(\frac{m\Delta - i\Delta}{\alpha}\right)\right). \end{aligned} \tag{2.9}$$

Defining a discrete version of the datacube as

$$\begin{aligned} f_{i,n,m} &= \int_\Lambda \iiint \text{rect}\left(\frac{x''}{\Delta}, \frac{y''}{\Delta}\right) \text{rect}\left(\frac{x'''}{\Delta}, \frac{y'''}{\Delta}\right) \\ &\quad \times f_0\left(x'' + i\Delta, y'' + n\Delta, \lambda' - \frac{m\Delta}{\alpha}\right) \\ &\quad \times h\left(x'' - x''' - \alpha\lambda', y'' - y''', \lambda' - \frac{m\Delta}{\alpha}\right) \\ &\quad \times dx''\, dy''\, dx'''\, dy'''\, d\lambda', \end{aligned} \tag{2.10}$$

so that $f_{i,n,m-i} = \Omega_{i,n,m}$. As a result, the measurement at the $(m,n)^{th}$ pixel is

$$\begin{aligned} g_{n,m} &= \sum_i t_{i,n}\, f_{i,n,m-i} \\ &= \sum_k t_{m-k,n}\, f_{m-k,n,k}. \end{aligned} \tag{2.11}$$

This equation provides a simple interpretation of the CASSI measurement of the discrete datacube. First, each spectral channel of the discrete 3D datacube is spatially modulated by an element-wise multiplication with the aperture code. This spatially modulated datacube is then sheared by passing through the dispersive element, which results in a shift in the first index. Finally, this spatially modulated and sheared datacube is integrated over the wavelength dimension to produce the 2D array of measurements on the detector. This measurement can also be interpreted as an operation on a sheared datacube $f'_{m,n,k} = f_{m-k,n,k}$ so that

$$g_{n,m} = \sum_k t_{m-k,n}\, f'_{m,n,k}.$$

If the 2D array of detector measurements is represented as a vector g, this expression can also be written as the matrix-vector product

$$g = \Phi f. \tag{2.12}$$

The matrix Φ in equation (2.12) is the sensing matrix that represents CASSI's operation on the discretized datacube. It is a non-negative matrix that maps voxels of the 3D sampled and sheared datacube to pixels of the detector array. As the number of pixels on the detector used for the measurement is smaller than the number of voxels in the discrete datacube, this measurement equation is under-determined. Numerical estimation methods must use the detector measurements, g, and a model of the sensing matrix, Φ, to recover an estimate of the discrete datacube, f.

Equation (2.10) can be interpreted as the object datacube imaged on to the aperture code, $f_0(x, y, \lambda)$, being filtered by the CASSI system transfer function, which ultimately determines the spatial and spectral resolution of the datacube estimate. The spatial resolution of the reconstructed datacube depends on (i) the point spread function (PSF), $h(x', x, y', y, \lambda)$, of the relay optics and dispersive element, (ii) the pixel size, Δ, (iii) the size of a feature on the coded aperture, and (iv) numerical estimation effects. However, ignoring numerical estimation effects, the spatial resolution is approximately given by the width and height of the smallest feature on the coded aperture. Using larger features on the coded aperture will produce larger transmissive areas, but will also rely more heavily on numerical estimation to estimate the spatial content of objects in the datacube that are being imaged on to opaque areas of the coded aperture.

The spectral resolution of the reconstructed datacube is the separation between spectral channels in the reconstructed datacube (in nm). The spectral resolution of the reconstructed datacube depends on (i) the amount of dispersion induced by the dispersive element, (ii) the PSF, (iii) the pixel size, (iv) the size of the smallest feature on the coded aperture, and (v) numerical estimation effects. The spectral resolution will be limited by the smallest feature size on the coded aperture. If the smallest code feature is assumed to be 2 detector pixels wide and the dispersion is in the horizontal direction, then imaging two adjacent monochromatic point sources of

close but distinct wavelengths on to the smallest code feature can potentially result in the spatio-spectral mapping of both point sources to the same pixel on the CCD. Thus, the spectral resolution of the system is determined by the amount of dispersion (in nm) across the 2 detector pixels. Furthermore, if the dispersion is non-linear, the spectral resolution of CASSI varies as a function of wavelength.

2.6 CASSI prototype 1 system description

To experimentally verify the CASSI spectral imaging concept, a proof-of-concept prototype as shown in figure 2.5 was constructed. The prototype consisted of (i) a coded aperture that is lithographically patterned as a chrome coating on a quartz mask, (ii) three lenses from Schneider Optics Inc. with an $F/\#$ of 1.4 and a focal length of 22.5 mm, (iii) an equilateral prism from Edmund Optics Inc. as a dispersive element, (iv) a CCD detector array from Photometrics with 1040×1392 pixels that are 4.65 μm square each, and (v) a $500 - 620$ nm bandpass filter that was placed in front of the imaging lens to remove the impact of stray light on the experimental measurements. Matlab routines were written to control and capture data on the CCD.

The aperture code used in all the experiments was based on an order 192 S-matrix code [30], with features that were four CCD pixels wide and four CCD pixels tall, and two completely closed rows of CCD pixels added between the code rows. The columns of the original S-matrix code were shuffled in a random but repeatable way. The code was originally designed for use with a coded aperture spectrometer [14, 15] and was not optimized for CASSI.

An equilateral prism was used instead of a grating because the grating produces overlapping diffractive orders while the prism only refracts the wavelengths into one order. Prisms also have large transmission efficiencies [31]. Given the system geometry and the low dispersion of the equilateral prism, the number of CCD columns

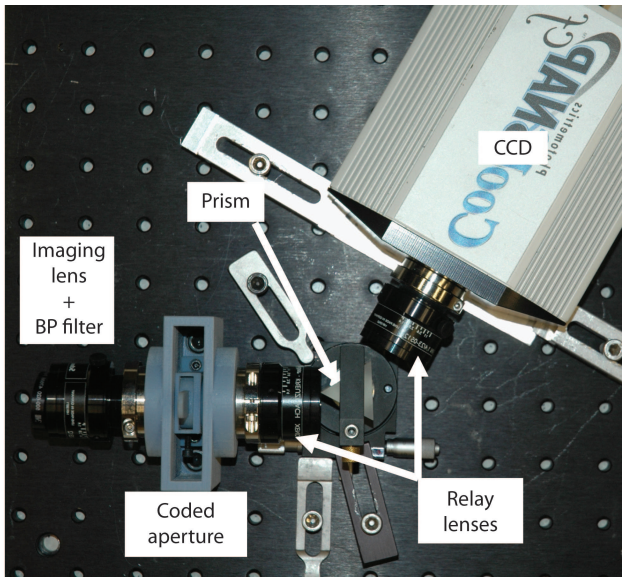

FIGURE 2.5: Experimental CASSI prototype. The prototype consists of (i) a coded aperture, lithographically patterned as a chrome coating on a quartz mask, (ii) three lenses from Schneider Optics Inc. with an $F/\#$ of 1.4 and a focal length of 22.5 mm, (iii) an equilateral prism from Edmund Optics Inc. as a dispersive element, (iv) a CCD detector array from Photometrics with 1040×1392 pixels that are 4.65×4.65 $(\mu\text{m})^2$ each, and (v) a $500 - 620$ nm bandpass filter that is placed in front of the imaging lens to remove the impact of stray light on the experimental measurements.

illuminated when white light was allowed to pass through the system was less than half the width of the CCD array.

2.7 GPSR numerical estimation method for CASSI

As mentioned previously, CASSI measures a 2D, spatio-spectral multiplexed projection of the 3D datacube representing the scene. Equation (2.12) is used to mathematically describe this *compressed* measurement process and it must be solved (inverted) to produce an estimate of the datacube, f. Since the number of detector pixels in the measurement g is significantly smaller than the number of voxels in the datacube, the measurement equation is under-determined. Thus, any numerical estimation method used to solve this equation to reconstruct the datacube must exploit properties or structures in the datacube known a priori that help to constrain the estimated

datacube to a certain solution space.

This section describes the first of three numerical estimation methods used to reconstruct a datacube from a snapshot CASSI measurement. It relies on the assumption that objects in the scene have piecewise smooth spatial structure, making the datacube compressible in the wavelet basis. Thus, it is very difficult to reconstruct a datacube of a scene that does not have any spatial structure, for example, a scene consisting of a random, dense, spatial distribution of point sources.

The datacube and its wavelet decomposition are related through the mathematical expression

$$f = \Psi b, \tag{2.13}$$

where b is a vector composed of the 2D wavelet transform coefficients for each spectral band concatenated to form one vector, and Ψ denotes the inverse 2D redundant wavelet transform applied to each spectral band to form the datacube f. Consequently, the CASSI detector measurement as described by equation (2.12) can be represented as

$$g = \Phi\Psi b + n, \tag{2.14}$$

where n is additive noise on the detector array.

If the datacube, f, consists of $\{N \times N\}$ spatial channels with N_λ spectral channels, it can be represented as a discrete cube with $\{N \times N \times N_\lambda\}$ elements. The corresponding detector measurements, g, can be represented as a matrix of size $\{N \times (N + N_\lambda - 1)\}$. The number of columns in this matrix reflects the fact that the detector measurement is a sum of coded images of the scene at each spectral channel, with each spectral image displaced by a column of pixels from the adjacent image. If f and g are represented as column vectors, the sensing matrix, Φ, can be represented as a matrix of size $\{[N\,(N + N_\lambda - 1)] \times (N^2\,N_\lambda)\}$. The vector of wavelet coefficients of the datacube, b, is of size $\{[N^2 N_\lambda\,(3\log_2(N) + 1)] \times 1\}$. The size of this

vector reflects the fact that the wavelet decomposition of the datacube is performed as a 2D undecimated (redundant) transform on each of the q spectral channels. The undecimated transform is used to ensure that the resulting method is translation invariant.

An estimate, $\hat{f}$, for the datacube can be found by minimizing this objective function

$$\hat{f}(\tau) = \Psi \left[\underset{b}{\operatorname{argmin}} \left\{ \frac{1}{2} ||g - \Phi\Psi b||_2^2 + \tau ||b||_1 \right\} \right]. \tag{2.15}$$

The solution to this convex unconstrained optimization problem is found by using a Matlab implementation of the iterative Gradient Projection for Sparse Reconstruction (GPSR) method developed by Figueiredo et al. [32], with Matlab code available online at http://www.lx.it.pt/~mtf/GPSR/. Inputs to the method are the vector of CASSI detector measurements, g, and a model of the sensing matrix, Φ, either as a matrix or as a Matlab function. A functional form is required when the number of detector measurements and the number of voxels in the datacube to be estimated are large enough to make it impossible to store the matrix Φ in memory. The quality of the model of the sensing matrix depends on how well it is experimentally calibrated and determines the quality of the estimated datacube.

The GPSR method searches for a datacube with a sparse representation in the wavelet basis; i.e. a b which contains mostly zeros and a relatively small number of large coefficients. The first term in equation (2.15) minimizes the ℓ_2 error between the measurements modeled from the estimate and the true measurement. The second term is a penalty term that encourages sparsity of the reconstruction in the wavelet domain and controls the extent to which spatially piecewise smooth estimates are favored. In this formulation, τ is a tuning parameter for the penalty term and higher values of τ yield sparser estimates of b. The objective function in equation (2.15) is known to be convex, thus having a unique minimum for a given choice of τ.

Empirically, while the use of redundant wavelet transform operations result in better datacube estimates than non-redundant wavelet transform operations, GPSR iterations with these operations involve large matrices and are significantly slower. Furthermore, the use of wavelet transforms means that the algorithm must estimate spatially square datacubes with dimensions that are powers of 2. This means that the GPSR algorithm does not take advantage of the separable nature of the CASSI system model along the vertical dimension, which can potentially allow the datacube estimation problem to be broken down into sub-problems and parallelized.

2.7.1 Simulation results

To test the CASSI sensing concept and the GPSR numerical estimation method, a Matlab simulation was conducted using a phantom datacube. For the purpose of the simulation, an approach that would result in easy visualization of the results of the reconstruction algorithm was desired. A datacube was generated by converting a 256×256 RGB color image of peppers, as shown in figure 2.6(a), to a synthesized datacube with a 15 channel spectrum defined for each pixel in the image.

The spectrum at each pixel corresponded to a weighted sum of RGB filter functions. Using the synthesized datacube, the detector measurement was simulated by passing the datacube through a simulated system forward model. Optical blurring, alignment and assembly issues that are typically encountered in the experimental measurement process were ignored. Figure 2.6(b) shows the random aperture code used in the measurement process. The feature sizes on the aperture code were assumed to be the same as the size of the detector pixels.

Using the aperture pattern and the simulated detector measurement shown in figure 2.6(c), the GPSR method was used to produce an estimate of the datacube. The resulting estimate was converted to a 2D image of RGB vectors using the RGB filter functions. The RGB image generated from the estimated datacube is shown in

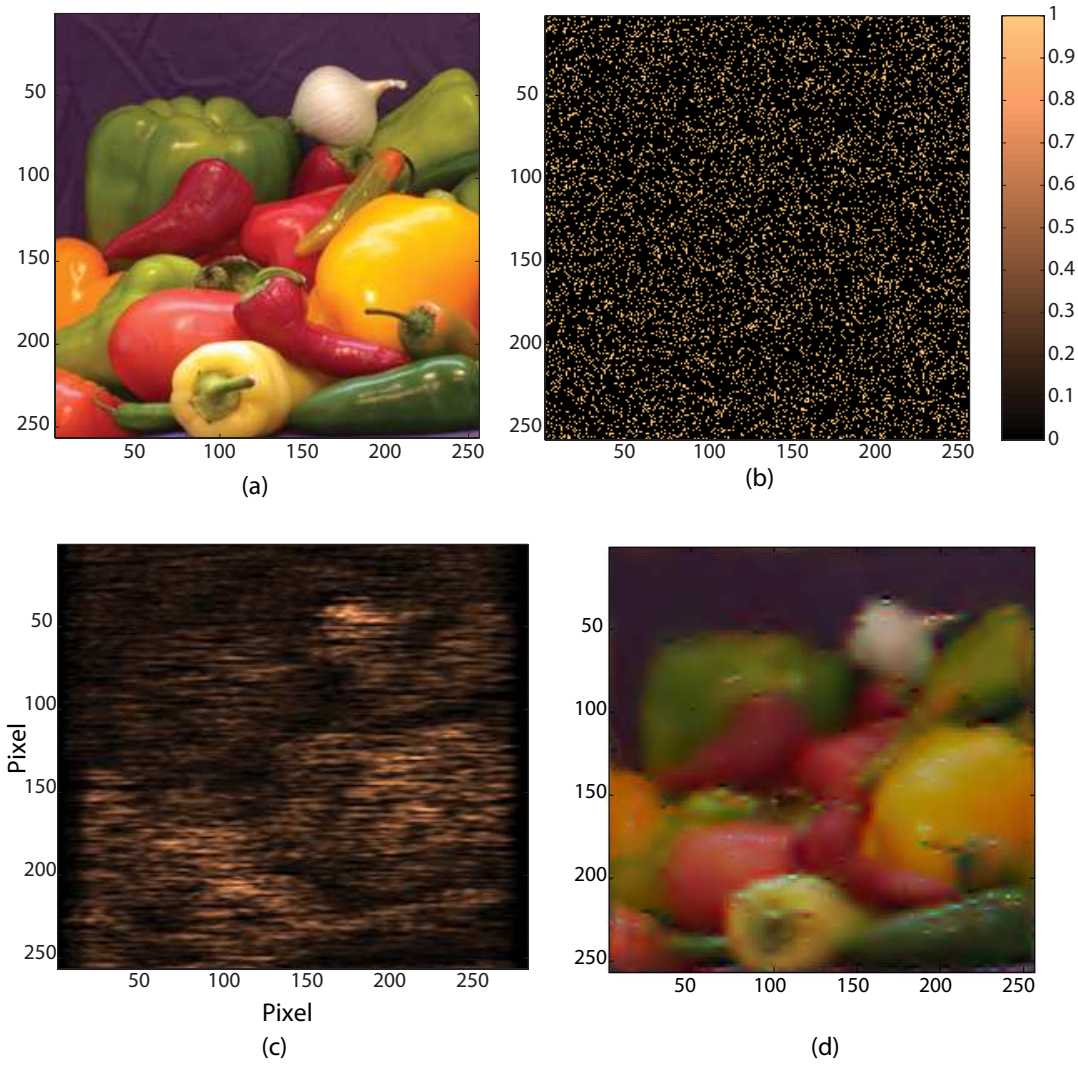

FIGURE 2.6: Testing the CASSI concept using through simulation using the GPSR method. (a) 256×256 RGB image used to generate the datacube. (b) 256×256 code pattern used for reconstruction of the simulated spectral datacube. The feature size was the same as the size of the detector pixels. (c) 256×270 simulated detector measurement of the spectral datacube. (d) 256×256 RGB image generated from the estimated spectral datacube.

figure 2.6(d). Although there are some spatial artifacts present, this simulation provided the first demonstration that the new CASSI design could potentially generate datacube estimates of natural scenes from snapshot CASSI detector measurements.

2.8 Spectral imaging of a static scene using GPSR

The first CASSI prototype described above was put together with off-the-shelf parts and not optimized in optical ray tracing software. Consequently, the results presented in this section are limited to capturing a snapshot spectral image of a relatively simple scene.

To generate an estimate of the datacube representing the spectral image of the scene, the GPSR method was provided with two inputs - (i) a snapshot CASSI detector measurement of the scene and, (ii) an image of the aperture code for its use in constructing a Matlab functional form of the sensing matrix, Φ. Instead of using an image of the ideal code pattern that was lithographically printed on the aperture code substrate, an image of the aperture code was captured on the detector after uniformly illuminating it with a 543 nm laser. Figure 2.7 shows the detector measurement of the aperture code pattern after propagation through the optics. It is important to note that while this approach attempts to capture some of the optical distortions in the system, it is only a first order, coarse approach to building a calibrated sensing matrix. A more detailed discussion of CASSI calibration is presented in chapter 3.

A scene consisting of two ping pong balls was constructed, as shown in figure 2.8. One ping pong ball was illuminated with a 543 nm laser and a white light source filtered by a green 560 nm narrow band filter. The other ball was painted red and illuminated with a white light source.

Figure 2.9 shows the CASSI detector measurement of the scene. Given the low linear dispersion of the equilateral prism, there is spatio-spectral overlap of the aperture code-modulated images of each ball. The $500-620$ nm bandpass filter placed in front of the imaging lens ensures that only the subset of the datacube corresponding to this band of wavelengths is measured by the instrument. The spatial features in

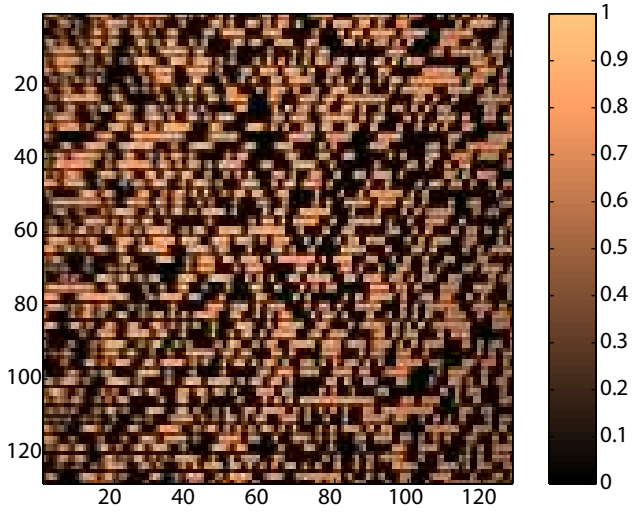

FIGURE 2.7: Image of the aperture code recorded on the CCD upon illumination of the entrance aperture with 543 nm light. The recorded pattern is used to build a functional form of the sensing matrix for the GPSR method to generate an estimate of the datacube.

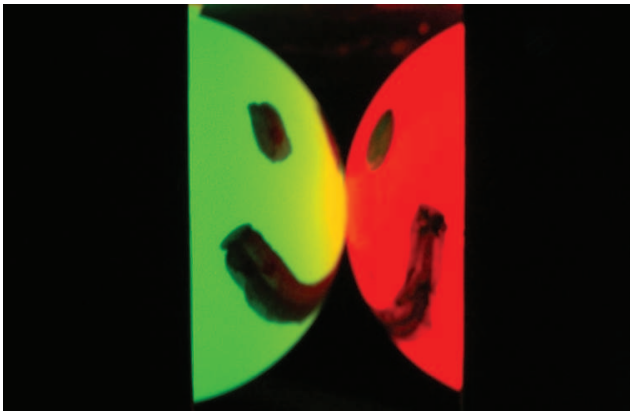

FIGURE 2.8: A scene consisting of a ping pong ball illuminated with a 543 nm green laser and a white light source filtered by a 560 nm narrow band filter (left), and a red ping pong ball illuminated with a white light source (right).

the scene include the spatially uniform bodies of each ball and the piecewise smooth features (eyes and smile) on each ball. The objects are also spectrally sparse, as illustrated by the fact that the spatial modulation of the aperture code on each ball in figure 2.9 is mostly visible.

Detector measurements of both the code pattern and the scene were digitally

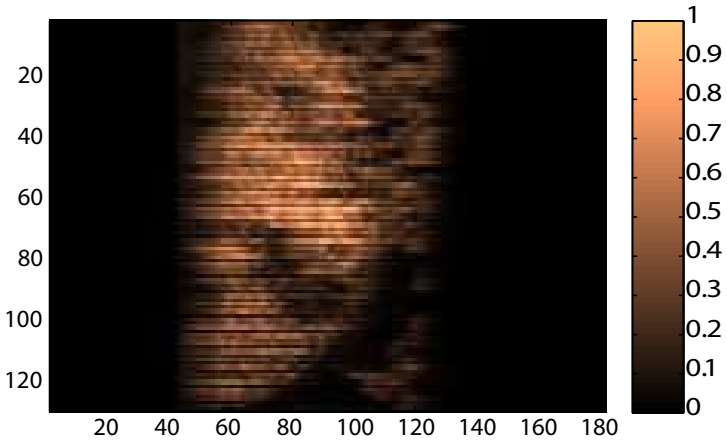

FIGURE 2.9: Detector measurement of the scene consisting of the two ping pong balls. Given the low linear dispersion of the prism, there is spatio-spectral overlap of the aperture code-modulated images of each ball.

downsampled by a factor of two in the row and column dimensions prior to performing a reconstruction of the datacube. This downsampling helped reduce the time needed by the reconstruction algorithm to generate an estimate of the datacube. The resulting $\{128 \times 128 \times 28\}$ datacube spanned a spectral range of 540 nm to 640 nm. 100 GPSR iterations required about 14 minutes of runtime on a desktop machine. The GPSR method was initialized with $\tau = 0.05$, a value that was determined via trial and error.

Figure 2.10 shows the spatial content of each of 28 wavelength channels between 540 nm to 640 nm. Note that the aperture code modulation on the spatial structure visible in figure 2.9 has been removed in all the wavelength channels and that the two balls are spatially separated. To validate the CASSI datacube estimate, point-wise spectra obtained using a non-imaging spectrometer (Ocean Optics USB2000) were used as a quantitative reference. Figure 2.11(a) shows the CASSI spectrum at a point on the green ping pong ball, while figure 2.11(b) shows the CASSI spectrum at a point on the red ping pong ball. The wavelength axis in the plots had to be calibrated due to the non-linear dispersion of the prism across the detector. This

calibration was performed by tracking the position of a point source while varying its wavelength. Figures 2.11(a) and (b) also show the reference spectra of each ball for comparison. The two reconstructed CASSI spectra closely match those generated by the reference spectrometer.

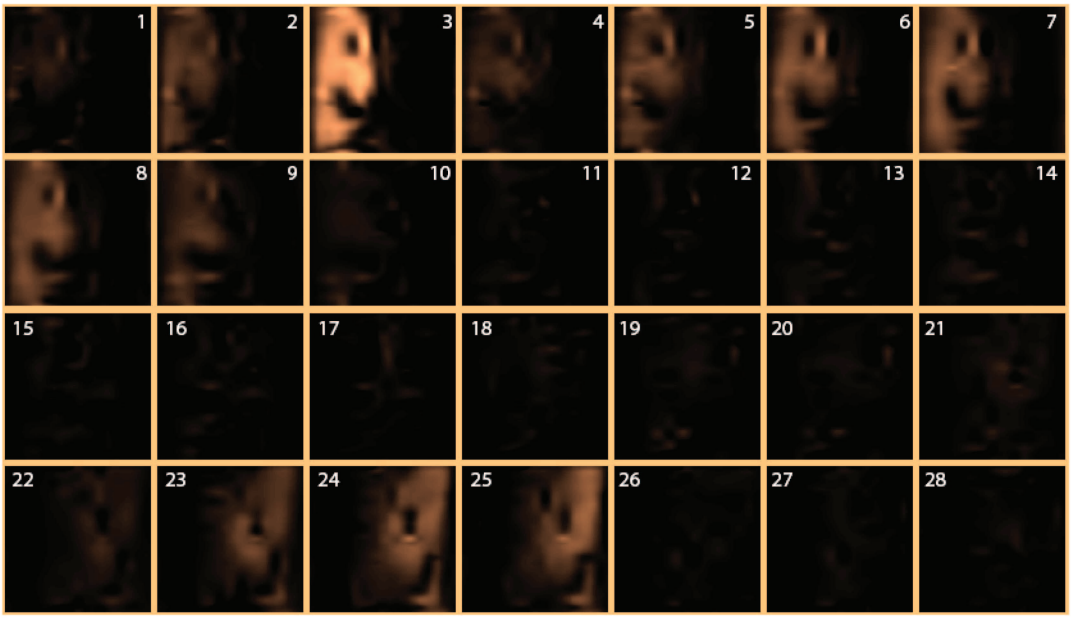

FIGURE 2.10: Spatial content of the scene in each of 28 spectral channels between 540 nm and 640 nm. The green ball can be seen in channels 3, 4, 5, 6, 7 and 8, while the red ball can be seen in channels 23, 24 and 25.

As mentioned earlier, the functional form of the sensing matrix used an image of the aperture code when the system was illuminated by a laser at 543 nm. The detector response of the imaging system at all other spectral channels was assumed to be identical and as such did not account for a wavelength dependent, anamorphic horizontal stretch of the image. The quality of the reconstructed datacube could certainly be improved if the sensing matrix captured some of the wavelength dependent optical distortions of the system. Such an approach is discussed further in chapter 3.

An important characteristic of any spectrometer or spectral imager is its spectral

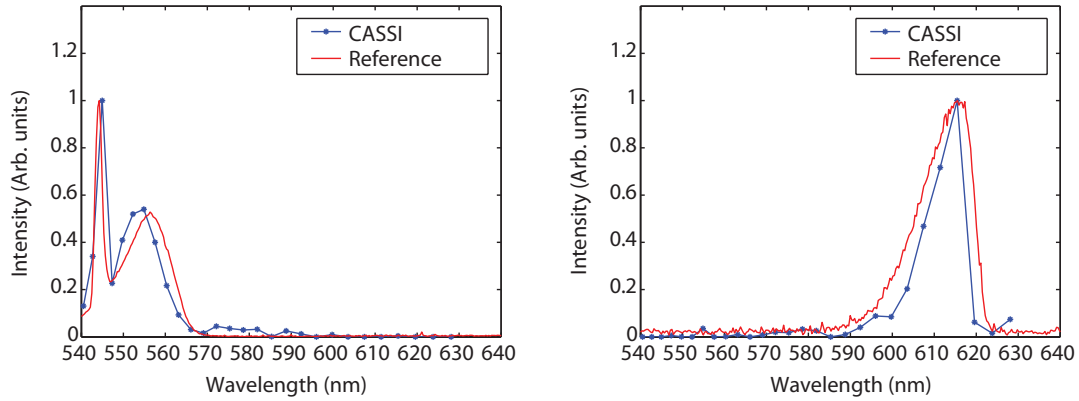

FIGURE 2.11: (a) Spectral intensity through a point on the ping pong ball illuminated with a 543 nm green laser and a white light source filtered by a 560 nm narrow band filter. (b) Spectral intensity through a point on the red ping pong ball illuminated with a white light source. Spectra from an Ocean Optics non-imaging reference spectrometer are shown for comparison.

resolution. Ignoring optical distortions such as the blurring and the smile distortion, the spectral resolution of CASSI is limited by the width of the smallest code feature, which was 4 detector pixels. The spectral resolution of the CASSI experimental prototype was determined by removing the bandpass filter in front of the imaging lens and illuminating the instrument with a 543 nm and 632 nm laser. The separation between images of the aperture code resulting from these wavelengths was approximately 100 pixels, corresponding to an average dispersion of 0.9 nm per pixel. Since the width of the smallest code feature was 4 detector pixels, the spectral resolution of the instrument was limited to 3.6 nm/spectral channel. This number represents the average spectral resolution since the dispersion of a prism is non-linear, which means that the spectral resolution varies with wavelength.

2.9 Conclusions

This chapter introduced spectral imaging and reviewed diverse spectral imager designs. A new CASSI design was introduced that uses a coded aperture, a single

dispersive element and a CCD array to measure a 2D spatio-spectral projection of a 3D discrete datacube representing the scene. The GPSR method was used to numerically estimate the datacube by assuming spatial sparsity of the scene in a wavelet basis.

A first generation prototype was used to provide experimental validation of the compressive spectral imaging concept. While the spatial resolution was only modest, it was enough to allow the identification of two distinct objects in the reconstructed scene. The spectral signatures of the scene as measured by the experimental prototype compared well with a reference non-imaging spectrometer. The next chapter will focus on the design of a custom-designed CASSI prototype with better image quality, more rigorous system calibration and testing alternative methods for numerically estimating datacubes from snapshot CASSI measurements.

3

Spectral imaging of a dynamic scene using CASSI

The first generation CASSI prototype discussed in the previous chapter had limitations, both in terms of its hardware as well as the software processing used to numerically estimate a datacube from a snapshot detector measurement. Optically, the instrument had a limited field of view because of severe vignetting at the edges of the field. Only a small region of the full aperture code could be imaged onto the detector without severe aberrations. There was also a severe wavelength dependent anamorphic distortion of the system response. Algorithmically, the GPSR numerical estimation method was computationally intensive.

With these limitations in mind, a second generation CASSI prototype was conceived and developed. The focus of this chapter is the design of this new custom-designed prototype and the use of its snapshot capability for spectral imaging of a dynamically changing scene at video rate. Section 3.1 describes the optical design, including the initial Zemax modeling of the prototype. Zemax is an optical ray tracing software that allows optimization of the shape and locations of optical surfaces based on the desired characteristics of the optical design. Section 3.2 describes the process of calibrating CASSI to provide a datacube estimation algorithm with

a system-specific model that accounts for additional factors that were absent in the mathematical model discussed in section 2.5 and were ignored in the previous chapter. Section 3.3 describes how the Two-step Iterative Shrinkage and Thresholding algorithm can be used to reconstruct a datacube from a CASSI measurement and demonstrates its use in recovering a spectral image of a static scene. Section 3.4 describes the Nested Adaptive Refinement Estimation method, which effectively utilizes the data gathered during the CASSI calibration process to build system-specific sensing matrices at multiple, nested scales and solves all of them to estimate the datacube. Section 3.5 describes the use of this algorithm to reconstruct a spectral image of the static scene, as well as a spectral image video of a dynamically changing scene.

3.1 CASSI prototype 2

The first subsection in this section describes all of the components used to construct the second generation CASSI prototype. The second subsection describes the optical design work conducted in Zemax to test the feasibility of the design.

3.1.1 System description

The second generation CASSI prototype was designed to operate over the visible spectrum between 450 and 650 nm. The final design, as shown in figure 3.1(a), consisted of (i) an objective lens (Computar $F/1.4$, 16 mm lens), (ii) a coded aperture, (iii) a bandpass filter (Omega Optical), (iv) an F/8 relay lens (Edmund Optics stock part C45762), (v) a double Amici prism as a dispersive element, and (vi) a 2D CCD detector array. The CCD detector (AVT Marlin, Allied Vision Technologies) was an 8-bit camera with 9.9×9.9 $(\mu\text{m})^2$ pixels and had its strongest response at 500 nm, with a relative response greater than 0.7 between $450 - 650$ nm.

Spectral dispersion was introduced using a custom designed double Amici prism

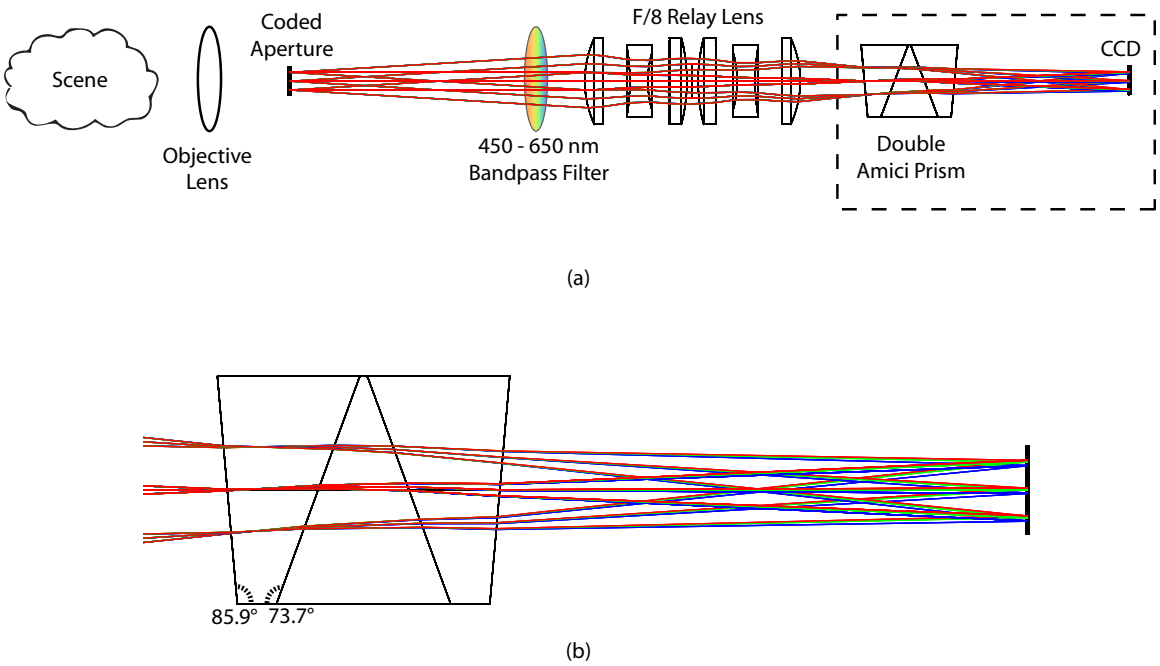

FIGURE 3.1: (a) Schematic of the direct view CASSI prototype. The prototype consists of (i) an objective lens (Computar $F/1.4$, 16 mm lens), (ii) a coded aperture, (iii) a bandpass filter, (iv) an F/8 relay lens (Edmund Optics part C45762), (v) a double Amici prism as a dispersive element, and (vi) a 2D CCD detector array (AVT Marlin with 9.9×9.9 $(\mu\text{m})^2$ pixels). (b) Ray bundles at different wavelengths, illustrated by the red, green and blue lines, are coincident when they first hit the double Amici prism but dispersed to different points on the CCD after passing through the prism.

that was made up of three prisms cemented together. Repeated use of Snell's law at each prism interface was used to design the double Amici prism to ensure that rays corresponding to the center wavelength would pass through the prism undeviated and this process is briefly reviewed here. Referring to figure 3.2, consider a monochromatic ray of light that is parallel to the optical axis and incident on the front face of the double Amici prism at an angle $\theta_1 = \phi_1$. The index of refraction of air $n_1 = 1$. The indices of refraction of each glass, n_2 and n_3, are well known functions of wavelength. After entering the first prism, the ray gets refracted at an

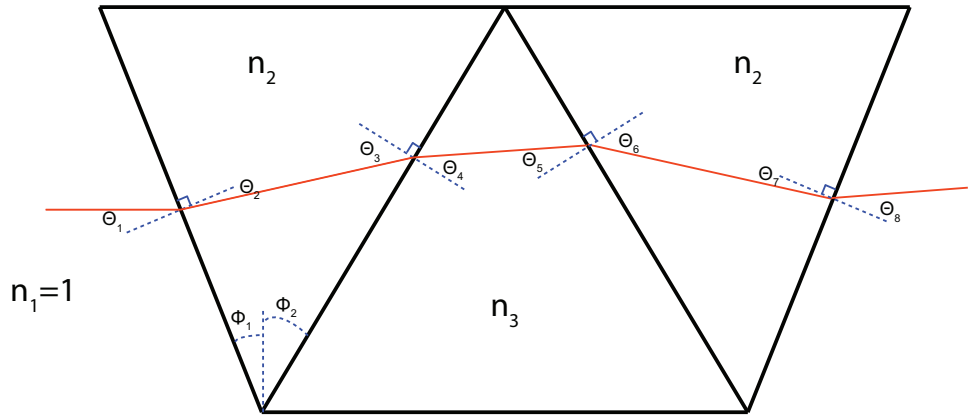

FIGURE 3.2: Diagram showing relevant variables when designing the double Amici prism. Upon picking the desired glass types with indices of refraction n_1 and n_2 to provide sufficient dispersion, angles ϕ_1 and ϕ_2 must be selected so that rays corresponding to the center wavelength pass through the prism undeviated.

angle θ_2, which is determined by Snell's Law

$$\theta_2 = sin^{-1}\left(\frac{n_1}{n_2} sin(\theta_1)\right).$$

After traveling through the first prism, the ray is incident on the second prism at an angle

$$\theta_3 = 90^o - (180^o - (\phi_1 + \phi_2) - (90 - \theta_2)).$$

At the interface between the first and second prism, the ray is refracted at an angle θ_4. This angle of refraction is determined by Snell's Law and is given by

$$\theta_4 = sin^{-1}\left(\frac{n_2}{n_3} sin(\theta_3)\right).$$

After traveling through the second prism, the ray is incident on the third prism at an angle

$$\theta_5 = 90^o - (180^o - (180^o - (90^o - \phi_2) - (90^o - \phi_2)) - (90^o - \theta_4)).$$

At the interface between the second and third prism, the ray is refracted at an angle θ_6. This angle of refraction is determined by Snell's Law and is given by

$$\theta_6 = sin^{-1}\left(\frac{n_3}{n_2} sin(\theta_5)\right).$$

After traveling through the third prism, the ray is incident on the last surface of the prism at an angle

$$\theta_7 = 90^o - (180^o - (\phi_1 + \phi_2) - (90^o - \theta_6)).$$

At the interface between the third prism and the air, the ray is refracted at an angle θ_8 determined by Snell's Law and is given by

$$\theta_8 = sin^{-1}\left(\frac{n_2}{n_1} sin(\theta_7)\right).$$

Upon picking the desired glass types, ϕ_1 and ϕ_2 must be selected so that rays corresponding to the center wavelength pass through the prism undeviated. The two angles are adjusted until the following conditions are satisfied

$$\theta_4 = \phi_2 \tag{3.1a}$$

$$\theta_8 = \theta_1. \tag{3.1b}$$

At the end of this analysis, SK2, a medium dispersion crown glass (low-medium index of refraction (n_d) and Abbe number (v_d) on the glass map as shown in figure 3.3), was selected as the glass type of the the first and third prisms of the double Amici prism. The second prism was made of SF4, a higher dispersion flint glass (high n_d and v_d on the glass map). Based on the indices of refraction of each prism and the conditions specified in equation (3.1), the interior angles of these prisms (see figure 3.1(b)) were chosen to ensure that rays corresponding to the center wavelength of 550 nm would pass through the prism undeviated. These angles were refined after optimization of the overall CASSI design in Zemax. Figure 3.1(b) shows a polychromatic, on-axis ray bundle passing through the custom-designed double Amici prism. Rays corresponding to adjacent wavelengths dispersed to either side of the optical axis, resulting in a direct view prism. This configuration was useful because all the components of the system could be placed in a line, which made system

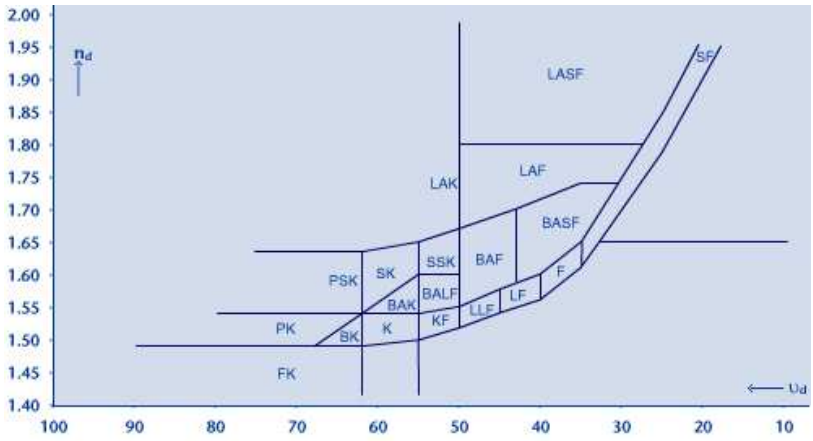

FIGURE 3.3: Schott glass map obtained from www.us.schott.com. The first and third prisms in the double Amici prism are made of a medium dispersion crown glass while the second prism is made of a high dispersion flint glass.

alignment much easier than that of the previous prototype discussed in chapter 2. Wavelength dependent anamorphic distortion was minimized because the path lengths of all wavelengths from the prism to the CCD were similar. Finally, the total amount of dispersion created by the double Amici prism could be adjusted by translating the prism between the relay lens and the detector array along the optical axis. This translation induced negligible change in the wavelength dependent PSF at the detector array. The prism assembly was fabricated by Shanghai Optics.

A relay lens (Edmund Optics stock part C45762) relayed the image from the plane of the coded aperture to the CCD. The bandpass filter (Omega Optical) limited the spectral range of the system to $450 - 650$ nm. The coded aperture was lithographically patterned as a chrome coating on a quartz substrate, with an anti-reflective coating on both sides designed for the $400 - 700$ nm range. The code used in all the experiments described in this chapter was a random 256×248 element binary pattern as shown in figure 3.4, with the smallest code feature being 2×2 CCD pixels (19.8 μm on each side). This limited the spatial size of the reconstructed datacube to a maximum of 256×248 spatial elements. The objective lens (Computar 16 mm

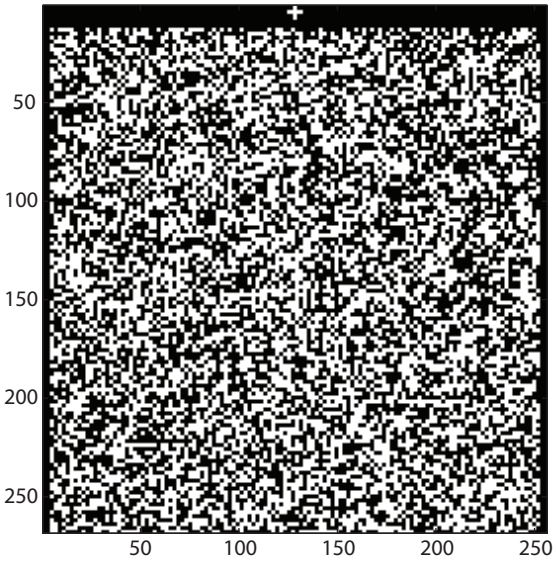

FIGURE 3.4: The aperture code used in the second CASSI prototype was a random 256×248 binary element pattern, with the smallest feature being 2×2 CCD pixels ($19.8\ \mu m$ on each side). The cross on the top of the code is used for system alignment and system calibration.

lens with an F/# of 1.4) imaged the scene on to the plane of the coded aperture. All of these system components were mounted on a metal rod/cage system, which served as the backbone of the system in front of the CCD. The coded aperture and the double Amici prism were placed inside rotation mounts that could be adjusted to ensure proper horizontal alignment of the image of the coded aperture and to ensure horizontal dispersion across the CCD.

3.1.2 Zemax modeling

Zemax was used to model and optimize the design and layout of the imaging optics and prism. These components included the coded aperture, the F/8 relay lens, the double Amici prism and the CCD, in that order. As part of the optimization of the design, spot diagrams, modulation transfer function plots and geometric image analysis were used to determine if the design was satisfactory. The relay lens was

not designed for use with any other optical elements. However, the prescription of the lens was modified with the introduction of the double Amici prism between the relay lens and the CCD, as shown in figure 3.1(a). After optimization of the design, the relay lens was positioned 39.3 mm from the coded aperture plane and 43.8 mm from the CCD. The overall design had a paraxial magnification of -1 and a working F/# of 8.04.

Spot diagrams

Figures 3.5, 3.6 and 3.7 show the spot diagrams of rays corresponding to wavelengths of 450, 550 and 650 nm respectively from 9 different field points on the coded aperture plane mapping to 9 different points on the detector array. In this particular configuration, the front of the prism was located 8.6 mm from the relay lens and the back of the prism was located 23.2 mm from the CCD. The RMS radii of the spots ranged between $7-9$ μm at 450 nm, $3-8$ μm at 550 nm, and $3-8$ μm at 650 nm. The Airy disk is the black circle at each field point. The radius of this disk is given by $1.22\lambda F/\#$, where λ is the wavelength and $F/\#$ is the F-number of the system. The radius of the Airy disk ranged between $4.4-6.4$ μm over $450-650$ nm. The system was not diffraction limited, as the spot sizes are larger than the size of the Airy disk, a result of the double Amici prism perturbing the imaging properties of the relay lens.

The system performance would have been ideal if all the spots were inside the Airy disk, which would mean that the system is diffraction limited. Being diffraction limited implies that the PSF of the imaging system is limited by the physical effects of diffraction rather than imperfections in the design. The performance of this configuration was clearly not ideal. This can be explained by the fact that rays from the relay lens passed through the direct view double Amici prism before hitting the detector, even though the relay lens was designed to relay the image from one image

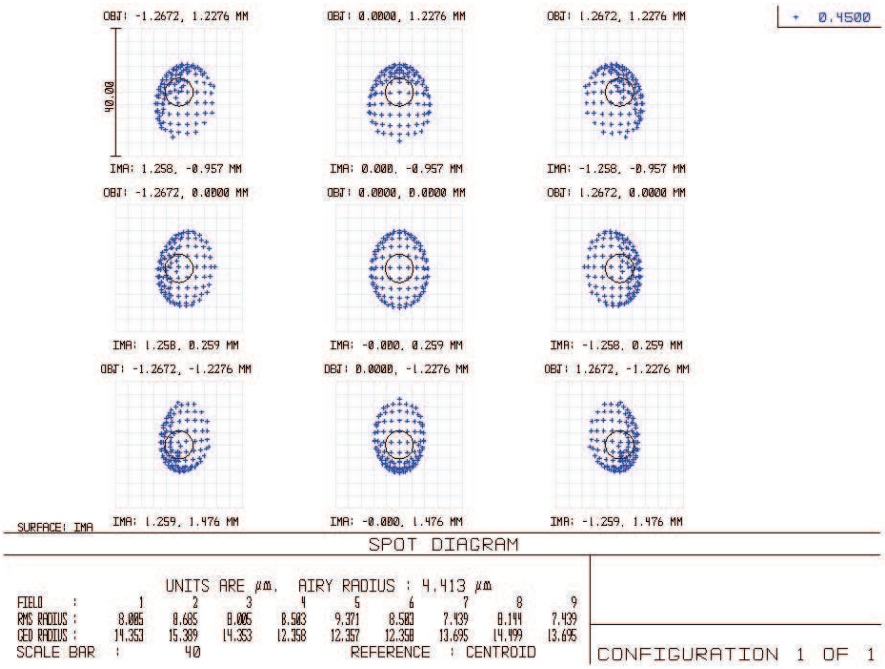

FIGURE 3.5: Zemax spot diagrams of 450 nm rays at 9 different field points at the detector. The RMS radii of the spots range between $7 - 9$ μm.

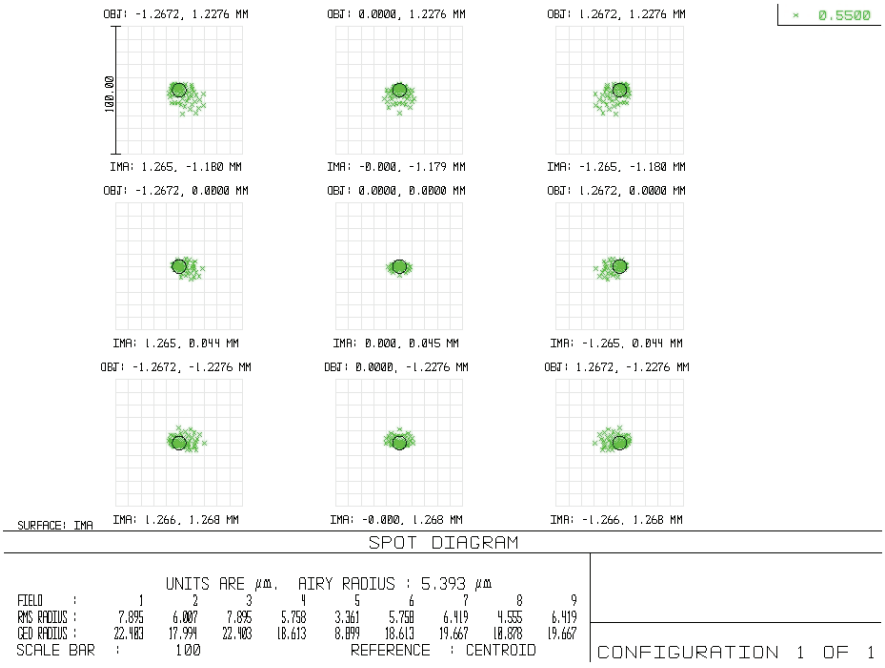

FIGURE 3.6: Zemax spot diagrams of 550 nm rays at 9 different field points at the detector. The RMS radii of the spots range between $3 - 8$ μm.

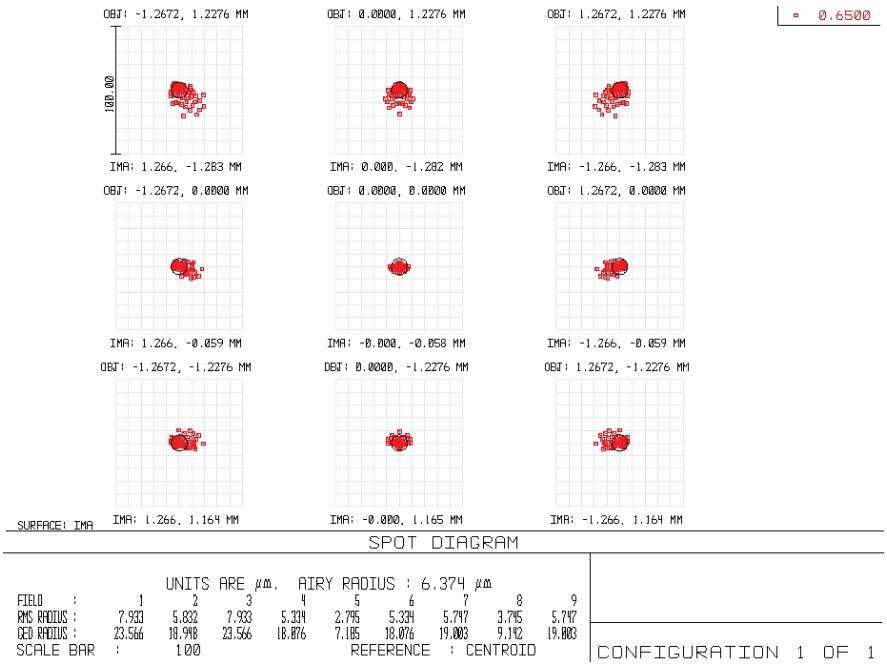

FIGURE 3.7: Zemax spot diagrams of 650 nm rays at 9 different field points at the detector. The RMS radii of the spots range between $3 - 8$ μm.

plane (in this case, the aperture code plane) to another image plane (in this case, the detector plane). The introduction of the extra glass surfaces between the relay lens and the detector led to the introduction of the aberrations in the form of the larger spot sizes.

Modulation transfer function (MTF) plots

Figure 3.8 shows the diffraction modulation transfer function data for all the field positions being simulated. This tells us how well spatial modulation on the object (in this case, the aperture code) is transferred to an image by the optics.

It was important to verify that even with relatively large spot sizes, there would be sufficient contrast between images of the open and closed mask features. Since the smallest mask feature was 2 detector pixels wide (19.6 μm = 0.0196 mm), the highest spatial frequency that needed to be resolved properly is given by $\frac{1}{2\times 0.0196} = 25.5\ \frac{cycles}{mm}$. The lowest value of the MTF at this spatial frequency is approximately 0.45, which

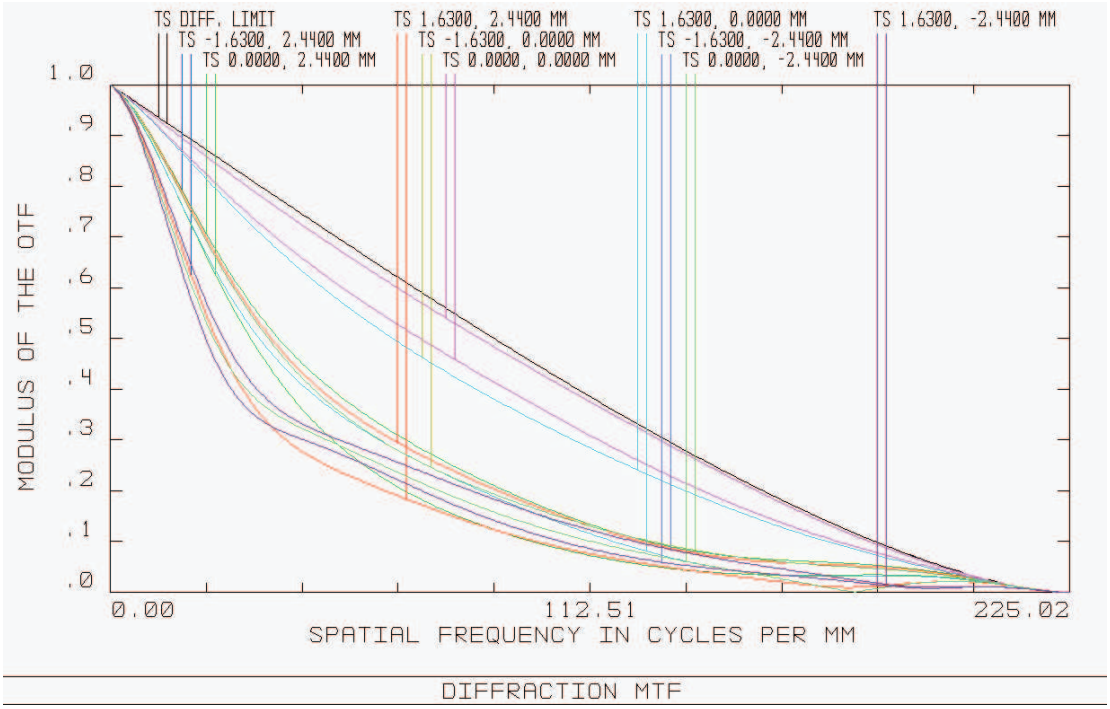

FIGURE 3.8: Zemax modulation transfer function (MTF) plots at each of the 9 field points. The information provided by these plots can be used to determine whether the contrast ratio between the open and closed features of the aperture code once they are imaged on to the detector array is sufficient.

indicates the contrast ratio between open and closed features on the aperture code once they are imaged on the detector. This ratio was large enough to distinguish the code features and as will be shown later, it was sufficient to recover a datacube in a simulation.

Geometric image analysis

Figure 3.9 demonstrates the imaging of three squares emitting at three different wavelengths (486 (blue), 587 (green) and 656 nm (red)) through the instrument. Figure 3.9(a) shows an image of the three squares after they were modulated by the aperture code. Figure 3.9(b) shows simulated detector plane images of the three squares after they are dispersed by the double Amici prism.

The effect of the non-diffraction limited PSF is noticeable as blurring around all

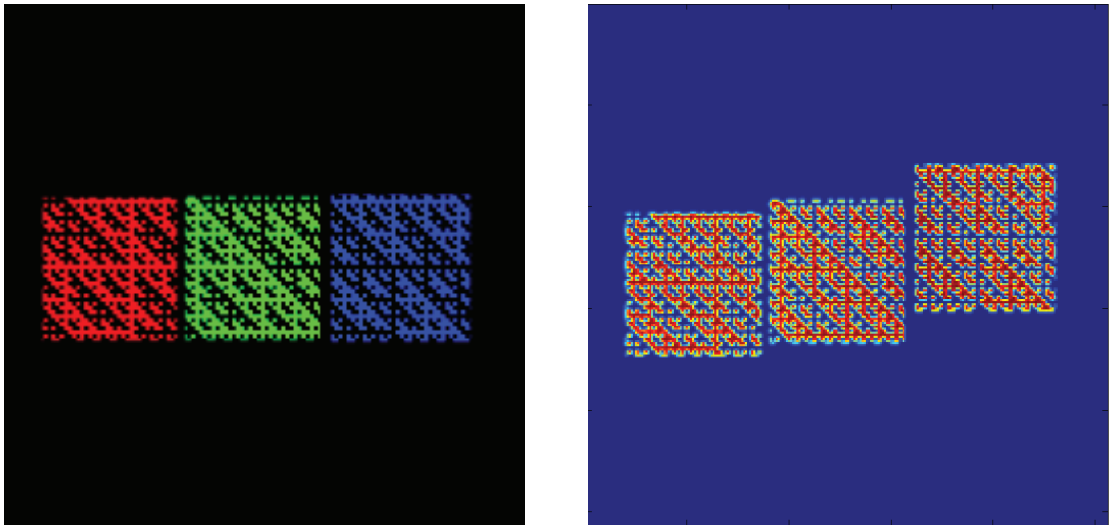

FIGURE 3.9: (a) Aperture code modulated image of the scene consisting of three squares radiating at 486 nm, 587 nm and 656 nm. (b) Zemax simulated images of the three squares after they are imaged through the double Amici prism on to the detector.

the mask features superimposed on the three squares. The effect of the dispersion by the prism is also noticeable, with the square emitting at 587 nm (close to the center wavelength of 550 nm) being only slightly displaced from its original location, while the two other squares are dispersed to either side of the 587 nm square.

Simulation after Zemax modeling

Figure 3.10(a) shows a 256×256 pixel random aperture code that was designed in Matlab, with each aperture code feature being 2 pixels wide. Reconstruction of the peppers image from a reconstructed datacube using this aperture code is shown in figure 3.10(b). The reconstruction was performed using 200 iterations of the GPSR algorithm. The system's response to a fully illuminated aperture code at 587 nm was modeled in Zemax and is shown in figure 3.10(c). The responses of the instrument at the other wavelengths in terms of blurring appeared to be virtually identical to that at 587 nm. There also did not appear to be any noticeable wavelength dependent

anamorphic distortion. Thus, the response at 587 nm was assumed to be invariant with wavelength. Reconstruction of the peppers image using this Zemax-modeled aperture code after 200 iterations is shown in figure 3.10(d).

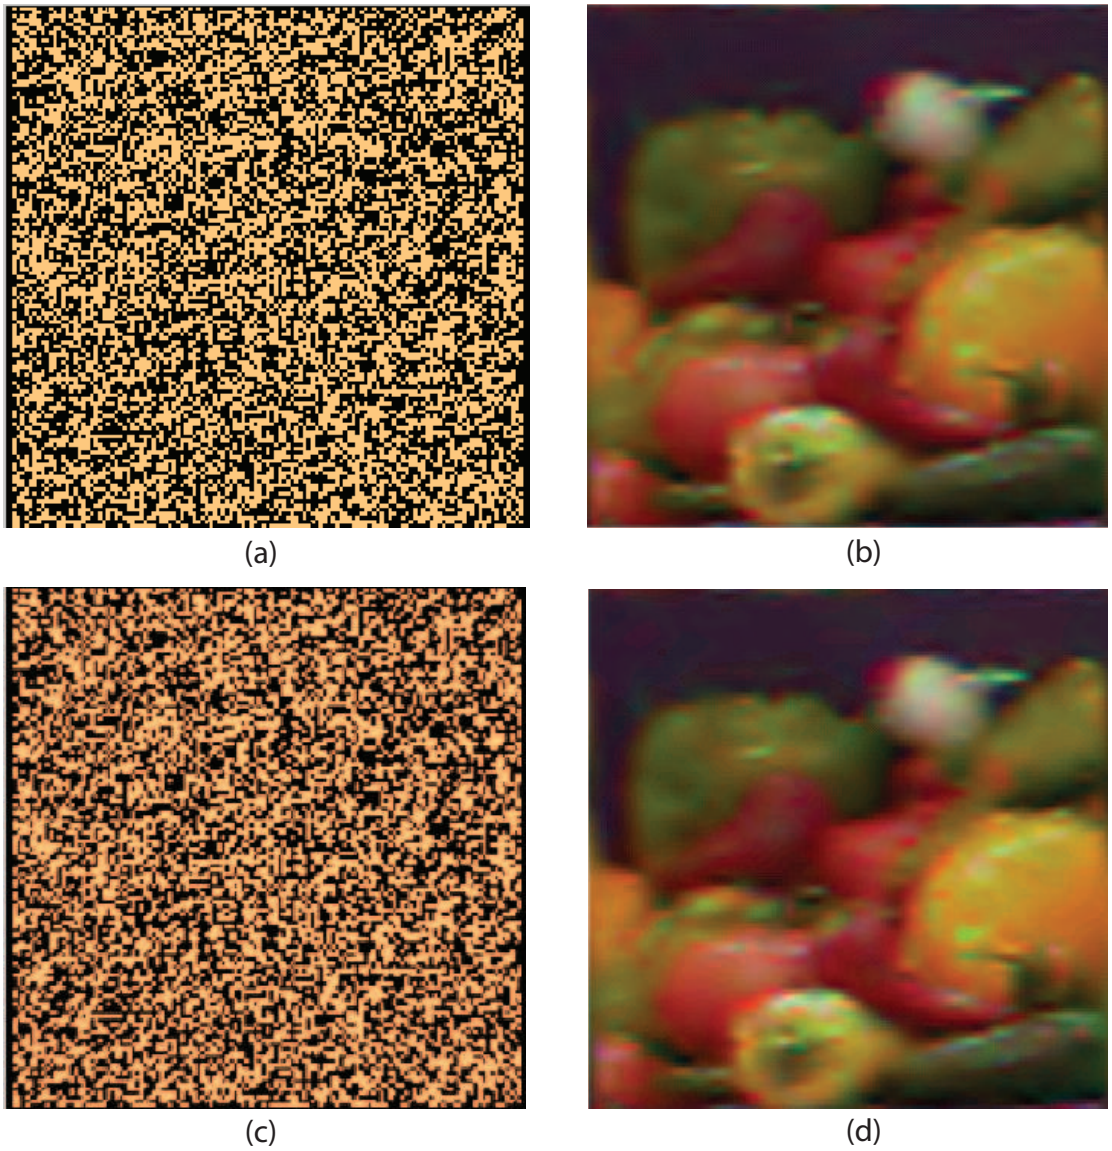

FIGURE 3.10: Simulation used to investigate the impact of optical blurring induced by the CASSI system PSF on the quality of the reconstructed datacube. (a) 256×256 binary aperture code, with features 2 pixels wide. (b) RGB image of peppers generated after reconstruction of the datacube using the aperture code in (a). (c) 256×256 aperture code after modeling in Zemax. (d) RGB image of peppers generated after reconstruction of the datacube using the aperture code in (c).

There was no noticeable difference between the reconstruction using this aperture code and the original aperture code. This suggested that there would not be any significant impact of the optical distortions present in the experimental prototype on the reconstruction of the datacube using GPSR.

3.2 CASSI calibration

After constructing the second generation CASSI prototype, a more rigorous calibration procedure was developed to improve the process of modeling the CASSI sensing matrix, Φ. The section begins with an examination of the structure of Φ and how it is affected by experimental non-idealities of the system. The following subsections detail the process of calibrating the experimental CASSI prototype.

3.2.1 Structure of the CASSI sensing matrix, Φ

It is useful to consider what the CASSI sensing matrix looks like. In the absence of any optical distortions, CASSI implements three operations on the voxelized datacube that is incident on the coded aperture plane. If this datacube has 8×8 spatial elements and 4 spectral channels (a total of 256 voxels), the detector makes a total of $8 \times (8 + 4 - 1) = 8 \times 11 = 88$ multiplexed measurements. The voxelized datacube is first spatially modulated ('punch' operation) by the coded aperture, then sheared by the dispersive element ('shear' operation) and then integrated along the wavelength dimension on the detector ('smash' operation). These three matrix operations are demonstrated in figure 3.11. The product of these three matrices represents the sensing matrix, Φ, that is shown in figure 3.12(a) and is a very coarse approximation to the true experimental sensing matrix.

As figures 3.5, 3.6 and 3.7 illustrate, the PSF of the second generation CASSI varies across the field and with wavelength, making it a shift variant blur. Figure 3.12(b) demonstrates the impact of capturing some of the blur as a modification of

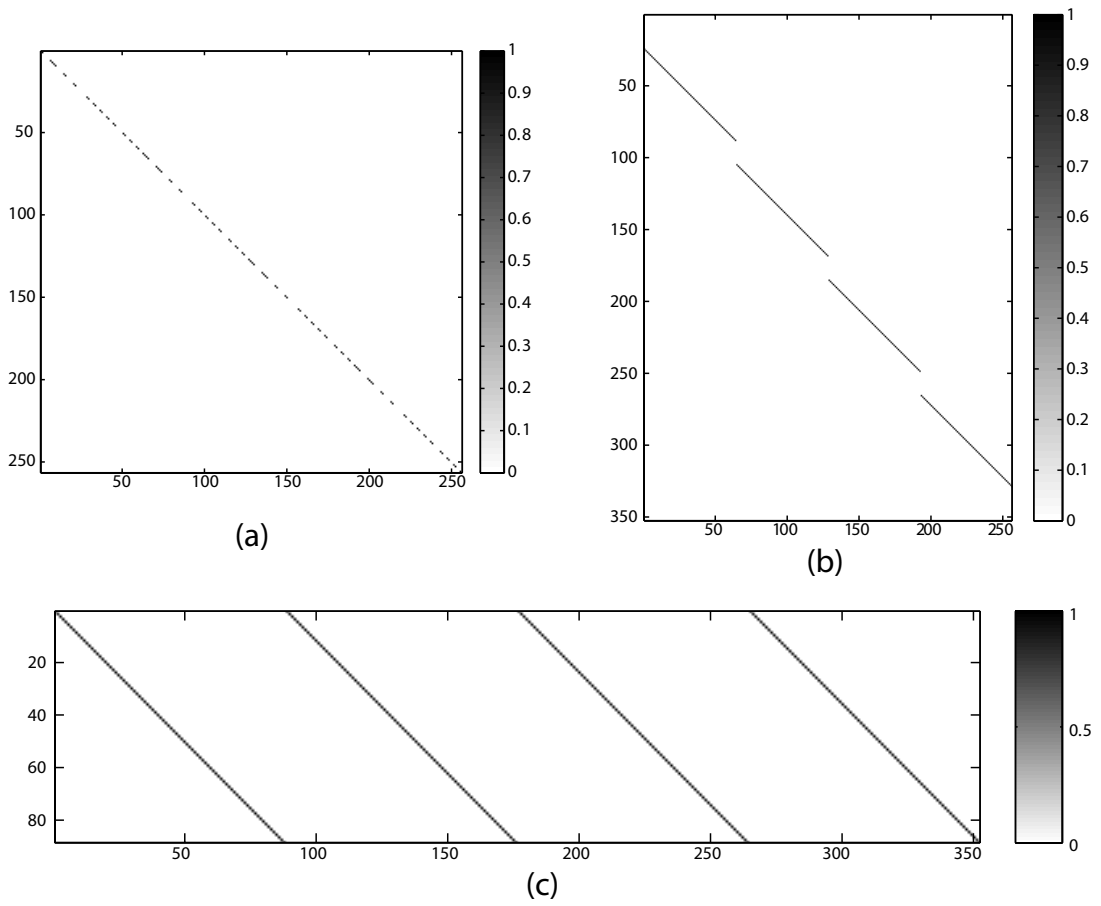

FIGURE 3.11: Matrices representing each step of the three step CASSI sensing process on a 256 voxel datacube. The square matrix in (a) represents the 'punch' operation by the coded aperture. The tall matrix in (b) represents the 'shear' operation by the dispersive element. The matrix in (c) represents the 'smash' operation involving the integration of the datacube along the wavelength axis by the detector.

the aperture code from a binary code as shown in figure 3.12(a) to a grayscale code. Note that the sensing matrix used by the GPSR algorithm to reconstruct a datacube of the ping pong balls in chapter 2 had this form. However, this was still quite far from the true experimental sensing matrix, as it did not incorporate the effect of a wavelength dependent, horizontal anamorphic distortion.

Figure 3.12(c) demonstrates the impact on Φ of approximating CASSI's shift variant blur as a combination of a shift invariant blur and a grayscale aperture code

that captures the spatial variation of the blur across the field.

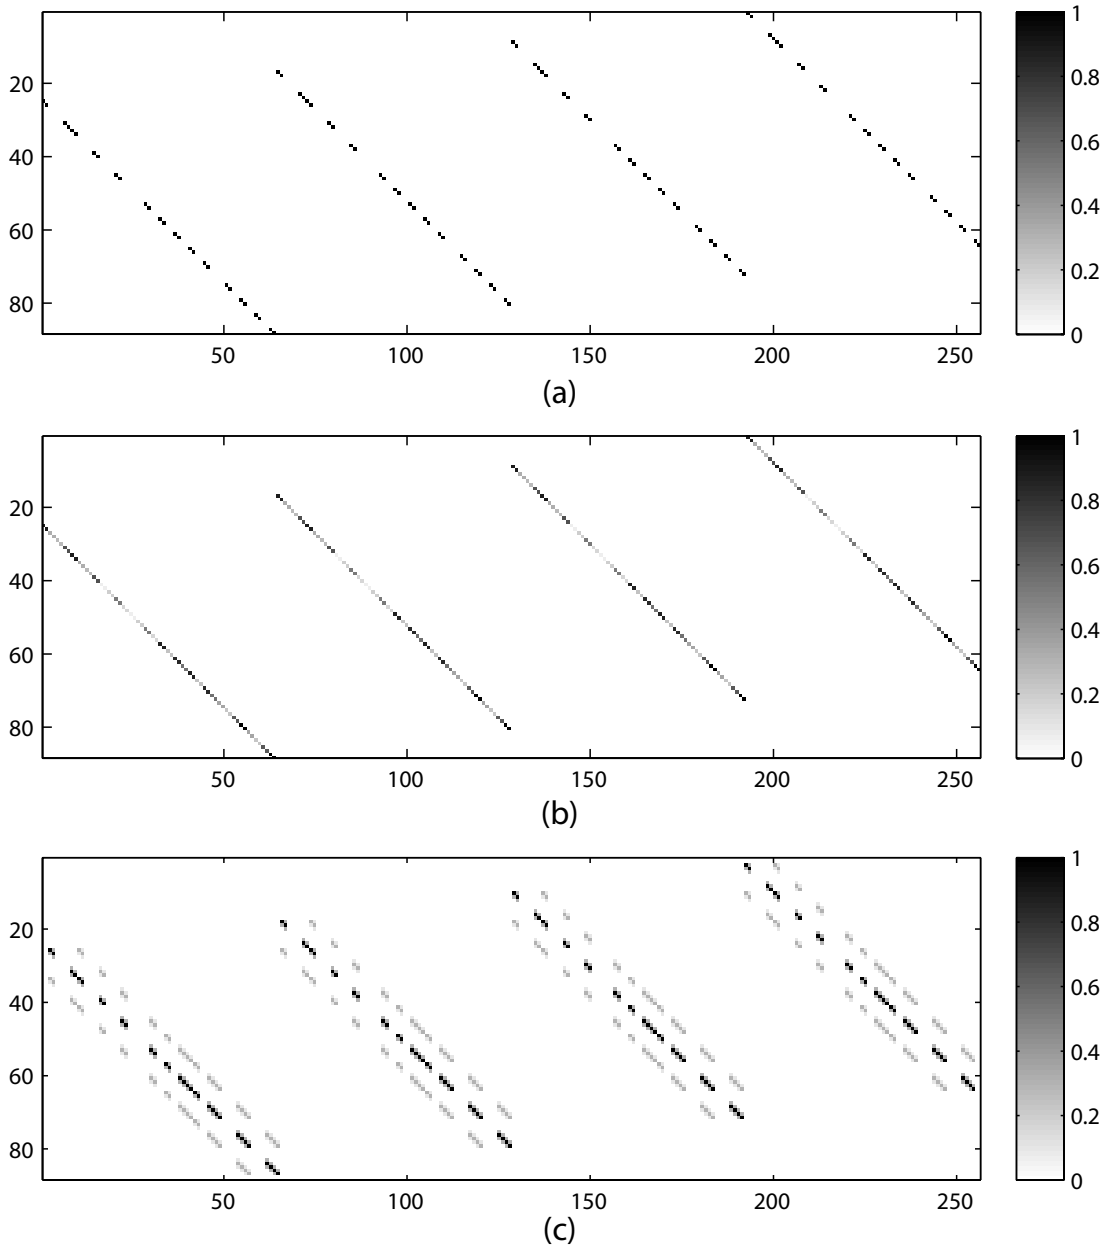

FIGURE 3.12: Effect of incorporating different amounts of optical blur into the CASSI sensing matrix, Φ. (a) Φ mapping a 256 voxel datacube to 88 detector pixels. All optical distortions are ignored. (b) Effect of modeling the spatially varying blur as a modification of the aperture code from a binary code to a grayscale code. (c) Effect of modeling the spatially varying blur as a a binary aperture code convolved with a shift invariant blur.

In deriving the matrix-vector expression shown in equation (2.12) for the CASSI system operator, three major assumptions were made: (i) the PSF $h(x, x', y, y', \lambda)$ is shift invariant, (ii) the dispersion by the double Amici prism is linear, and (iii) that there is one-to-one mapping between elements of the aperture code to the detector pixels. These assumptions made it easy to interpret the CASSI measurement as a linear process. Experimentally, however, the PSF varies across the field, the dispersion is non-linear over CASSI's spectral range, and there are subpixel misalignments between the aperture code features and detector pixels. These factors are non-idealities that deviate Φ from the ideal form in figure 3.12(a) and the coarsely calibrated form in figure 3.12(b). The form of Φ in figure 3.12(c) is a significantly closer approximation to the experimental Φ, although it ignores sub-pixel misalignments and the non-linear dispersion by the double Amici prism.

3.2.2 Calibration process

As mentioned in section 2.8, the quality of the datacube estimate generated by any numerical estimation algorithm is strongly correlated with how well Φ is characterized. Element-wise calibration of Φ for the second generation CASSI would require the capture of the detector response to a spatio-spectral scan of a point source. As this process would require considerably more equipment, a slightly simplified calibration process was performed that did not fully capture the spatio-spectral, shift-variant PSF that is suggested by the spot diagrams in section 3.1.2. Sub-pixel misalignments of the aperture code with the detector and the non-linear dispersion of the double Amici prism could be accounted for.

To experimentally characterize Φ, detector images of the aperture code were taken at 231 equally spaced wavelengths from 440 to 670 nm after uniformly illuminating the coded aperture with each monochromatic wavelength of light within and around the bandpass ($450 - 650$ nm) of the system. The images of the aperture code were

grayscale (i.e. not binary) and could be used to construct a Φ of the form shown in figure 3.12(b). The image of the cross on the top of the aperture code could be used to estimate the spectrally-varying, shift-invariant part of the PSF. This additional information could be used to construct a Φ that would closely resemble the form shown in figure 3.12(c), and would be a significantly closer approximation to the true, experimental CASSI sensing matrix.

The calibration process was conducted with careful efforts to reduce the impact of dark noise on the CCD and the non-uniform spectral intensity of the calibrating light source. A summary of key processes performed during calibration of CASSI includes:

1. **Illumination control**: Every effort was made to illuminate the aperture code with the light from a monochromator as uniformly as possible.

2. **Shot-noise reduction**: At each wavelength, 10 CCD frames were captured and averaged to reduce the impact of shot and readout noise.

3. **Background subtraction**: At each wavelength, 10 dark frames were captured at the same exposure time as the bright frames and averaged. The averaged dark frame at each wavelength was then subtracted from its corresponding bright calibration frame.

4. **Exposure time adjustment**: To improve the signal-to-noise ratio (SNR) of the aperture code image at each wavelength, the exposure time at each wavelength was scaled so that the mean counts over the coded aperture in the CCD measurements at all wavelengths were similar.

5. **Light source spectral intensity distribution**: Light from the source at each wavelength was measured with a photodiode having a known responsivity

curve to obtain a calibration curve for the non-uniform spectral intensity of the light source.

The wavelength dependent images of the aperture code obtained after background subtraction were normalized by the non-uniform exposure time curve and the non-uniform spectral intensity distribution of the light source.

3.2.3 Calibration results

The procedure described above was performed to capture 231 monochromatic images of the aperture code in 1 nm increments between $440 - 670$ nm. The calibration process yielded two important results to help characterize CASSI. Figure 3.13 shows 33 monochromatic images of the coded aperture, each displaced by one column of CCD pixels. As the wavelength increases, the image of the aperture code on the detector shifts from right to left due to the dispersion by the double Amici prism. As described above, these calibration images could be used to construct the CASSI sensing matrix, Φ. For the results described in this chapter, these 33 wavelengths were defined as the centers of 33 spectral channels that any numerical estimation algorithm would attempt to estimate the content of.

Although the center wavelengths of the chosen spectral channels were separated by one column of CCD pixels, each channel had a different bandwidth due to the non-linear dispersion of the double Amici prism. The purple dots in figure 3.14 track the position of the cross on top of the monochromatic image of the coded aperture as a function of wavelength. In essence, this non-linear curve discretely captures the double Amici prism's dispersion coefficient $\gamma(\lambda)$, first introduced in equation (2.3). The dispersion at the blue end of the spectrum is much greater than in the red. The black crosshairs in figure 3.14 identify the 33 wavelengths that define the centers of the 33 spectral channels.

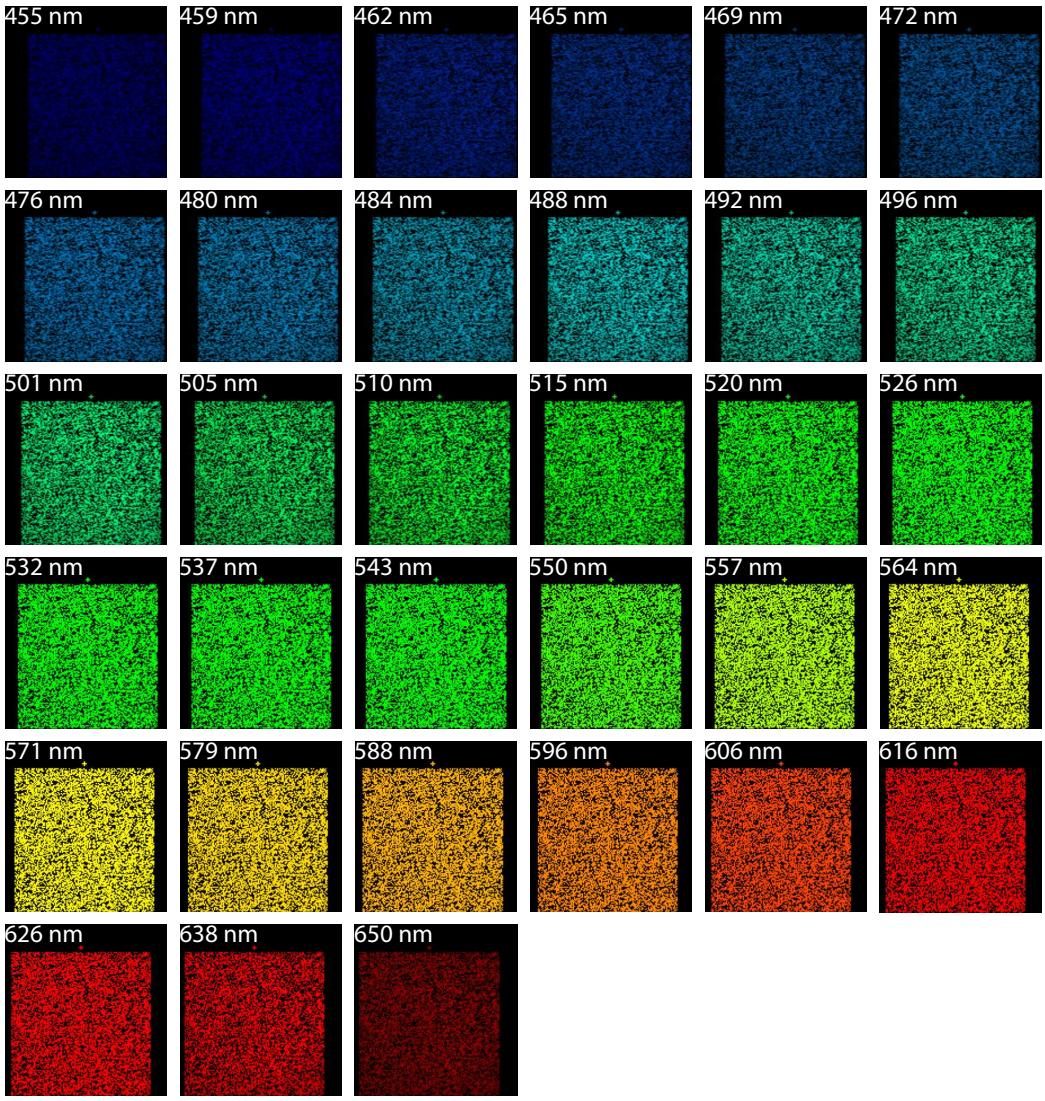

FIGURE 3.13: Demonstrating CASSI's response to uniform illumination at 33 different wavelengths. These wavelengths define the centers of CASSI's 33 spectral channels.

3.3 TwIST numerical estimation method for CASSI

Like the GPSR method introduced in section 2.7, the Two-step Iterative Shrinkage/Thresholding (TwIST) algorithm, introduced by Bioucas-Dias et al. [33], may be used to solve CASSI's under-determined linear system of equations in (2.12) by minimizing a convex objective function. The Matlab code for TwIST is available on-

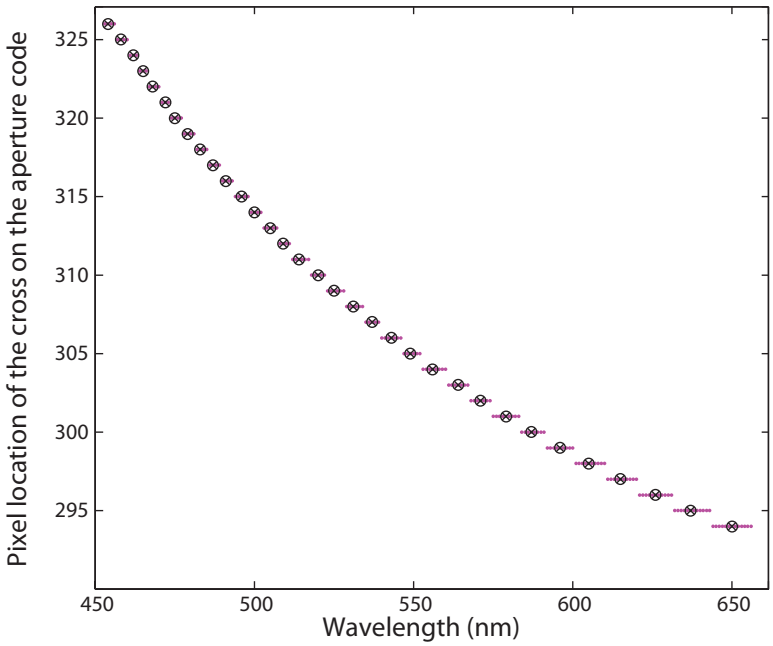

FIGURE 3.14: The purple dots track the location of the cross on the top of the aperture code as a function of wavelength and demonstrate the non-linear dispersion by the double Amici prism. The black crosshairs identify the 33 wavelengths that define the centers of CASSI's 33 spectral channels.

line at http://www.lx.it.pt/∼bioucas/TwIST/TwIST.htm. In the context of image reconstruction for a CASSI system, TwIST may be adapted to estimate a datacube as a solution to the non-linear, unconstrained minimization problem

$$\hat{f}_{TwIST}(\tau,\Gamma) = \left[\underset{f}{\operatorname{argmin}} \left\{ \frac{1}{2}\|g - \Phi\, f\|_2^2 + \tau\Gamma(f) \right\}\right], \tag{3.2}$$

where the choices for the regularization function $\Gamma(f)$ include, but are not limited to, the ℓ_1 norm. As an alternative to ℓ_1 regularization of wavelet coefficients that was previously used with GPSR, the total variation (TV) regularizer defined as

$$\|f\|_{TV} = \sum_k \sum_{i,j} \sqrt{[f(i+1,j,k) - f(i,j,k)]^2 + [f(i,j+1,k) - f(i,j,k)]^2}, \tag{3.3}$$

may be used. The TV terms penalizes potential solutions with higher discrete horizontal and vertical gradients in the spatial dimension. When the datacube is spatially piecewise smooth, the gradient of a spectral channel in the datacube is sparse.

With this choice of regularizer, the TwIST method searches for a datacube estimate, $\hat{f}_{TwIST}$, with a sparse representation in the spatial gradients of each of its spectral slices. The first term in equation (3.2) minimizes the ℓ_2 error between the measurements modeled from the estimate and the true measurement. The second term is a penalty term that encourages sparsity on the TV basis of each spectral slice and controls the extent to which piecewise smooth estimates are favored. Just as with the GPSR formulation, τ is a tuning parameter for the penalty term and higher values of τ yield spatially smoother estimates of f.

Empirically, the estimation of the datacube using TV instead of wavelet basis operations has at least two important advantages. First, computing the TV of each spectral slice of the datacube is much faster than computing the redundant wavelet transform, making TwIST iterations much faster than GPSR iterations. Second, TV has no restriction on the spatial dimensions of the datacube. As a result, the separability of the CASSI system model in the vertical dimension can be exploited. For example, estimating an $N_x \times N_y \times N_\lambda$ datacube from an $N_x \times (N_y + N_\lambda - 1)$ detector measurement can be performed by stitching two $\frac{N_x}{2} \times N_y \times N_\lambda$ datacubes estimated by dividing the detector measurements into two $\frac{N_x}{2} \times (N_y + N_\lambda - 1)$ arrays and processing these arrays on separate computers in parallel.

3.3.1 Spectral imaging of a static scene using TwIST

To test the use of TwIST for spectral image recovery, a snapshot CASSI measurement was captured of a static scene that was constructed under fluorescent room lights. The scene consisted of four plastic objects of different colors and shapes. The objects included a red apple, yellow banana, green pineapple and a blue stapler, which were all set against a black backdrop as shown in figure 3.15(a), which was captured using a Canon RGB digital camera.

As mentioned in section 3.1.1, the aperture code used by CASSI was an array

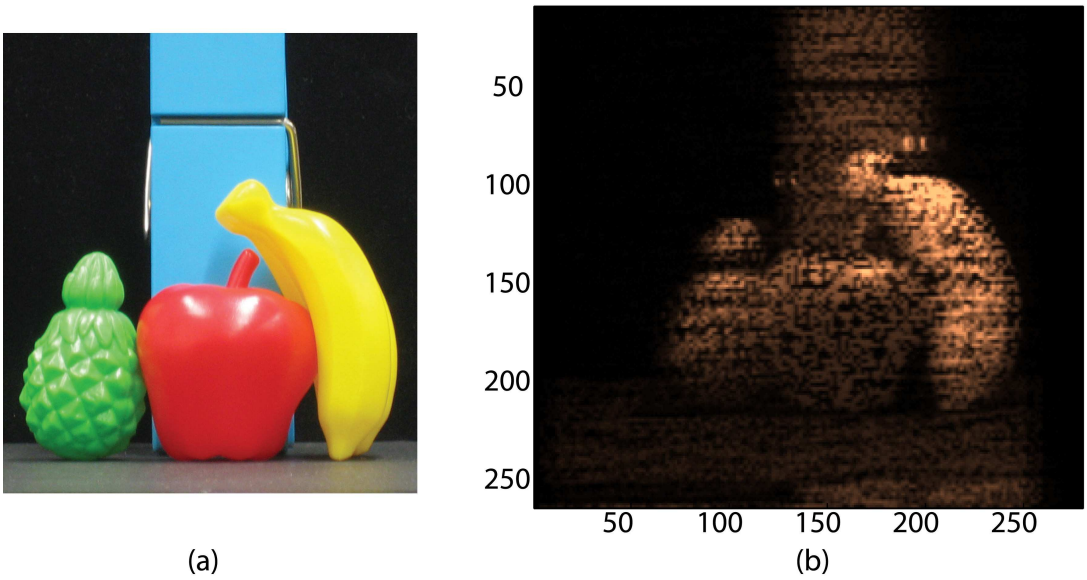

FIGURE 3.15: (a) A scene consisting of colorful objects illuminated by fluorescent room lights (b) the corresponding CASSI detector measurement of the scene.

of 256×248 elements with each element the same size as the detector pixel. The feature size on the aperture code was twice as large as the detector pixel on each side. The CASSI detector measurement consisted of 256×280 pixels, as shown in figure 3.15(b). It is easy to observe the spatio-spectral overlap of the aperture code modulated images of the four objects due to the dispersion. The 82 ms exposure was taken under fluorescent lights in the ceiling of the lab where the experiment was performed.

To generate an estimate of the datacube representing the spectral image of the scene, the TwIST method was provided with two inputs including (i) the snapshot CASSI detector measurement of the scene and, (ii) a Matlab functional form of the sensing matrix, Φ, constructed using the 33 grayscale images of the aperture code at distinct wavelengths captured during the calibration process. The resulting sensing matrix had a form that closely resembled that shown in figure 3.12(b) and was used to model the mapping from a $256 \times 248 \times 33$ voxel 3D datacube to a 256×280 pixel

2D detector measurement.

Figure 3.16 shows the spectral image estimate with 33 spectral channels generated over 50 iterations of the TwIST algorithm using the TV regularizer described by equation (3.3). The spatial content of each of 33 spectral channels between 454 nm to 650 nm is shown. Note that the modulation by the aperture code on the spatial structure visible in figure 3.15(b) has been removed from all the spectral channels. The GPSR algorithm with ℓ_1 regularization of the redundant wavelet transform coefficients of each spectral slice was unable to handle this spectral image estimation problem involving more than 2 million unknown voxels. The value of the regularization parameter τ was chosen to be 0.1 based on visual inspection of the reconstructed datacube at the end of 50 TwIST iterations. There was no noticeable change beyond this point. The initial estimate of the datacube was defined as $f_{initial} = \Phi^{\mathrm{T}} g$. The algorithm would take many more iterations to converge to a solution if it were initialized with a random initial estimate of the datacube.

To validate the datacube estimate generated using the CASSI measurement, the spectral information obtained by a non-imaging spectrometer was used as a quantitative reference. Specifically, the spectral signature of each object in the scene was measured using a non-imaging, point-wise spectrometer (Ocean Optics USB2000). The tip of the fiber connected to the spectrometer was brought close to each object to obtain a point-wise spectral signature. Figure 3.17 shows the spectra obtained from the reference spectrometer. The spectra suggest that the color of each object in the scene simply modulates the spectrum of the fluorescent room lights. Note that the spectra are fairly peaky, making the spectral content of the scene relatively sparse. These spectral signatures were then integrated into the 33 spectral channels used to recover a datacube from the CASSI measurement using TwIST. The plots in figure 3.18 demonstrate the agreement between the reference spectral signatures (in blue dots) and the spectral signature of the selected spatial points on each object

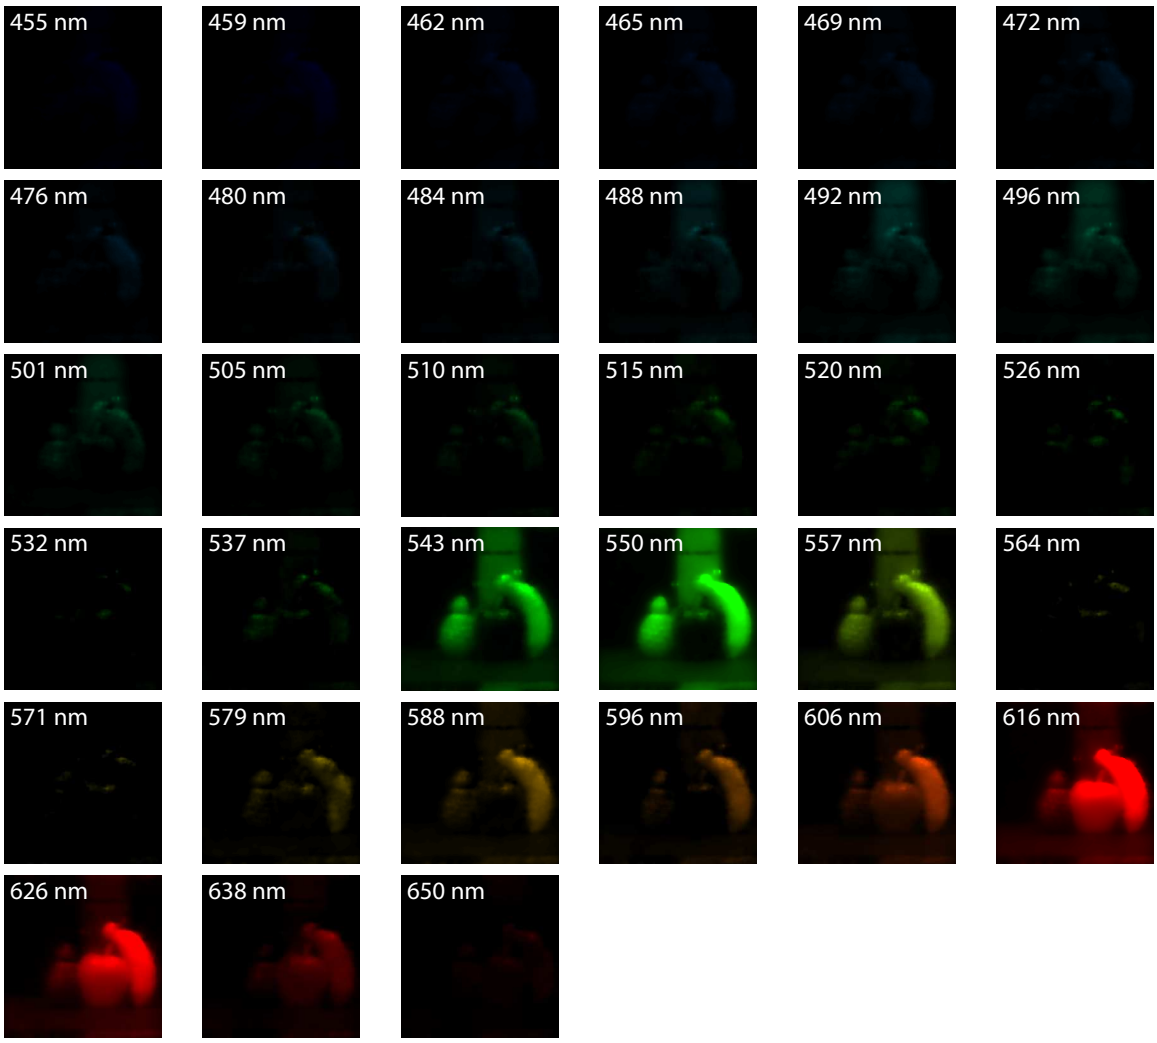

FIGURE 3.16: Datacube estimate of the scene in figure 3.15(a) recovered from the CASSI measurements in figure 3.15(b) using the TwIST algorithm.

in the computed estimates (in magenta dots), with TwIST. The spectra are normalized by the total intensity per object. The relative intensity between the objects is not shown, as the intensity of the blue stapler is much smaller in comparison to the intensity of the red apple in the datacube estimates. Note that differences in the system responses of CASSI and the reference spectrometer were not accounted for in this comparison.

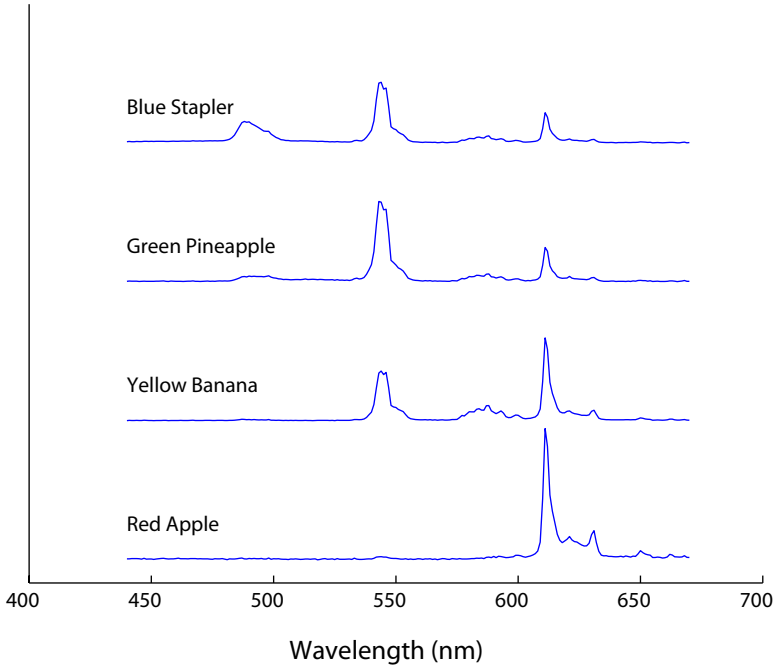

FIGURE 3.17: Spectra of the four objects shown in figure 3.15(a) at equally spaced 1 nm wavelengths, as measured by a non-imaging, reference spectrometer.

3.4 NeAREst numerical estimation method for CASSI

Datacube reconstruction using the GPSR and TwIST algorithms, as discussed in sections 2.8 and 3.3.1, is based on solving a single system of linear equations described by equation 2.12). As the number of voxels in the datacube to be reconstructed exceeds the number of pixels in the CCD measurement, the system in equation (2.12) is under-determined, having infinitely many solutions. To overcome the algebraic limitation on the number of voxels that can be reconstructed and the corresponding detail they reveal, both algorithms assume additional information to favor a particular solution over the others. The best solution is found by reformulating the system of linear equations into a convex, unconstrained optimization problem with an additive regularization term. However, the effective use of these methods is limited to situations that allow some human interaction in the reconstruction process to determine a data-dependent regularization parameter, such as one specifying the degree

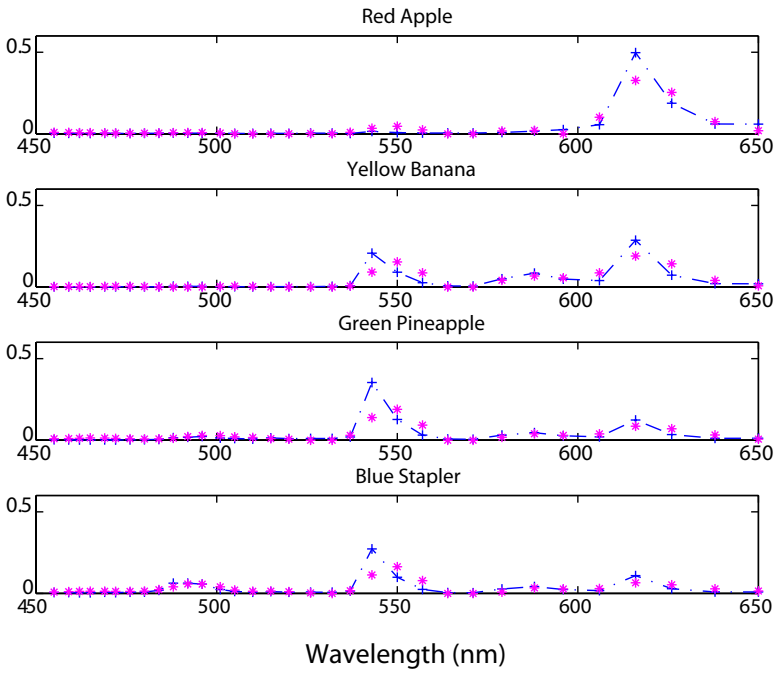

FIGURE 3.18: Spectra of the bodies of the four objects binned into 33 channels. The purple dots are the CASSI spectral signatures at points on the objects after datacube recovery using the TwIST method. The dashed blue curves are the spectral signatures measured by the reference spectrometer after being integrated into 33 spectral channels with equivalent bandwidths to the CASSI spectral channels.

of sparsity of the datacube.

An alternative method to reconstruct a datacube from a CASSI measurement is called Nested Adaptive Refinement Estimation, or NeAREst for short [34], developed by Dr. Xiaobai Sun and Dr. Nikos Pitsianis in the Computer Science department at Duke University. As of February 2010, the Matlab code for NeAREst is yet to be released to the public. Unlike the GPSR and TwIST algorithms, this algorithm does not require the specification of the value of a manually tuned regularization parameter. This property makes it convenient for spectral imaging of a dynamically changing scene, a situation where a sequence of datacubes must be reconstructed from a corresponding sequence of CASSI video frames. The use of NeAREst for recovery of spectral images of both static and dynamic scenes will be demonstrated later in section 3.5.

3.4.1 A nested set of discretized systems

The NeAREst method utilizes the fact that propagating rays of light can be bundled into different granularity levels, both spatially and spectrally. Spectrally, the bandwidths of the spectral channels may be coarser or finer. Spatially, features on the aperture code may be represented in terms of their true size of 2 CCD pixels each or at a finer level where the aperture code can be considered as consisting of virtual features that are the same size as the detector pixels. NeAREst builds and solves a nested set of systems of equations at multiple spatio-spectral scales. Each system is a discretization of the continuous light propagation model in equation (2.4) at a particular scale. As shown in figure 3.19, the vector of detector measurements, g, can then be represented as a matrix vector product between Φ_s and f_s at a particular scale s. Across all scales, the system of equations has the same form, but varies in the number of degrees of freedom, i.e. the number of voxels in f_s to be estimated. Pictorially, this amounts to varying the size of the sensing matrix and thus controlling the size of the datacube that is estimated, while retaining the size of the measurement vector, g. In principle, the propagation model in equation (2.4) embeds all the NeAREst multi-scale systems and the systems are coupled by a cascading discretization framework. In practice however, one must also incorporate system-specific parameters, which are captured through the calibration process, into the nested systems. As mentioned in section 3.2, the calibration process is necessary to account for the non-linear dispersion by the double Amici prism and the optical blur in the system.

First consider the discretization scheme used to construct the sensing matrix at a given scale, *without* accommodating for CASSI system specifics. The power spectral

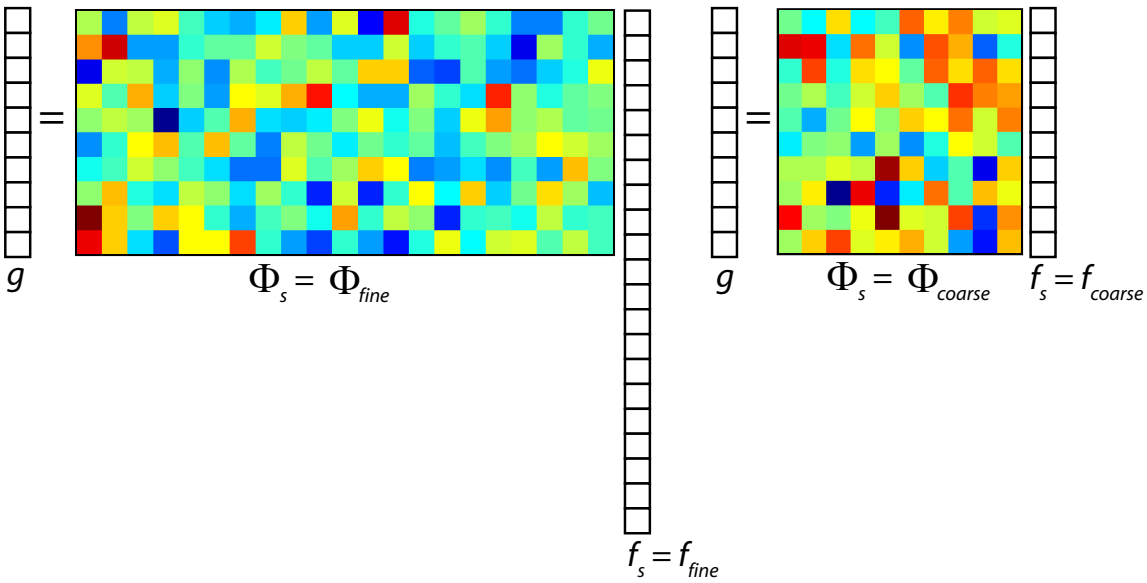

FIGURE 3.19: NeAREst solves a nested set of systems of equations at multiple spatio-spectral scales. Here two systems at a fine and a coarse scale are represented pictorially. Across different scales, the system varies in the number of degrees of freedom. The matrix Φ_s at a coarse scale has fewer columns than that at a finer scale, but the former is not necessarily a sub-matrix of the latter. The datacube estimate must satisfy the systems at all scales. Thus, a solution to the system at a fine scale is disqualified if it does not yield a solution at a coarse scale.

density at the plane of the CASSI detector is

$$g(x, y, \lambda) = \sum_{i,j} \iint_{A_{i,j}} G_{i,j}(x, y; x', y', \lambda)\, dx'\, dy', \tag{3.4}$$

with

$$G_{i,j}(x, y; x', y', \lambda) = T_{i,j}(x', y')\, f(x', y', \lambda)\, h(x, x' - \gamma(\lambda), y, y', \lambda),$$

where $\gamma(\lambda)$ is the dispersion at wavelength λ and $A_{i,j}$ is the $(i,j)^{th}$ virtual sub-aperture at a particular scale. For example, at the finest scale, the size of the sub-aperture is the same size as the detector pixels, while at a coarser scale, the size of this sub-aperture may be the same size as the smallest aperture code element, i.e. 2×2 detector pixels. In equation (3.4), $T_{i,j}(x', y') \geq 0$ describes the effective aperture code at the $(i,j)^{th}$ sub-aperture. Note that even when the sub-aperture

is of the same size as the original code features, it is not necessarily equal to the binary-valued function in equation (2.2), which explains why equation (3.4) has a different form to equation (2.3). The intensity measured at the $(m,n)^{th}$ detector pixel can be described as

$$g_{n,m} = \sum_k \int_{\Lambda_k} \iint_{P_{n,m}} g(x,y,\lambda)\, dx\, dy\, d\lambda, \tag{3.5}$$

where $P_{n,m}$ denotes the spatial support of pixel (m,n) and Λ_k denotes the k^{th} spectral channel, which is determined by the calibration process and the scale at which the system of equations is being constructed.

Thus, at a scale characterized by $\{A_{ij}\}$ as the sub-apertures and $\{\Lambda_k\}$ as the spectral channels, there is a system of equations relating a discrete 3D spatio-spectral datacube $f(x_i, y_j, \lambda_k)$, for some $(x_i, y_j, \lambda_k) \in A_{i,j} \times \Lambda_k$, to the complete 2D array of CASSI measurements,

$$g_{n,m} = \sum_{i,j,k} Q(m,n;i,j,k) f(x_i, y_j, \lambda_k), \tag{3.6}$$

where

$$Q(m,n;i,j,k) = \iint_{P_{n,m}} \int_{\Lambda_k} \iint_{A_{i,j}} T_{i,j}(x',y')\, h(x, x'-\gamma(\lambda), y, y', \lambda)\, dx'\, dy'\, d\lambda\, dx\, dy.$$

$Q(m,n;i,j,k)$ may be approximated by applying numerical quadrature for integration, so that

$$Q(m,n;i,j,k) \approx w_{m,n} w_{i,j} w_k T_{i,j}(x'_i, y'_j) h(x'_i - \gamma(\lambda_k), x_m, y'_j, y_n, \lambda_k), \tag{3.7}$$

where $w_{m,n}$ and $w_{i,j}$ are the quadrature weights associated with the spatial discretization and w_k is the quadrature weight associated with the spectral discretization.

Equation (3.6) and the approximation to Q in equation (3.7) together describe the discretized system of equations at a given scale characterized by the spatio-spectral

discretization $\{A_{ij}\}$ and $\{\Lambda_k\}$. In a CASSI system, pixelization at the detector is uniform on a 2D cartesian grid. The discretization of the aperture at a given scale is also assumed to be uniform on a cartesian grid. Under such conditions, $w_{m,n}w_{i,j}$ is constant over all $(m, n; i, j)$.

3.4.2 Accommodation of CASSI system specifics

The discretization scheme presented in the previous subsection must be refined to produce a sensing matrix at each scale that accounts for the non-linear dispersion by the double Amici prism and the optical blur in CASSI. These system specifics are characterized through the calibration process described in section 3.2. The spectral channels are determined so that the centers of each channel, at the finest scale, are separated by one column of detector pixels. Thus, the position of a channel corresponds to a fixed dispersion in terms of detector pixels relative to a fixed spectral channel. At a coarser scale, the position of a spectral channel on the detector is linearly related to the channel index k as

$$\gamma(\lambda_k) = \alpha_\lambda k, \tag{3.8}$$

where α_λ represents the number of pixels of dispersion per spectral channel. For example, if the centers of two neighboring spectral channels are two pixels apart in dispersion, $\alpha_\lambda = 2$. Note that although equation (3.8) is similar to the linear assumption in equation (2.9), it holds even though the dispersion introduced by the double Amici prism is non-linear. This is under the assumption that the spectral channels are not equally partitioned in bandwidth. The non-linearity of dispersion is captured by the non-uniform spectral quadrature weights w_k across the spectral channels.

Upon substituting equation (3.8) into equation (3.7),

$$Q(m + \alpha_\lambda k, n; i, j, k) \approx c_k \, T_{i,j}(x'_i, y'_j) \, h(x'_i, x_m, y'_j, y_n, \lambda_k), \tag{3.9}$$

where $c_k = w_{m,n} w_{i,j} w_k$ varies only with the channel index, k, under the assumption of uniform spatial partitions on the detector and the aperture, as indicated in the previous subsection.

The calibration process is also required to account for the PSF, h, of the imaging system from the coded aperture to the detector array, which consists of the relay lens and the double Amici prism. As previously mentioned in section 3.2.2, h is known to vary both spatially and spectrally, but complete characterization of this function would require considerably more equipment. Thus, h is approximated by two factors, one that is spatially shift invariant, and another that captures the spatial variation at each spectral channel. Specifically, the operator Q is expressed in terms of the calibrated factors,

$$Q(m + \alpha_\lambda k, n; i, j, k) \approx c_k \, T_{i,j}^{[k]}(x'_i, y'_j) \, h^{[k]}(x'_i - x_m, y'_j - y_n). \tag{3.10}$$

Here, $h^{[k]}$ is assumed to be the spatially shift invariant part of the PSF that depends on the imaging optics and the k^{th} spectral channel. $T_{i,j}^{[k]}$ is the effective grayscale code calibrated at the corresponding channel and compensates for the difference between $h(x, y, \lambda_k)$ and $h^{[k]}(x, y)$.

In contrast to equations (3.6) and (3.7), equation (3.6) and the approximation to Q in equation (3.10) together describe the discretized system of equations at a given scale characterized by the spatio-spectral discretization $\{A_{ij}\}$ and $\{\Lambda_k\}$ after accounting for the calibration process. Simulations by Dr. Pitsianis and Dr. Sun revealed to them that the reconstruction quality of a datacube using the calibrated sensing matrices at each scale was much superior to that of a datacube reconstructed without accounting for the calibration process.

3.4.3 Solving the nested systems

Equation (3.6) and the approximation to Q in equation (3.10) describe the mapping of the 3D discrete datacube $f(x_i, y_j, \lambda_k)$ at a given scale to a measurement on each detector pixel $g_{n,m}$ for a calibrated CASSI. If the 2D array of detector measurements is represented as a vector g, then at a given scale, s, the expression in equation (3.6) can be represented as the matrix vector product

$$g = \Phi_s f_s = \sum_k \Phi_{s_k} f_{s_k}, \tag{3.11}$$

where f_{s_k} is a 2D slice of the 3D datacube at scale s associated with the k^{th} spectral channel and Φ_{s_k} is a sub-matrix of Φ_s relating f_{s_k} to its contribution to the detector measurements. Across the different discretization scales, the system varies in the number of degrees of freedom. The matrix Φ_s at a coarse scale has fewer columns than that at a finer scale, but the former is not necessarily a sub-matrix of the latter. NeAREst solves the linear system at each scale by transforming it into Csiszar's version of the Kullback-Leibler (KL) divergence as the objective function

$$\arg\min_{f_s \geq 0} \sum_{m,n} \{ g \cdot \log\left(\frac{g}{\Phi_s f_s}\right) - g + \Phi_s f_s \}_{n,m}, \tag{3.12}$$

with a slight modification to permit zero elements in g. In equation (3.12), the second term is the algebraic difference between g and $\Phi_s f_s$, while the first term is the weighted difference in the bitwise representation, with $g_{n,m}$ as the weights. The KL objective function is solved using the Richardson-Lucy (RL) iterative method [35].

Unlike the objective functions that are minimized using the GPSR and TwIST algorithms, the objective function for NeAREst in equation (3.12) does not require the specification of a regularization parameter. Nonetheless, there are two regularization means in effect. By design, the RL method excludes solutions with negative

values, thus imposing a non-negativity constraint. In contrast, when processing CASSI measurements with the GPSR and TwIST algorithms, this non-negativity condition would have to be enforced by brute force truncation. While this constraint substantially narrows the space of possible solutions, it does not necessarily make the solution unique [34]. The second means of regularization is the use of the nested structure across different scales, governed by the physical model in the continuous form. A solution to the system at a fine scale is disqualified if it does not yield a solution at a coarse scale. Thus, in contrast to GPSR and TwIST which minimize a scalar valued objective function, NeAREst minimizes a vector valued objective function because the resulting estimate for the datacube has to satisfy the objective function at all the scales.

3.4.4 Exploiting the nested structure of NeAREst

NeAREst reduces the total number of iterations needed to reconstruct a datacube through two algorithmic schemes. The first is a relaying scheme that passes the datacube solution at a coarse scale to the next finer scale as the initial guess. The second scheme is an acceleration scheme that utilizes an efficient approximation to the gradient of the KL divergence within a few iterations being performed at a given scale.

The computational cost of each iteration of NeAREst is smaller than that of the GPSR and TwIST methods discussed earlier. With any method being used to process CASSI measurements, the computational cost per iteration is at least as much as that for the matrix-vector multiplication with the sensing matrix. NeAREst has the advantage that the size of the sensing matrix, Φ_s in equation (3.11), is smaller at a coarser scale, lowering the computational cost per iteration at coarser levels.

In comparison, the GPSR and TwIST methods access the full sensing matrix for the matrix-vector multiplication step at every iteration. The GPSR method also

requires the intermediate representation of the datacube in an alternate basis. The computational cost of an iteration increases depending on the number of levels of wavelet decomposition used, as well as the choice of wavelet basis used (e.g. redundant versus non-redundant decomposition).

3.5 CASSI spectral imaging using NeAREst

3.5.1 Spectral imaging of a static scene

The NeAREst method was used to reconstruct a datacube of the static scene first presented in section 3.3.1. Like figure 3.18, figure 3.20 demonstrates the agreement of the spectral signatures of each of the four objects in the scene obtained using NeAREst (magenta dots) and a non-imaging reference spectrometer (blue dots). Once again, the differences in the system responses of CASSI and the reference spectrometer were not accounted for in this comparison.

Compared to TwIST, the spectral image estimation process with NeAREst requires substantially fewer iterations, with each iteration requiring substantially fewer arithmetic operations. No regularization parameter is required. Figure 3.21(a) shows a color restoration of the TwIST estimate of the scene. Figure 3.21(b) shows a color restoration of the scene estimate by NeAREst. In general, the NeAREst restoration is better at revealing the boundaries, shadows and variations in the reflectivity of the objects in the scene. The spectral speckles, among other artifacts, are mostly due to the distribution of the features on the aperture code. The aperture code has relatively long rows or columns of white (open) or black (closed) areas, as shown in figure 3.4. In these areas, an object consisting of multiple spectral components is not well modulated. The red apple looks sharp because its spectral support is relatively narrow.

Thus, comparing the performance of TwIST and NeAREst has provided some insight into aperture code design for future improvement. In general, the distribution

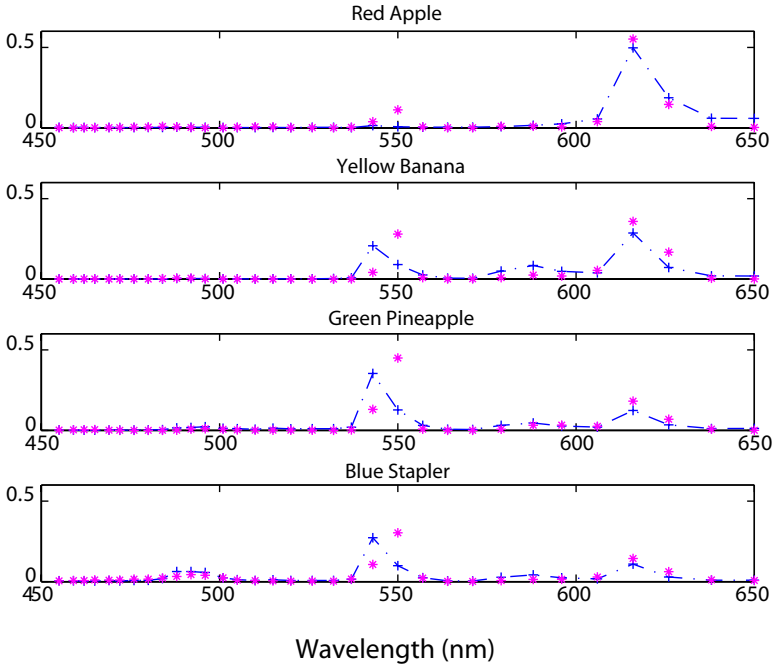

FIGURE 3.20: Spectra of the bodies of the four objects binned into 33 channels. The purple dots are the CASSI spectral signatures at points on the objects after datacube recovery using the NeAREst method. The dashed blue curves are the spectral signatures measured by the reference spectrometer after being integrated into 33 spectral channels with equivalent bandwidths to the CASSI spectral channels.

FIGURE 3.21: (a) Color restoration of the scene estimate with TwIST, (b) Color restoration of the scene estimate with NeAREst.

of the open and closed features on the aperture code should be as uniformly random as possible. Thus, multiple element rows, columns or blocks of these features should

be avoided.

It should be noted that an estimate of the datacube generated by NeAREst was provided as an initial estimate of the datacube to the TwIST algorithm for denoising. However, the details became blurred and disappeared within 3 iterations, mostly due to the penalty applied by TwIST on the regional gradient variation.

3.5.2 Spectral imaging of a dynamically changing scene

A spectral imager with a snapshot capability has a distinct advantage in spectral imaging of fast-changing phenomena. In such situations, temporal scanning to capture a datacube may include artifacts in the datacube due to the motion or spectral variation over time of objects in the scene.

CASSI was used to capture a spectral video of birthday candles at video rate (30 frames per second). Five candles with colored flames were arranged on a table, against a black backdrop, as shown in figure 3.22(a). The experiment was performed under fluorescent lights in the lab. 300 raw CASSI frames were captured over 10 seconds at 30 frames per second. Note that CASSI could also use a CCD with a faster frame rate as long as there was enough light collected in each frame. The CASSI detector array measurement included 256×280 pixels, as shown in figure 3.22(b). It is easy to observe the spatio-spectral overlap of the aperture code modulated images of the four objects due to the dispersion.

'Media1.mov' is a video of the birthday candles with their live flames as viewed by a Canon SD300 digital camera and CASSI. Note that the flame on the blue candle was often bright enough to saturate the 8-bit CASSI CCD. This would result in errors in the spectral reconstruction of that flame as the various spectral channels are multiplexed on top of each other. Also note that between frames $128 - 141$, the blue and purple flames briefly combine, during which time the left edge of the purple flame also saturates the CCD.

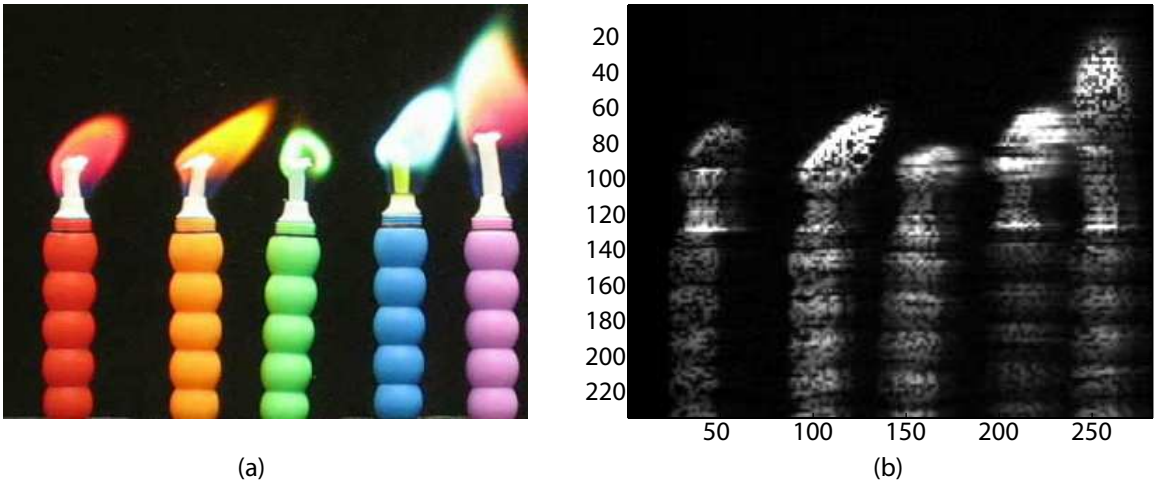

FIGURE 3.22: Single frame excerpts of birthday candles with colored flames as viewed by (a) a Canon SD300 digital camera, and (b) from the CASSI CCD detector array. See 'Media1.mov' for a video of all the frames.

The use of the NeAREst method described in section 3.4 enabled frame-by-frame reconstruction of a spectral video of this scene (see 'Media2.mov'). A datacube for each of the 300 video frames was reconstructed with the same set of sensing matrices at two different scales, without the need to specify a frame-dependent regularization parameter. At the finest scale, CASSI could measure 33 spectral channels, each separated by a column of detector pixels. Two other discrete systems were built at a coarser scale, each with 17 spectral channels. A combination of the solutions at the coarse level, based on a quadrature rule, was used as the initial guess at the fine scale. The reconstruction was carried out in single-threaded MATLAB on a SunFire X4100 M2 workstation with a Dual-Core AMD Opteron 2220 running at 2.8 GHz. The processing of each frame started fresh, and took approximately 45 seconds per frame. It should be noted that the temporal continuity of the live candles may also have been exploited for faster reconstructions of the video frames by using the datacube solution to the previous frame as the initial guess of the current one.

Figure 3.23 shows an example of a datacube estimate for frame 107. It shows

the spatial content of 33 spectral channels between 455 nm and 650 nm. Note that the spatial modulation by the aperture code on each object visible in figure 3.22(b) has been effectively removed from all the spectral channels. As expected, frames $127-141$ in 'Media2.mov' show the blue flame and the left edge of the purple candle flame appearing in virtually all the spectral channels. as these flames were burning brightly enough to saturate pixels on the CASSI detector array. Figure 3.24 shows enlarged versions of spectral channels 11, 22, 27, and 30 to demonstrate the spatial detail that the NeAREst algorithm is able to recover.

To validate the datacube estimate generated with CASSI data, spectra obtained with a non-imaging spectrometer (Ocean Optics USB2000) was used as a quantitative reference. The tip of the fiber connected to the spectrometer was brought very close to the body of each candle to obtain a point-wise spectral signature. Figure 3.25 shows the spectra obtained from the reference spectrometer.

These spectral signatures were then integrated into the 33 spectral channels used for our numerical estimation of spectral images from CASSI. The plots in figure 3.26 demonstrate the agreement between the reference spectral signatures (in blue dots) and the spectral signature of the candle bodies in the computed estimates using NeAREst (in magenta dots). The spectra are normalized by the total intensity per object. Note that differences in the system responses of CASSI and the reference spectrometer have not been accounted for in this comparison. The agreement between the spectra of the candle bodies obtained from the CASSI and the reference spectrometer suggests that CASSI may have also measured the spectra of the candle flames correctly. However, it was not possible to measure reference spectra for the candle flames as they were changing dynamically.

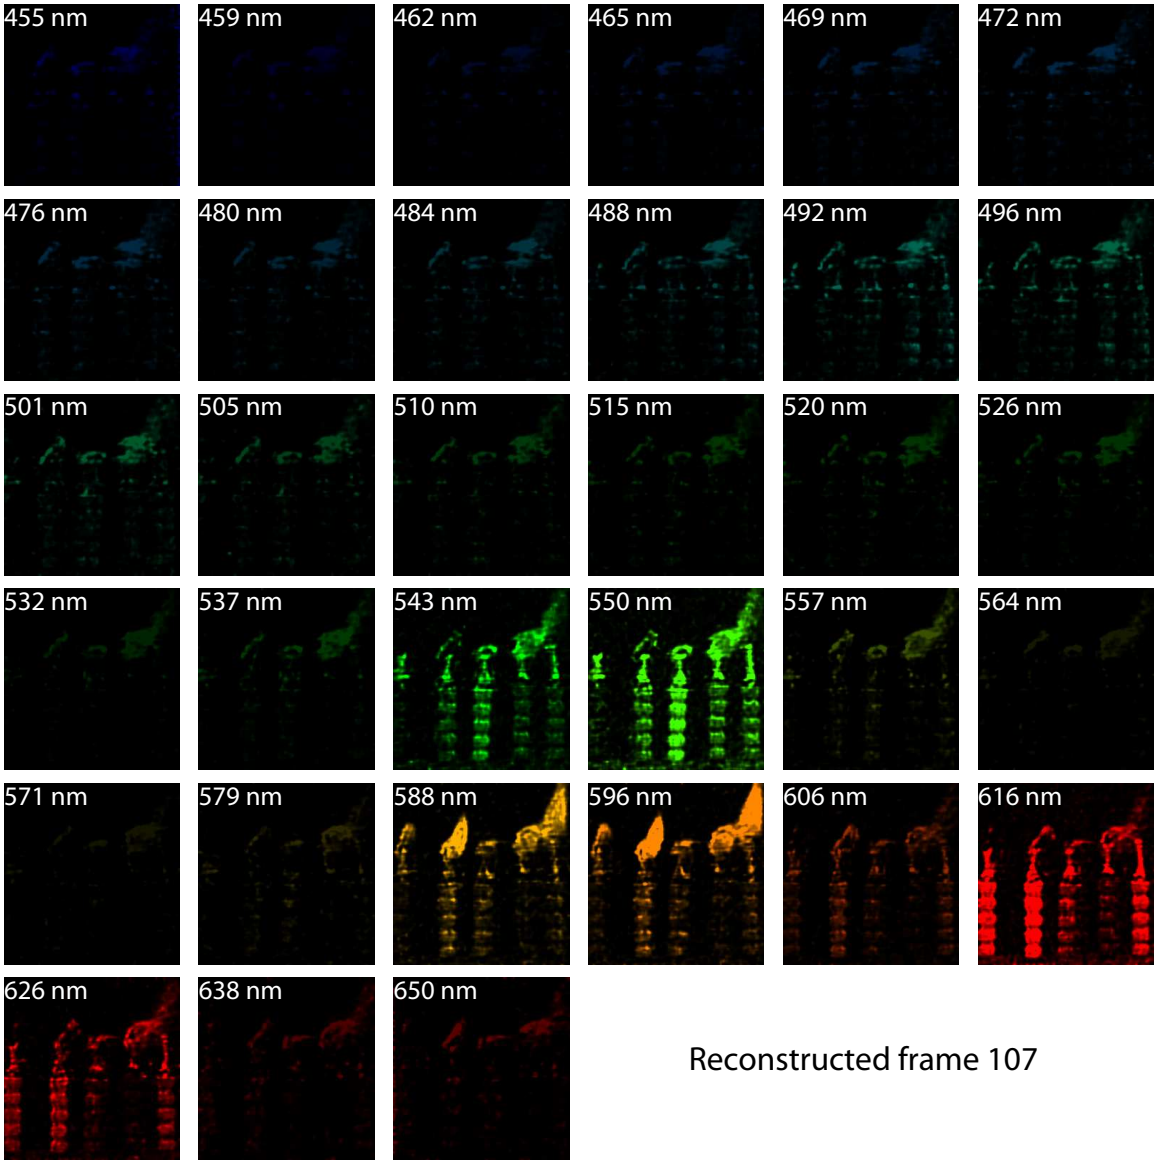

FIGURE 3.23: Spectral image estimate of CASSI frame 107 generated using the NeAREst algorithm. The spatial content in each of 33 spectral channels between 455 and 650 nm is shown. See 'Media2.mov' for a video of all the frames.

3.6 Conclusions

In this chapter, a second generation CASSI prototype was presented that was a significant improvement over the first prototype discussed in chapter 2. The new prototype uses a double Amici prism to introduce spectral dispersion while maintaining a direct

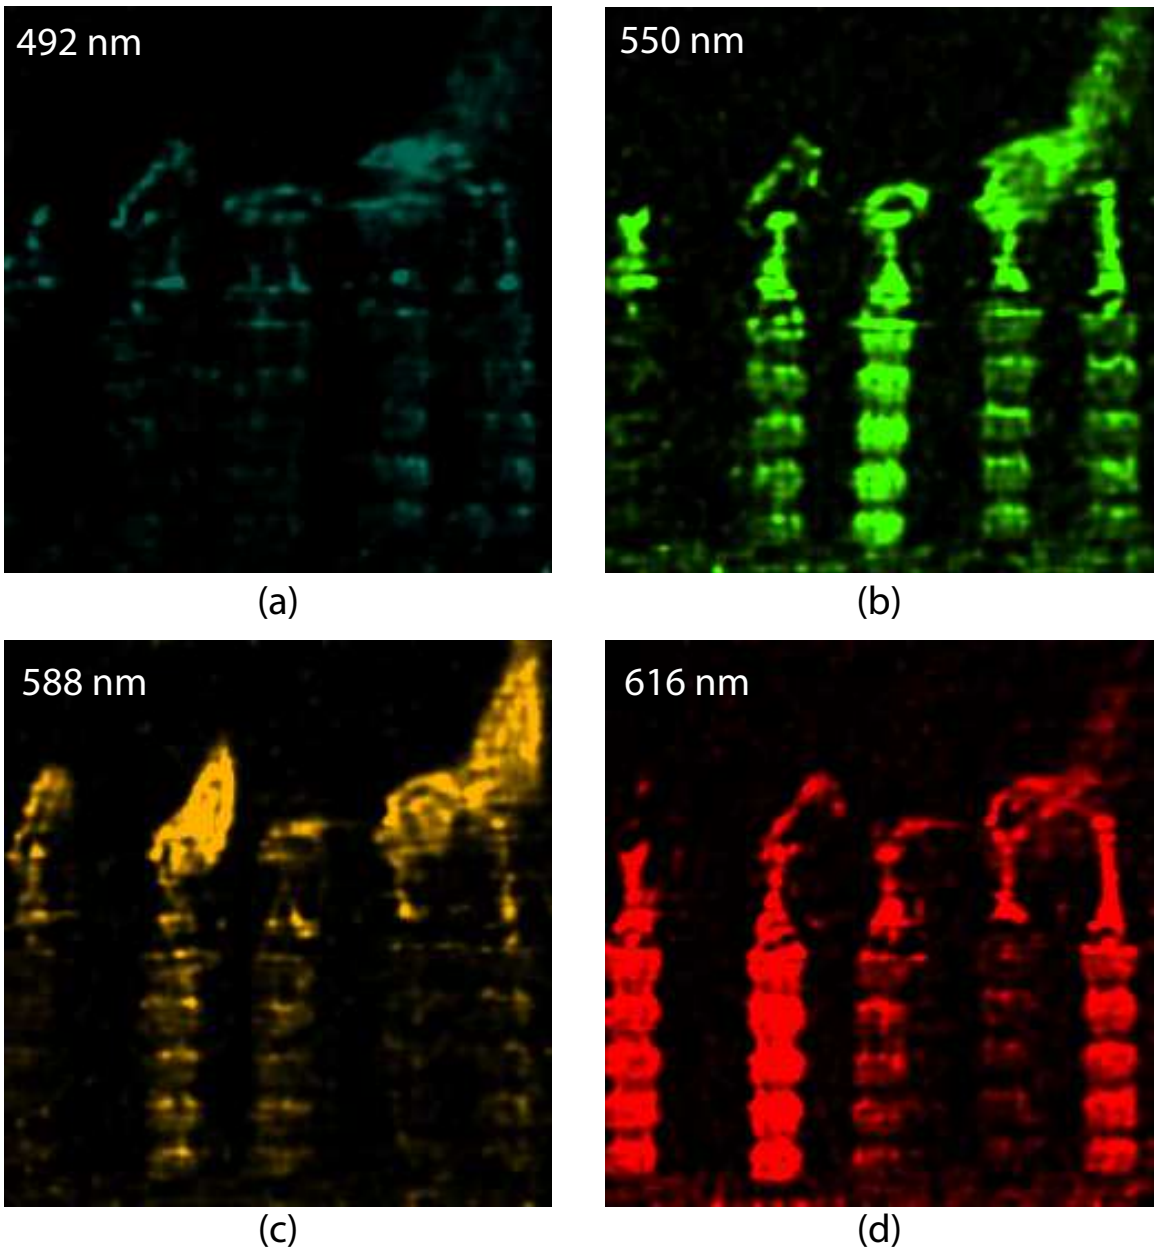

FIGURE 3.24: Enlarged images of the birthday candles at spectral channels with center wavelengths of 492 nm, 550 nm, 588 nm and 616 nm.

view configuration of the overall system. The angles and glass types of the double Amici prism were chosen through optimization of the design of the prototype from the coded aperture plane to the CCD using Zemax ray tracing software.

The TwIST and NeAREst methods were introduced as alternatives to the GPSR

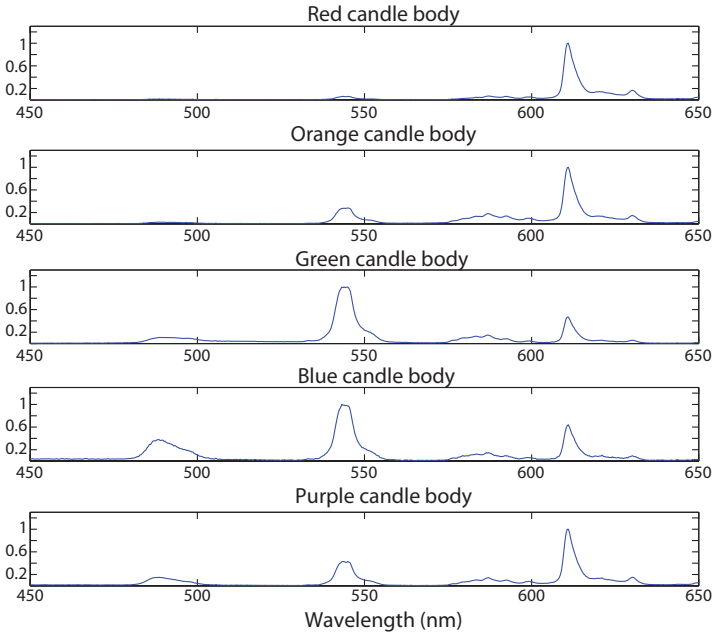

FIGURE 3.25: Spectra of the bodies of the five candles at equally spaced 1 nm wavelengths, as measured by a non-imaging, reference spectrometer.

method to reconstruct a datacube from a snapshot CASSI measurement. In contrast to the GPSR and TwIST methods, the NeAREst method did not require the specification of a regularization parameter for reconstruction. This made NeAREst the algorithm of choice to process raw frames from CASSI to generate a spectral image video of a dynamically changing scene captured at video rate (30 frames per second). A rigorous calibration process was described that facilitated the construction of CASSI system-specific sensing matrices, which improved the quality of a reconstructed datacube.

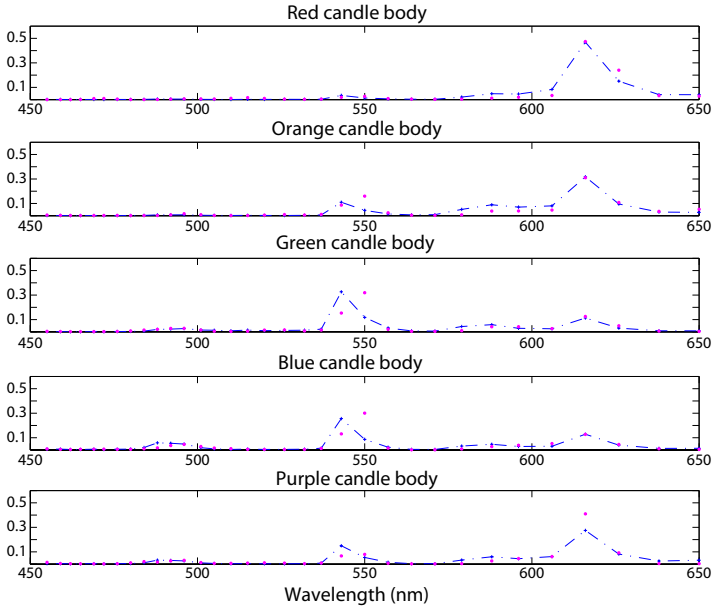

FIGURE 3.26: Spectra of the bodies of the five candles binned into 33 channels. The purple dots are the CASSI spectral signatures at points on the candles after datacube recovery using the NeAREst method. The dashed blue curves are the spectral signatures measured by the reference spectrometer after being integrated into 33 spectral channels with equivalent bandwidths to the CASSI spectral channels.

4

Compressive sampling with CASSI

Although the results presented in the previous chapters support CASSI's potential for its use for snapshot spectral imaging, it is natural to wonder about the constraints imposed on CASSI by its snapshot measurement process. Recall that the scenes imaged by CASSI in the previous two chapters were sparse or compressible on some basis, both spatially as well as spectrally. The theorems of compressive sampling place bounds on the number of sparse basis coefficients in the 3D datacube that can be accurately estimated from the considerably lower dimensional CASSI measurements. The objective of this chapter is to study CASSI from the perspective of compressive sampling theory and to provide insights into potential improvements to the design.

The chapter begins with a background on compressive sampling in section 4.1. In section 4.2, CASSI is presented as an instrument to physically implement a compressively sampled measurement of the datacube. The performance of CASSI as a compressive sampling instrument is then investigated through Matlab simulations in section 4.3. With insights gained from the previous sections, an alternative CASSI design with better performance in noiseless simulations is proposed in section 4.4.

The noise sensitivity of CASSI and its alternative is compared in section 4.5. Finally, the implications of the results for a CASSI using a dual disperser architecture are discussed in section 4.6.

4.1 Background on compressive sampling

Upon arriving at equation (2.12), CASSI was noted to be an example of a system that forms an under-determined system of equations where the total number of detector measurements, m, is significantly lower than the number of unknown voxels, N, in the datacube to be estimated (i.e. $g = \Phi f$, $\Phi \in \mathbb{R}^{M \times N}$ with $M << N$), and as such is an ill-posed problem. If direct, linear inversion is attempted, this problem has an infinite number of solutions since the associated null space is non-trivial.

However, the theory of compressive sampling (CS) [36] suggests that an under-determined problem of the form in (2.12) is well posed for inversion if two practically feasible assumptions are satisfied. These include, (1) the signal of interest (the datacube) is sparse or compressible in some orthonormal basis, Ψ and, (2) the measurement system (CASSI) is designed so that the linear projection implemented by the sensing matrix, Φ, does not significantly damage the salient information in any sparse or compressible signal through the dimensionality reduction.

In general, for S-sparse signals, a stable sensing matrix is one that allows the system matrix $H = \Phi\Psi$ to satisfy the Restricted Isometry Property (RIP). Mathematically, a matrix H is said to satisfy the RIP of order S if there is a restricted isometry constant $\delta_S \in (0, 1)$ such that for any S-sparse vector b

$$(1 - \delta_S)\|b_\Omega\|_2^2 \leq \|H_\Omega b_\Omega\|_2^2 \leq (1 + \delta_S)\|b_\Omega\|_2^2. \tag{4.1}$$

This amounts to the matrix H_Ω preserving the lengths/structure of S-sparse vectors as tightly as possible, where Ω represents the set of indices on which the S-sparse signal is supported and H_Ω is a sub-matrix of H obtained by selecting the columns

of H corresponding to the indices in Ω [36]. This is equivalent to stating that all subsets of S columns taken from H are nearly orthogonal (they cannot be exactly orthogonal, since there are more rows than columns).

From the perspective of CS, a "good" system matrix should satisfy equation (4.1) for the largest possible S that is as close to M as possible. Certain classes of sensing matrices including those whose entries are independent and identically distributed realizations of a Gaussian or Bernoulli random variable are known to produce system matrices that satisfy the RIP, provided that the number of measurements made, M, is large enough relative to the sparsity S [37]. In fact, a Gaussian sensing matrix is universal in the sense that the resulting system matrix $H = \Phi\Psi$ will satisfy the RIP with high probability regardless of the choice of orthonormal basis Ψ [38]. Unfortunately, it is not possible to physically realize a sensing matrix with random, negative values while maintaining a snapshot measurement process.

When implementing a snapshot CASSI measurement, the CASSI sensing matrix, Φ, is an example of a matrix that is constructed by choosing uniformly at random M rows of an orthonormal basis, and produces *incoherent measurements*. With such a sensing matrix, it has been conjectured that the resulting system matrix $H = \Phi\Psi$ satisfies the RIP in equation (4.1) with overwhelming probability if the number of measurements, M, satisfies the sufficient condition that

$$M \geq C \cdot \mu^2(\Phi, \Psi) \cdot S \cdot logN, \tag{4.2}$$

for some constant C and where $\mu(\Phi, \Psi)$ is the mutual coherence defined as

$$\mu(\Phi, \Psi) = \sqrt{N} \max_{\substack{1 \leq i \leq M \\ 1 \leq j \leq N}} |\langle \phi_i, \psi_j \rangle| \,. \tag{4.3}$$

Here ϕ_i and ψ_j represent the i^{th} row of Φ whose rows are orthonormal, and the j^{th} column of Ψ, respectively [39]. This quantifies the maximum correlation between Φ

and Ψ. An alternative definition of mutual coherence is given by

$$\mu(H) = \max_{i \neq j, 1 \leq i,j \leq N} \left\{ \frac{h_i^T h_j}{\|h_i\|_2 \|h_j\|_2} \right\}, \tag{4.4}$$

where h_i is the i^{th} column of $H = \Phi\Psi$ [40]. This can be interpreted as the maximum off-diagonal element of the Gram matrix of H whose columns are normalized. This definition quantifies the correlations (coherences) between the columns of H based on the angles between the columns.

Given that the number of detector measurements is large enough relative to the sparsity S and satisfies equation (4.2), the system matrix H satisfies equation (4.1) with overwhelming probability and a highly accurate estimate of the datacube as a solution to the measurement equation $g = Hb$ may be found by searching for

$$\hat{b} = \arg\min_b ||b||_0 \text{ such that } g = \Phi f = \Phi\Psi b = Hb, \tag{4.5}$$

which amounts to finding a solution with the lowest total number of non-zero sparse coefficients. Unfortunately, searching for such a solution is numerically unstable and NP-complete, requiring an exhaustive search of all $\binom{N}{s}$ possible locations of the nonzero entries in b [38]. The computational intractability of this combinatorial search is bypassed [41] with a search for

$$\hat{b} = \arg\min_b ||b||_1 \text{ such that } g = \Phi f = \Phi\Psi b = Hb, \tag{4.6}$$

where $||b||_1 = \sum_{i=1}^{S} |b_i|$. When no detector noise is present, $\hat{b}$ matches b exactly. In the presence of detector noise n, the measurement equation is modified to $g = Hb + n$. Given $||n||_2 \leq \epsilon$, the problem becomes a search for [39]

$$\hat{b} = \arg\min_b ||b||_1 \text{ such that } ||g - Hb||_2 \leq \epsilon. \tag{4.7}$$

This problem and equation (4.6) are both convex and may be solved using non-linear, iterative convex optimization techniques [42]. The solution to equation (4.7) is optimal because exactly sparse signals are recovered exactly. In addition, the theory shows that strong stability exists for compressible signals in the presence of noise [43]. This is of great practical value in processing natural signals measured in the real world where, more often than not, signals tend to be compressible rather than sparse and are measured in the presence of noise.

The GPSR method discussed in chapter 2 to process CASSI measurements essentially rephrases the problem in equation (4.7) as an unconstrained optimization problem. The TwIST method discussed in chapter 3 can also be used to solve a similar unconstrained optimization problem, except with a TV regularizer. A theoretical proof of exact reconstruction for TV regularization currently only holds in the case of Fourier measurements of signals with exactly sparse gradients in the absence of noise [44]. No such proof of stability exists in the presence of noise and non-exact sparsity of the gradients.

4.2 CASSI as a compressive sampling instrument

While the previous chapters provided empirical evidence for CASSI's ability to recover a datacube from a snapshot measurement, they left open questions about the optimality of the system design and the limits of its performance. In this section, CASSI is shown to be a means of physically implementing a compressively sampled measurement. A discretization framework is outlined for Matlab simulations to be discussed in the next section, the mathematical structure of the CASSI sensing matrix Φ is presented, the choices of sparsity bases for CASSI are analyzed and the RIP is established for a system matrix that is generated using Φ and the choice of the sparsifying basis, Ψ. Although the previous chapters established the relevance of incorporating optical distortions in numerical processing of experimental data, much

can be learned by studying a hypothetical CASSI without optical distortions and its associated sensing matrix, Φ.

4.2.1 Discretization framework

Let the discrete datacube, f, to be recovered from CASSI measurements be composed of N_λ spectral slices where each spectral slice has $N_x \times N_y$ pixels so that the total number of voxels in the datacube is $N_f = N_x \times N_y \times N_\lambda$. Assuming that a dispersive element introduces a 1 pixel shift for each spectral channel, the 2D detector array makes $M = N_x \times (N_y + N_\lambda - 1)$ measurements. Assume that the size of each pixel in a spectral slice and the size of each feature on the coded aperture are identical to the detector pixel size.

Now, let f be rearranged such that $f = [f_1; f_2; f_3; \cdots; f_{N_\lambda}]$, where the semicolon denotes vertical concatenation of column vectors such that f is a column vector of length $N_f = N_x \times N_y \times N_\lambda$. The $N_x \times N_y$ pixels in each $f_i \in \mathbb{R}^{N_x \times N_y}$ are ordered by vectorizing the i^{th} spectral slice column-wise. Given this discretization and ordering scheme, the CASSI sensing matrix has the form $\Phi = [\Phi_1 \Phi_2 \cdots \Phi_{N_\lambda}] \in \mathbb{R}^{M \times N_f}$. In $\Phi_1 \in \mathbb{R}^{M \times (N_x \times N_y)}$ the bottom $(N_x \times N_y)$ rows form a diagonal sub-matrix obtained by putting along the diagonal a vector that is constructed by vectorizing elements of the coded aperture column-wise. $\Phi_i \in \mathbb{R}^{M \times (N_x \times N_y)}$ are obtained by shifting up the diagonal sub-matrix in Φ_1 by $(i-1) \times N_y$ rows. Figure 3.12(a) showed an example of Φ with $N_x = N_y = 8$ and $N_\lambda = 4$. If the total number of 1's in Φ_i is defined as M_i, then the transmission rate for the i^{th} spectral channel is defined as $T_i = M_i/(N_x \times N_y) \times 100$. In addition, since the same aperture code modulates all the spectral channels, $T_i = T$, which can be modified by changing the overall transmissivity of the coded aperture.

The size and structure of the CASSI system matrix, H, depends on the choice of the sparsifying basis Ψ. Datacube reconstruction using the GPSR method in section

2.7 and the TwIST method in section 3.3 is performed by treating each spectral channel independently and representing each spectral channel in a 2D sparse basis separately. This is based on the reasonable assumption that the correlations between spectral channels is not significantly large since the dispersion introduced by the prism is not large. The sparsity of the whole datacube is then defined as the sum of the sparsities in all the spectral channels. Given this 2D reconstruction scheme, the CASSI system matrix is defined as

$$\begin{aligned} H &= \Phi\Psi \\ &= [\Phi_1\Phi_2\cdots\Phi_{N_\lambda}]\ \text{bldiag}(\Psi_{2D},\Psi_{2D},\cdots,\Psi_{2D}) \\ &= [\Phi_1\Psi_{2D}\ \Phi_2\Psi_{2D}\ \cdots\ \Phi_{N_\lambda}\Psi_{2D}] \\ &= [H_1H_2\cdots H_{N_\lambda}], \end{aligned} \tag{4.8}$$

where Ψ_{2D} denotes the basis of choice for sparsifying a spectral channel.

4.2.2 Choice of basis

Examples of choices of Ψ_{2D} include the discrete cosine transform (DCT) and the wavelet transform, in which case $b_i \in \mathbb{R}^{N_x\times N_y}$ has the same dimensions as $f_i \in \mathbb{R}^{N_x\times N_y}$ and the system matrix $H \in \mathbb{R}^{M\times N_f}$. Alternatively, Ψ_{2D} may be a redundant dictionary like the redundant wavelet transform, in which case $b_i \in \mathbb{R}^N$ where $N > N_x \times N_y$ and $H \in \mathbb{R}^{M\times N}$. Although the TV basis used during reconstruction with TwIST also treats each spectral channel independently, it is an implicitly sparse basis because the inverse TV operator is not explicitly used for reconstruction.

For simplicity of analysis, assume that a 2D non-redundant wavelet basis is used for the 2D sparse basis Ψ_{2D}. The diagonal sub-matrix of each Φ_i has diagonal elements that take on 0 or 1 randomly. Hence, Φ_i samples a uniformly random set of pixels in f_i and equivalently selects a random set of rows from the orthonormal wavelet basis Ψ_{2D}. The selected random set of basis vectors are the measurement

basis vectors for the 2D sparse representation b_i of the i^{th} spectral slice f_i. Let S_i be the sparsity of the i^{th} spectral slice such that $S = \sum_{i=1}^{N_\lambda} S_i$. Based on the discussion presented in section 4.1, a sufficient condition for accurate reconstruction of the i^{th} spectral slice is $M_i \geq C_i \cdot \mu^2(\Phi_i, \Psi_{2D}) \cdot S_i \cdot (\log(N_x N_y))$, where M_i is the number of 1's in Φ_i.

It is possible to have a large $\mu(\Phi_i, \Psi_{2D})$ if an image contains high frequency features because the wavelet representation of the image would involve a wavelet function that has a high correlation with the Kronecker delta basis. Thus, such a situation would require a large M_i as a sufficient condition. Based on experience gained through numerical simulations, however, such incoherence is overly sufficient, and much smaller M_i can produce reconstructions with high accuracy.

It is also possible to achieve the best possible incoherence $\mu(\Phi_i, \Psi_{2D}) = 1$ if the 2D DCT basis is used for Ψ_{2D}. Thus, using the the 2D DCT basis instead of the 2D non-redundant wavelet basis implies that a smaller number of measurements M_i are required for accurate reconstruction of the i^{th} spectral slice. However, choosing the 2D DCT basis may increase the number of nonzero coefficients S_i. This implies that the sparsity is degraded, which consequently requires a larger number of measurements. Based on experience, it is difficult to exactly predict which choice of basis generally produces better reconstructions using the upper bounds because the bounds are sufficient conditions but not necessary. Although the M_i predicted by the upper bounds is smaller using DCT basis than if the wavelet basis is used, the wavelet basis generally produces better results with approximately equal number of measurements.

These examples demonstrate that the choice of basis for CASSI sensing matrices is less general than the choice of basis for Gaussian or Bernoulli random sensing matrices that are known to be universal in the sense that they do not change the

incoherence significantly when combined with various bases.

4.2.3 Restricted Isometry Property for the CASSI system matrix

To maximize the reconstruction fidelity of datacubes for a variety of scenes with varying amounts of sparsity, the CASSI system matrix, H, should satisfy the RIP in equation (4.1) for the largest possible S. Given that datacube reconstruction using GPSR and TwIST has previously been performed by treating each spectral channel independently and representing each channel in a 2D sparse basis separately, it is important to study how the CASSI multiplexing process influences the restricted isometry property of the CASSI system matrix, H.

The CASSI measurement system equation without noise can be written as

$$\begin{aligned} g = \Phi f &= H\, b \\ &= \sum_{i=1}^{N_\lambda} H_i b_i = \sum_{i=1}^{N_\lambda} \Phi_i f_i \end{aligned} \tag{4.9}$$

using equation (4.8). Then, assuming that M_i is large enough, which may be obtained by controlling the transmission rate of the aperture code T, the RIP can be established for the sub-matrix of H corresponding to the i^{th} spectral channel. In particular, there exists a restricted isometry constant δ_{S_i} such that

$$\begin{aligned} (1-\delta_{S_i})\|b_{\Omega_i}\|_2^2 &\leq \|H_{\Omega_i} b_{\Omega_i}\|_2^2 \\ &= \|(\Phi_i \Psi_{2D})_{\Omega_i} b_{\Omega_i}\|_2^2 \leq (1+\delta_{S_i})\|b_{\Omega_i}\|_2^2 \end{aligned} \tag{4.10}$$

for any S_i-sparse signals in the the i^{th} spectral channel and Ω_i denotes the set of indices on which the S_i-sparse signal is supported.

Assuming that CASSI is designed such that each H_i satisfies the S_i-RIP, the

multiplexing of all the spectral channels influences the RIP for H. In particular,

$$\begin{aligned}
\|H_\Omega b_\Omega\|_2^2 &= b_\Omega^T H_\Omega^T H_\Omega b_\Omega = f_\Omega^T \Phi_\Omega^T \Phi_\Omega f_\Omega \\
&= \sum_{i=1}^{N_\lambda}\sum_{j=1}^{N_\lambda} f_{\Omega_i}^T \Phi_{\Omega_i}^T \Phi_{\Omega_j} f_{\Omega_j} \\
&= \sum_{i=1}^{N_\lambda} f_{\Omega_i}^T \Phi_{\Omega_i}^T \Phi_{\Omega_i} f_{\Omega_i} + \sum_{i\neq j}\sum f_{\Omega_i}^T \Phi_{\Omega_i}^T \Phi_{\Omega_j} f_{\Omega_j} \\
&= \sum_{i=1}^{N_\lambda} b_{\Omega_i}^T H_{\Omega_i}^T H_{\Omega_i} b_{\Omega_i} + \sum_{i\neq j}\sum f_{\Omega_i}^T \Phi_{\Omega_i}^T \Phi_{\Omega_j} f_{\Omega_j} \\
&= \sum_{i=1}^{N_\lambda} \|H_{\Omega_i} b_{\Omega_i}\|_2^2 + \sum_{i\neq j}\sum f_{\Omega_i}^T \Phi_{\Omega_i}^T \Phi_{\Omega_j} f_{\Omega_j}, \qquad (4.11)
\end{aligned}$$

where Ω denotes the set of indices on which the S-sparse signal b is supported. Using equation (4.10),

$$\begin{aligned}
&\sum_{i=1}^{N_\lambda} \|H_{\Omega_i} b_{\Omega_i}\|_2^2 \geq \sum_{i=1}^{N_\lambda} (1-\delta_{S_i}) \|b_{\Omega_i}\|_2^2 \\
&\geq (1 - \max_i \delta_{S_i}) \sum_{i=1}^{N_\lambda} \|b_{\Omega_i}\|_2^2 \geq (1-\delta) \|b_\Omega\|_2^2 , \qquad (4.12)
\end{aligned}$$

where the restricted isometry constant $\delta = \max_i \delta_{S_i}$. Similarly, using equation (4.10) produces

$$\begin{aligned}
&\sum_{i=1}^{N_\lambda} \|H_{\Omega_i} b_{\Omega_i}\|_2^2 \leq \sum_{i=1}^{N_\lambda} (1+\delta_{S_i}) \|b_{\Omega_i}\|_2^2 \\
&\leq (1 + \max_i \delta_{S_i}) \sum_{i=1}^{N_\lambda} \|b_{\Omega_i}\|_2^2 \leq (1+\delta) \|b_\Omega\|_2^2 . \qquad (4.13)
\end{aligned}$$

Combining equations (4.12) and (4.13) with (4.11) results in

$$(1-\delta)\|b_\Omega\|_2^2 \leq \|H_\Omega b_\Omega\|_2^2 \leq (1+\delta)\|b_\Omega\|_2^2 + \sum_{i\neq j}\sum f_{\Omega_i}^T \Phi_{\Omega_i}^T \Phi_{\Omega_j} f_{\Omega_j}, \qquad (4.14)$$

where the inequality on the left hand side holds because all the elements of Φ_i and f_i for all i are non-negative. Now assuming that there exists a constant $c \ll 1$ and if

$$\sum_{i \neq j} \sum f_{\Omega_i}^T \Phi_{\Omega_i}^T \Phi_{\Omega_j} f_{\Omega_j} = \sum_{i \neq j} \sum b_{\Omega_i}^T H_{\Omega_i}^T H_{\Omega_j} b_{\Omega_j} \leq c \left\| b_\Omega \right\|_2^2 , \tag{4.15}$$

then equation (4.14) becomes

$$(1-\delta) \left\| b_\Omega \right\|_2^2 \leq \left\| H_\Omega b_\Omega \right\|_2^2 \leq (1+\delta) \left\| b_\Omega \right\|_2^2 + c \left\| b_\Omega \right\|_2^2 \leq (1+\delta_c) \left\| b_\Omega \right\|_2^2 . \tag{4.16}$$

By choosing $\delta_{CASSI} = \max(\delta, \delta_c)$, the inequalities can be made symmetric so that

$$(1-\delta_{CASSI}) \left\| b_\Omega \right\|_2^2 \leq \left\| H_\Omega b_\Omega \right\|_2^2 \leq (1+\delta_{CASSI}) \left\| b_\Omega \right\|_2^2 . \tag{4.17}$$

From equations (4.14), (4.16), and (4.17), it is easy to see that the constant δ_{CASSI} depends on the magnitude of the non-negative cross-term $\sum \sum_{i \neq j} f_{\Omega_i}^T \Phi_{\Omega_i}^T \Phi_{\Omega_j} f_{\Omega_j}$. This suggests that the cross-term may lead to degradation in the reconstruction quality of the datacube because $\delta_{CASSI} \geq \delta$, which implies that the dimensionality reduction implemented by the CASSI system matrix, H, is not able to preserve the isometry with S-sparse signals b.

4.2.4 Reconstruction with a 3D sparsifying basis

The analysis above reveals that the existence of the cross-term in equation (4.12) is caused by the use of the 2D sparse representation basis. Now consider the case where a 3D sparse basis is used as Ψ. Recall that the rows of Φ do not overlap because there is at most only one nonzero, which is 1, per column. Thus the rows are orthogonal by construction. A by-product of the orthogonality of the rows of Φ is the orthogonality of the rows of $H = \Phi\Psi_{3D}$

$$h_i^T h_j = \left(\sum_{k \in \mathcal{K}_i} \psi_k \right)^T \left(\sum_{l \in \mathcal{K}_j} \psi_l \right) = 0, \tag{4.18}$$

where $\mathcal{K}_i$ denotes the set of the indices of the columns of Ψ_{3D} that correspond to 1s in the i^{th} row of Φ, and ψ_i denotes the i^{th} column of Ψ_{3D}. Note that since $\mathcal{K}_i \cap \mathcal{K}_j = \emptyset$ by construction of Φ, the inner products between any two different rows are 0 as shown in equation (4.18) and this implies that the cross-term in equation (4.14) is 0.

While the concerns associated with the cross-term vanish with the use of a 3D sparsifying basis like a 3D wavelet transform, the number of nonzero coefficients S may increase and the incoherence, which is defined as $\mu(\Phi_{CASSI}, \Psi_{3D})$ instead of $\mu(\Phi_i, \Psi_{2D})$ may change as well. Experience suggests that it is generally very difficult to strictly separate, quantify and compare the effects of the change of basis from 2D to 3D on the reconstruction fidelity because of the simultaneous changes in the incoherence, sparsity, and the value of the cross-term.

4.3 Matlab simulations investigating CASSI performance

With the discretization framework outlined in the previous section, Matlab simulations were conducted to study the performance of CASSI as a compressive sampling system. In all the simulations, each spectral channel was treated independently and represented in a 2D non-redundant Haar wavelet basis. Given simulated CASSI measurements, the TwIST method was used to recover a datacube, $\hat{f}$ by solving

$$\hat{f}(\tau) = \Psi \left[\underset{b}{\text{argmin}} \left\{ \frac{1}{2} \|g - H\, b\|_2^2 + \tau ||b||_1 \right\} \right]. \tag{4.19}$$

The regularization parameter τ, in every case, was empirically determined so that the peak signal-to-noise ratio (PSNR) would be maximized. The PSNR was defined as

$$PSNR = 10 \cdot \log_{10} \left(\frac{MAX_f^2}{MSE} \right), \tag{4.20}$$

where MAX_f is the maximum value in the true datacube, $f \in \mathbb{R}^{N_x \times N_y \times N_\lambda}$, used for simulation and MSE is the mean squared error between the true datacube and the

reconstructed datacube defined as

$$MSE = \frac{1}{N_x N_y N_\lambda} \sum_{i=0}^{N_x-1} \sum_{j=0}^{N_y-1} \sum_{k=0}^{N_\lambda-1} \left[f(i,j,k) - \hat{f}(i,j,k) \right]^2. \tag{4.21}$$

3 different datacubes along with aperture codes at 7 different overall transmissions were used throughout this simulation study. Each datacube consisted of $128 \times 128 \times 8$ voxels. All three datacubes, as shown in figure 4.1, included a weak background and had roughly 1670 Haar wavelet coefficients each.

6 of the aperture codes codes were random patterns with overall transmissions of 12.5%, 30%, 50%, 70%, 90% and 100%. In addition, a code pattern called the sparse code was designed with 12.5% overall transmission. Unlike its random transmission counterparts, the open features on the sparse code, represented by the 1's in Φ, were positioned so that there would be no multiplexing of the spectral channels as they are imaged onto the detector array. As a result, the cross-term in equation (4.14) becomes 0. Figure 4.2 plots the number of non-zero detector measurements M as a function of T for any of the three datacubes considered.

Through Matlab simulations to be presented in the following subsections, features of CASSI's system design and its performance were investigated. In particular, the simulations provided insight and conclusions about three questions:

1. What is the importance of coding by the optical hardware?

2. Is there an optimal aperture code pattern for CASSI?

3. Is the performance of CASSI scene dependent?

4.3.1 Importance of coding

Coding introduced by the coded aperture and the dispersive element is essential for any hope of recovery of a sparse datacube from the small number of linear multiplexed projections implemented by the CASSI measurement process. The pattern

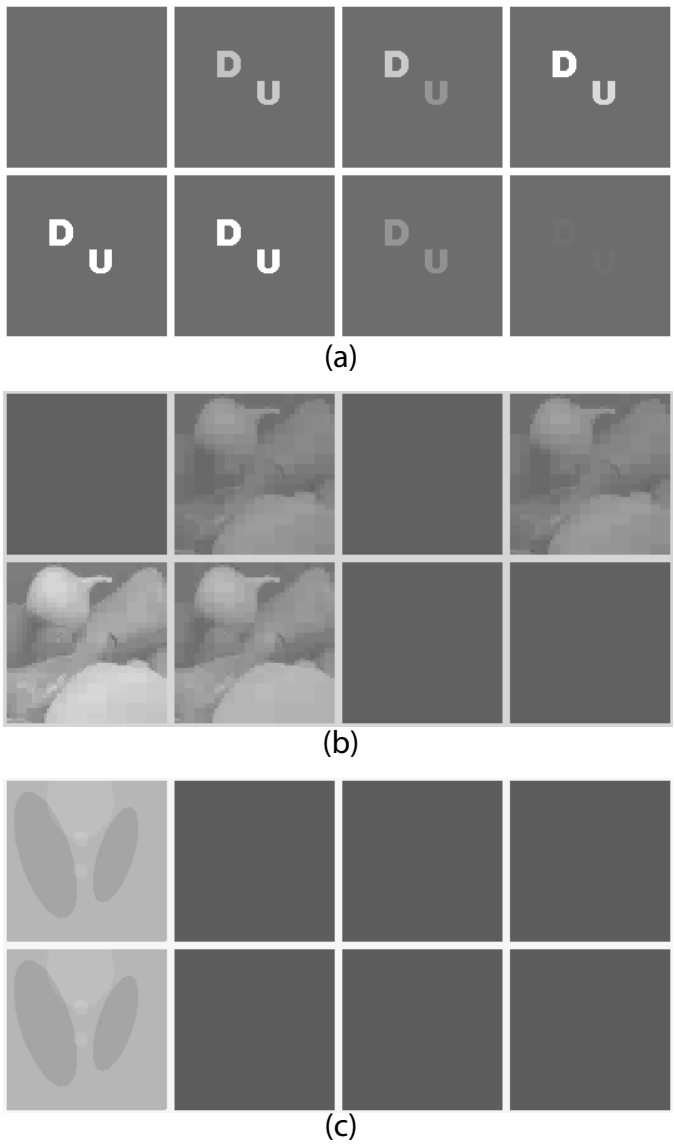

FIGURE 4.1: Three $128 \times 128 \times 8$ datacubes used to study the performance of CASSI through Matlab simulations. (a), (b) and (c) show the 'DU' datacube, 'peppers' datacube, and 'phantom' datacube, respectively. The scenes are assumed to contain a weak background that is generated by adding a small constant value to all the spectral slices. For comparison purposes, the datacubes are synthesized such that they produce roughly equal number of 2D Haar wavelet coefficients ($S = 1666$ for (a), $S = 1674$ for (b), and $S = 1676$ for (c)).

and overall transmission of the aperture code determine the amount of multiplexing on the detector pixels and also help to determine the spatial resolution of the

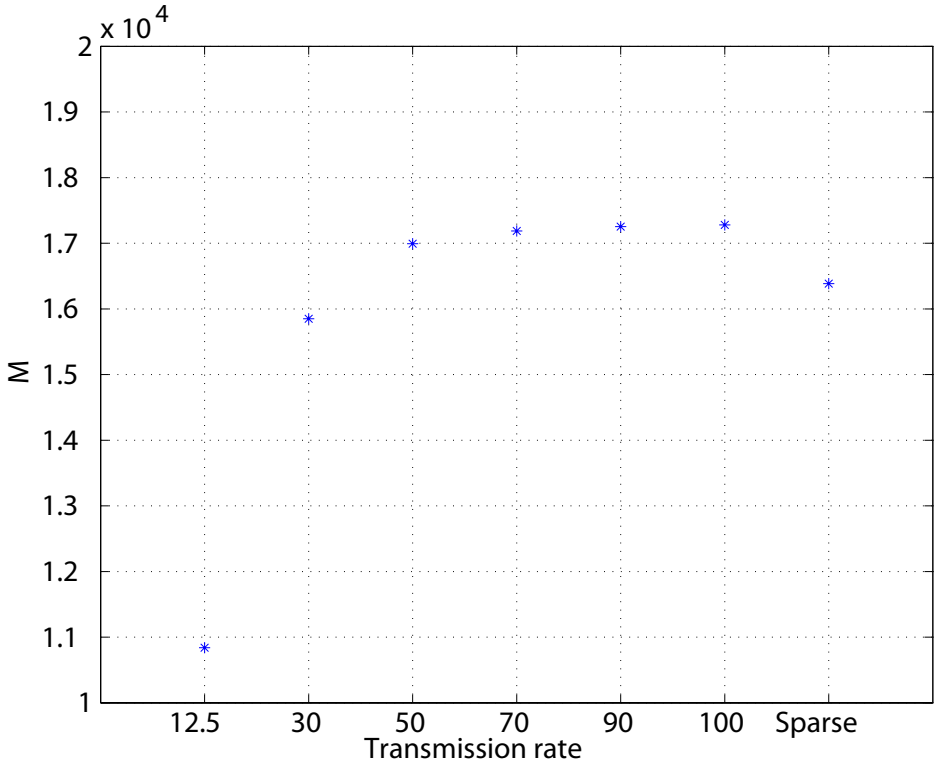

FIGURE 4.2: Number of nonzero CASSI detector measurements as a function of transmission rate T. The sparse code has an overall transmission of 12.5%.

reconstructed datacube. *Both* the spatial modulation introduced by the aperture code and the spectral coding introduced by the dispersive element are necessary for CASSI to have a chance to reconstruct an accurate estimate of a datacube. Without this dispersion, all the spectral channels are modulated by the same aperture code pattern and imaged onto the same detector pixels (see figure 4.3(b)), making it impossible to distinguish them, as shown in figure 4.3(a). The reconstructed datacube has a PSNR of just 13.30. On the other hand, the importance of the aperture code is demonstrated in figures 4.4(a), (c) and (e) by the fact that using an aperture code with 100% transmission rate, i.e. a fully open aperture, results in the worst PSNR for all of the datacubes considered even when dispersion exists.

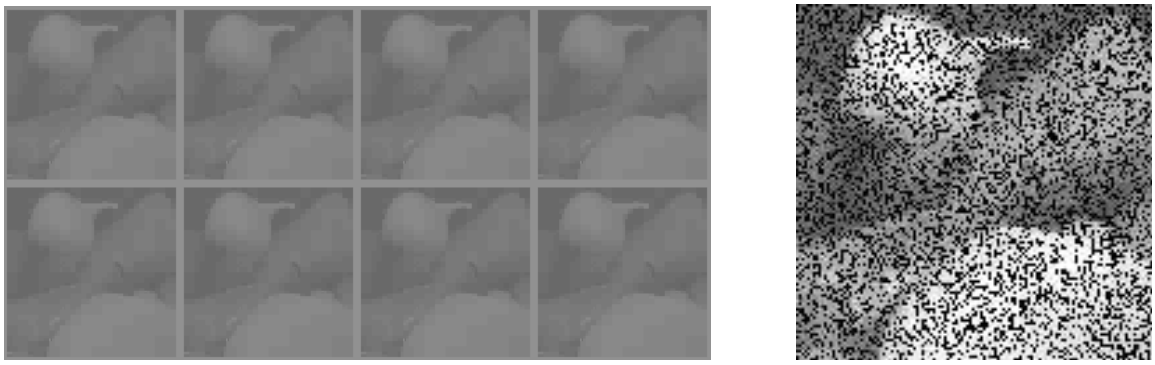

FIGURE 4.3: Reconstructions of the 'peppers' datacube with 70% transmission rate without spectral dispersion (spectral coding): (a) shows the reconstruction (PSNR=13.30), and (b) shows its associated detector measurements whose size is $128 \times (128 + 8 - 1)$.

4.3.2 Optimality of the aperture code

Given that the aperture code is an essential component for datacube recovery from CASSI measurements, it is natural to wonder if there is an optimal transmission rate and code pattern in general. To answer this question, CASSI measurements were simulated for each of the 3 datacubes shown in figure 4.1 at each transmission rate considered. Figure 4.5 plots the PSNR of each reconstructed datacube as a function of the transmission of the aperture code.

The figure illustrates that even though the sparse code does not produce any cross-term, it produces reconstructions with very poor spatial quality because the number of non-zero detector measurements for each spectral channel M_i is low. The spatial quality of reconstructions can be improved by increasing the transmission T of the random aperture code, which effectively increases M_i for each i. Although this also results in an increase in the size of the cross-term, the increased transmission results in more detector measurements and consequently a more spatially accurate estimate of the datacube. However, as illustrated in figure 4.5, increasing the transmission of the aperture code beyond a certain point does not produce any further increase in

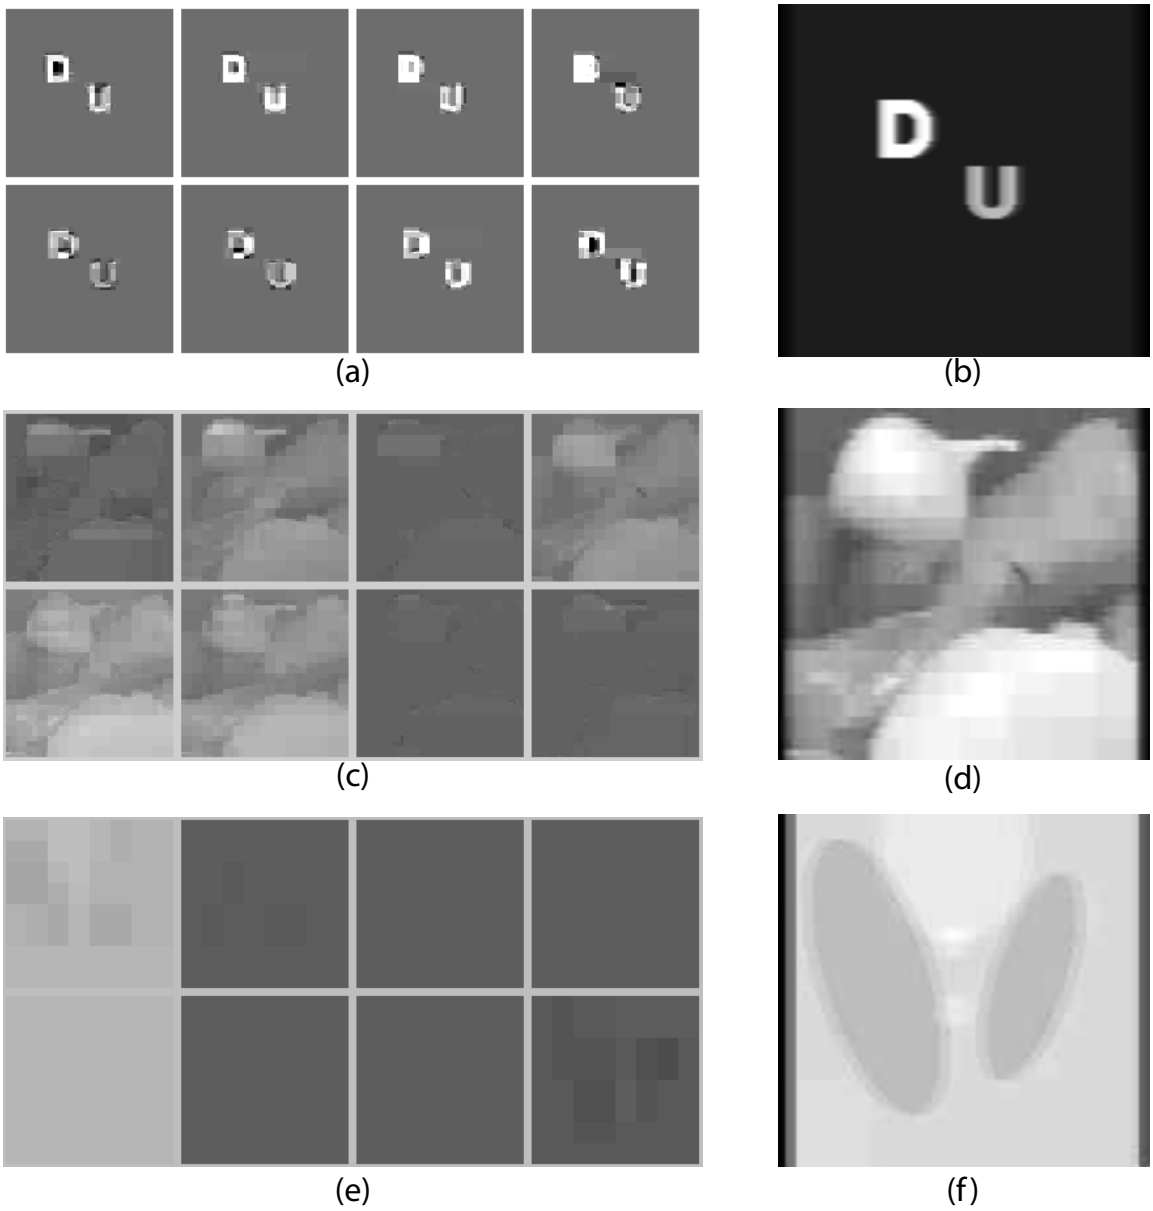

FIGURE 4.4: Reconstructions of the three datacubes in figure 4.1 with 100% transmission rate (no spatial coding by coded aperture): (a) PSNR=22.71, (c) PSNR=24.10 and (e) PSNR=30.91 show the reconstructions of the scenes in figures 4.1(a), (b) and (c) respectively. (b), (d) and (f) show the associated detector measurements.

the PSNR. In fact, the PSNR begins to fall until it deteriorates to its lowest possible value, namely that achieved when a fully open aperture is used. The transmission

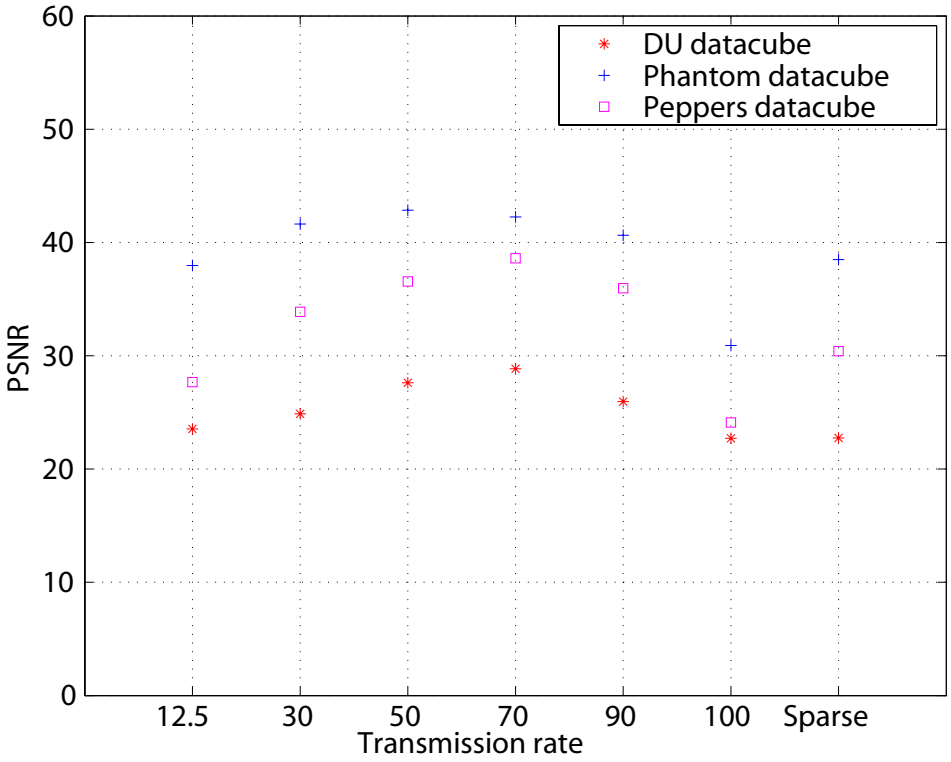

FIGURE 4.5: PSNR of the reconstructed datacubes as a function of the transmission rate of the aperture code.

rate at which the PSNR begins to fall is also scene dependent. Figure 4.5 shows that, at least for the datacubes considered, this is observed to happen between a transmission rate of 50% and 70%.

Reconstructions were also performed using 10 different random aperture patterns for the 'DU', 'phantom' and 'peppers' datacubes with transmission rates of 70%, 50% and 70% respectively. These transmission rates were chosen because they produced the reconstructions with the highest PSNR as shown in figure 4.5. As figure 4.6 shows, there is only marginal variation in the PSNR as a function of the pattern of the aperture code. This implies that as long as the open and closed features of the coded aperture are distributed uniformly in a random manner over the full aperture area, the pattern itself does not significantly matter. This conclusion was also made back in section 3.5.

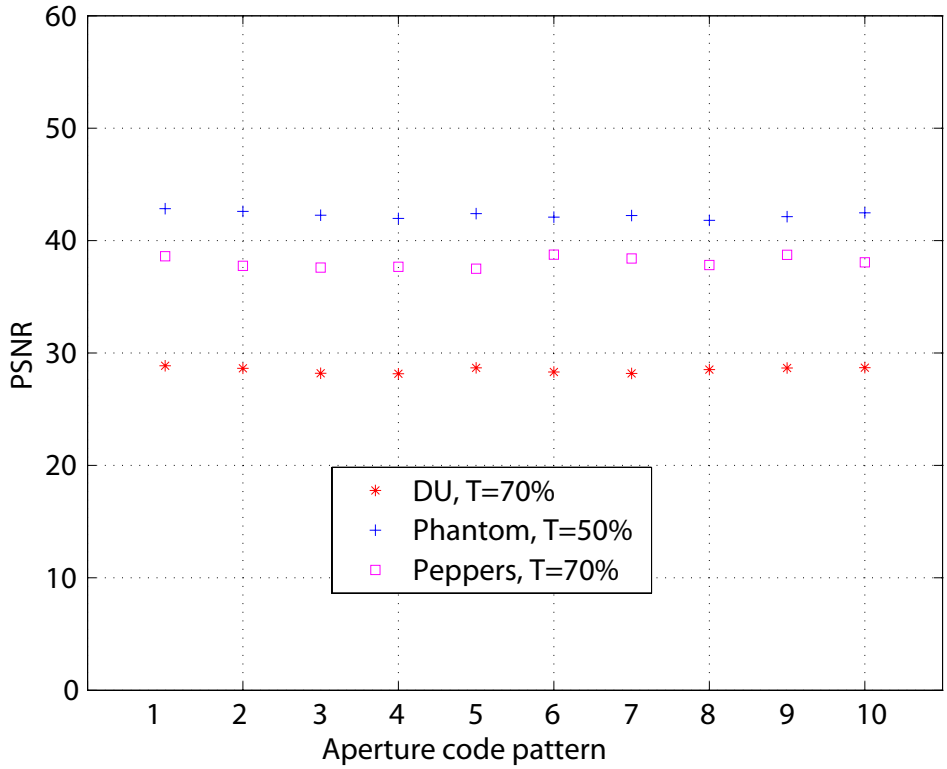

FIGURE 4.6: Variation of the PSNR of the reconstructed datacubes as a function of random aperture code pattern.

4.3.3 Scene dependence of CASSI performance

An important property of a system designed to recover a sparse signal from a compressed measurement is the ratio of the number of measurements made, M, to the sparsity of the signal, S, needed to accurately recover any signal with given sparsity S *regardless* of the locations of the S non-zero coefficients in the signal. Simply put, the ratio, M/S, should be as small as possible without dependence on the locations of the non-zero basis coefficients in S sparse signals. The smaller the ratio, the more efficient the measurement system. Thus, an ideal CASSI design is one that can be used to recover datacubes of any scene with equal PSNR for the smallest possible ratio M/S.

Before investigating the value of this ratio for CASSI, it is worth noting that the

number of non-zero CASSI detector measurements is typically scene dependent. As an extreme example, if the spatial distribution of a scene maps to the closed features on the coded aperture, represented by 0's in Φ, then the detector will not make any measurements whatsoever. Thus, for the purpose of comparison, all the 3 datacubes that were simulated produce an approximately equal number of non-zero detector measurements that is limited only by the transmission rate T of the aperture code. This was due to the presence of the (weak) background in all the spectral channels.

As another example, even though the total number of non-zero detector measurements produced by the 'DU' datacube matches that produced by the other two datacubes, the number of pixels that the 'D' and 'U' map to limits the "effective" number of non-zero measurements. As a result, removing the background produces a reconstruction with virtually the same PSNR (PSNR=28.37) and visual quality as that of a scene with the background (compare figure 4.7(a) producing PSNR=28.37 with figure 4.7(c) producing PSNR=28.77).

As figure 4.5 demonstrates, the performance of CASSI is scene dependent even when the number of non-zero detector measurements M is up to 10.3 times the number of non-zero coefficients S, which is approximately 1670. Figure 4.8 illustrates the best estimates of the 'phantom' and 'peppers' datacubes recovered from simulated CASSI measurements. This scene dependence can be related to the coherence of the CASSI sensing matrix Φ. Unlike a sensing matrix constructed through independent and identically distributed draws from a Gaussian or Bernoulli distribution that allows for accurate recovery with high probability provided that $M_i \geq C_i \cdot S_i \cdot \log(N_x N_y / S_i)$, the CASSI sensing matrix Φ requires $M_i \geq C_i \cdot \mu^2(\Phi_i, \Psi_{2D}) \cdot S_i \cdot \log(N_x N_y)$ for all i spectral channels for accurate recovery of the datacube that is completely independent of the scene.

Simulations were attempted to determine the approximate ratio of the number of non-zero CASSI detector measurements and the number of nonzero transform coeffi-

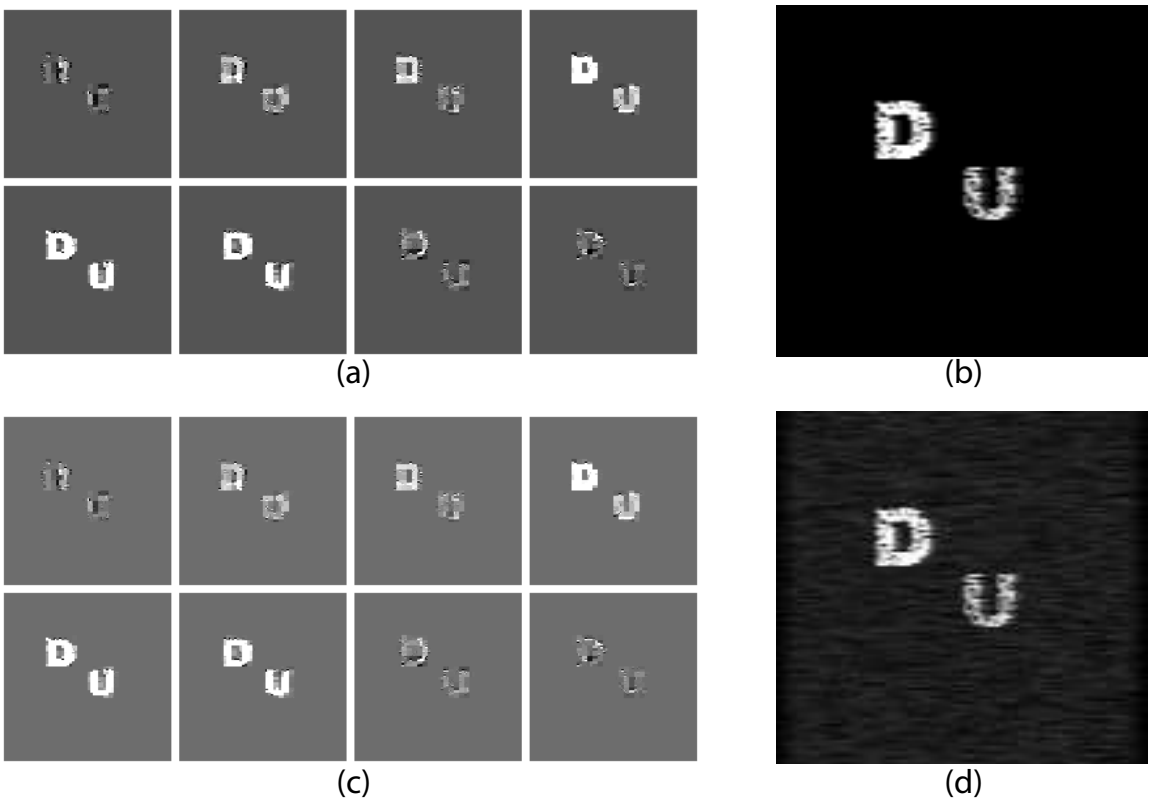

FIGURE 4.7: Comparison of reconstructions with and without the weak background. (a) PSNR=28.37 and (c) PSNR=28.77 show the reconstructions of the datacubes without and with the background, respectively. (b) and (d) show the associated detector measurements.

cients (i.e., M/S) required to accurately reconstruct datacubes that are independent of the scene. Table 4.1 shows the PSNR for 6 different datacubes designed such that their content distributions were constrained to produce roughly 100 2D wavelet coefficients. For these datacubes, the number of detector measurements required to produce a scene-independent reconstruction quality (PSNR $\simeq$ 58) was found to be at least around 5900, a factor of about 59 over S in each datacube. Note that there is a large variation in the PSNR below an aperture code transmission rate of 20%.

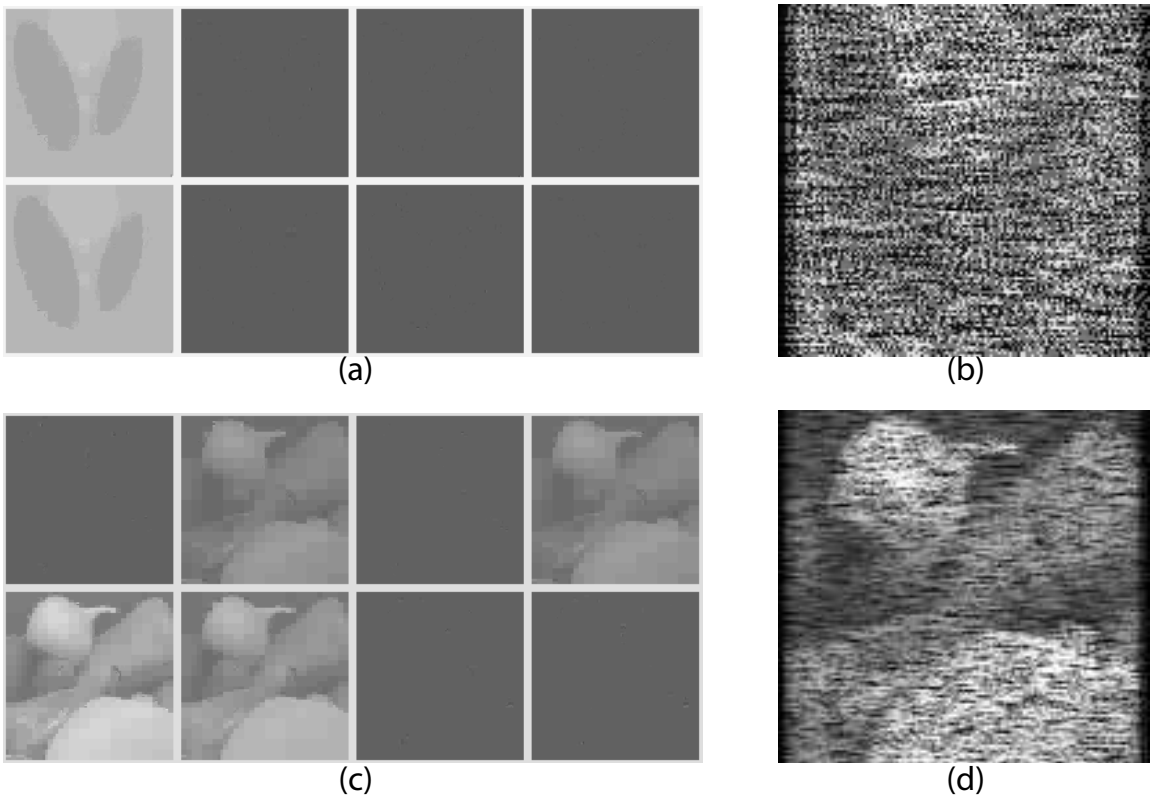

FIGURE 4.8: Best reconstructions of the 'phantom' and 'peppers' datacubes with 50% and 70% transmission rates, respectively. (a) PSNR=42.84 and (c) PSNR=38.61 show the best reconstructions of the datacubes, respectively. (b) and (d) show their corresponding detector measurements.

	S	M (T=5)	**PSNR** (T=5)	M (T=10)	**PSNR** (T=10)	M (T=20)	**PSNR** (T=20)	M (T=30)	**PSNR** (T=30)
Scene 1	96	1603	32.06	3139	50.79	5911	57.31	8423	61.32
Scene 2	104	1603	30.25	3139	37.20	5911	58.07	8423	62.18
Scene 3	96	1603	33.28	3139	38.48	5911	58.80	8423	61.98
Scene 4	102	1603	28.72	3139	52.64	5911	60.14	8423	63.94
Scene 5	100	1603	26.74	3139	33.85	5911	57.37	8423	61.59
Scene 6	102	1603	39.03	3139	48.58	5911	55.59	8423	59.53

Table 4.1: Determination of the minimal ratio of M and S (M/S) required for scene-independent accurate reconstructions. T denotes the transmission rate.

4.4 An alternative design for reducing the coherence of Φ

In the previous section, the scene dependence of CASSI's performance was attributed to the insufficiency of measurements, M_i, in each spectral channel, i. However, the

bound on the number of measurements required is not an accurate predictor of CASSI's performance for a given specific scene. For example, although the best PSNR achieved for the 'DU' datacube as 28.77, the best PSNRs achieved for the 'peppers' and 'phantom' datacube were well above 35. Thus, a global property like the mutual coherence (see equations (4.3) and (4.4)) between Φ and Ψ, which is used to set bounds on the number of measurements needed for accurate recovery is *not* useful to predict the reconstruction quality of specific scenes when the number of measurements is insufficient. Clearly the bound on the number of measurements is only a sufficient condition, but not a necessary one for accurate datacube recovery.

An alternative explanation must exist for why some scenes can be reconstructed well, while others like the 'DU' datacube are reconstructed poorly. A likely answer lies in the lack of the mutual orthogonality of the columns of the CASSI sensing matrix Φ or, as a consequence, the columns of the CASSI system matrix H. For H to preserve the lengths of S-sparse vectors b through dimensionality reduction tightly, as required by the RIP, the subset of columns of H that the S components of b map to should be as mutually orthogonal as possible.

The CASSI design is such that the structure of the sensing matrix, Φ (e.g., see figure 4.10(a)), may contain several identical columns. There may also be some entirely zero columns. It is not desirable to have such columns because there can be sparse vectors that are in the null space of Φ. This is in contrast to the well-known good CS sensing matrices (random Gaussian and Bernoulli) that have columns that are (statistically maximally) independent of each other and hence satisfy the RIP with a small restricted isometry constant. This suggests that the sensing matrix for an alternative, potentially improved CASSI design that has less scene dependent performance should have columns that are as mutually orthogonal as possible.

One solution is to relax the snapshot requirement of the CASSI measurement process. Then capturing multiple images of the scene while pushbrooming CASSI

by small amounts leads to a greater number of measurements, which makes recovery of the datacube using the effective sensing matrix less ill-posed. This approach is currently under investigation. Relaxing the snapshot requirement may also allow the implementation of positive *and* negative values in the effective sensing matrix, which may lead to a system matrix with a smaller restricted isometry constant.

At this point, it is important to note once again that it is not possible to physically realize a sensing matrix with random, negative values while maintaining a *snapshot* measurement process. The non-negative nature of a physically implementable matrix, along with the non-negative nature of the scene significantly restrict the possibility of creating an ideal CS sensing matrix. Since the space in which the columns of any physically implementable Φ reside is limited to the non-negative orthant, the angles between the columns are more likely to become highly coherent as the degree of dimensionality reduction becomes larger.

One potential snapshot design alternative to the conventional CASSI design discussed so far is to implement a coding scheme that simultaneously generates multiple shifted images of the scene at the coded aperture plane. The optical layout for the conventional CASSI described so far in this chapter and in chapters 2 and 3 is shown in figure 4.9(a). In this design, the objective lens forms a single image of the scene at the coded aperture plane. In the alternative design, the conventional CASSI is modified by replacing the objective lens with an array of smaller lenslets, as demonstrated in figure 4.9(b).

The lenslet array effectively generates multiple images of the scene, each of which is laterally shifted at the coded aperture plane depending on the position of the lenslet in the array. The coded aperture and the prism then encode each datacube, associated with each image at the coded aperture plane. Note that since the coded aperture position is physically fixed, each datacube is encoded in a different manner corresponding to the shifts of the image. These encoded datacubes are then

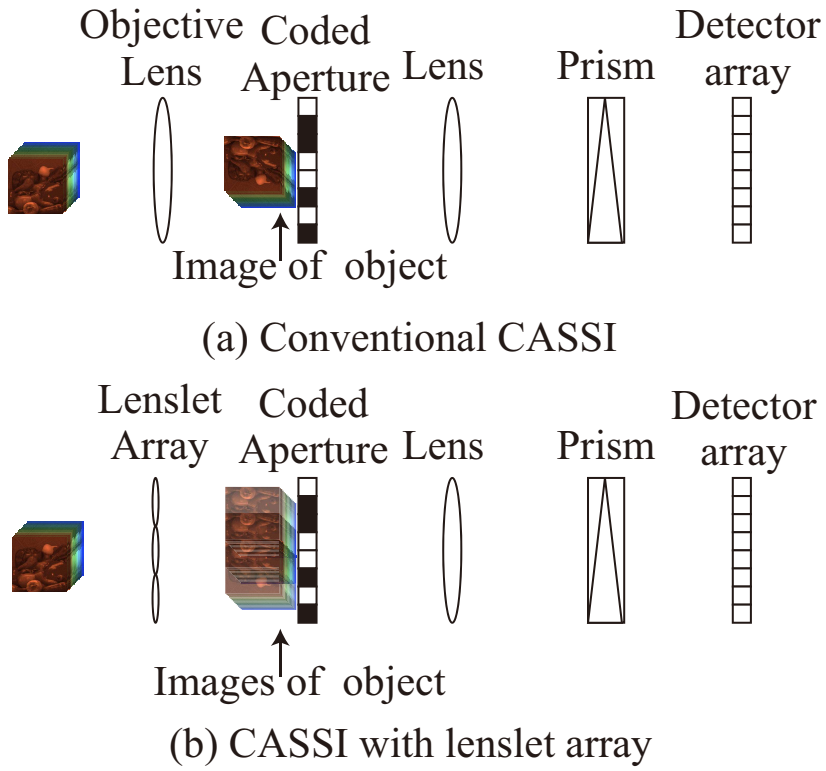

FIGURE 4.9: Improving CASSI performance by replacing the objective lens with a lenslet array of equal focal length. The lenslet array leads to multiple overlapping images of the scene on the coded aperture prior to dispersion by the prism.

multiplexed at the detector array (e.g., see figure 4.11(d)). For a system using a lenslet array with N_L lenslets, the detector measurements g may be mathematically modeled as

$$\begin{aligned} g &= \Phi \sum_{k=1}^{N_L} S_k f \\ &= \Phi_{new} f, \end{aligned} \tag{4.22}$$

where $S_k \in \mathbb{R}^{N_f \times N_f}$ represents a matrix that laterally (along the x- and y-axis) shifts the datacube $f \in \mathbb{R}^{N_f}$ corresponding to the k^{th} lenslet, Φ is the sensing matrix for the conventional CASSI as defined in section 4.2.1, and Φ_{new} is the new effective sensing matrix for the modified design.

With this new coding scheme, the number of 1's in each column of Φ_{new} is pro-

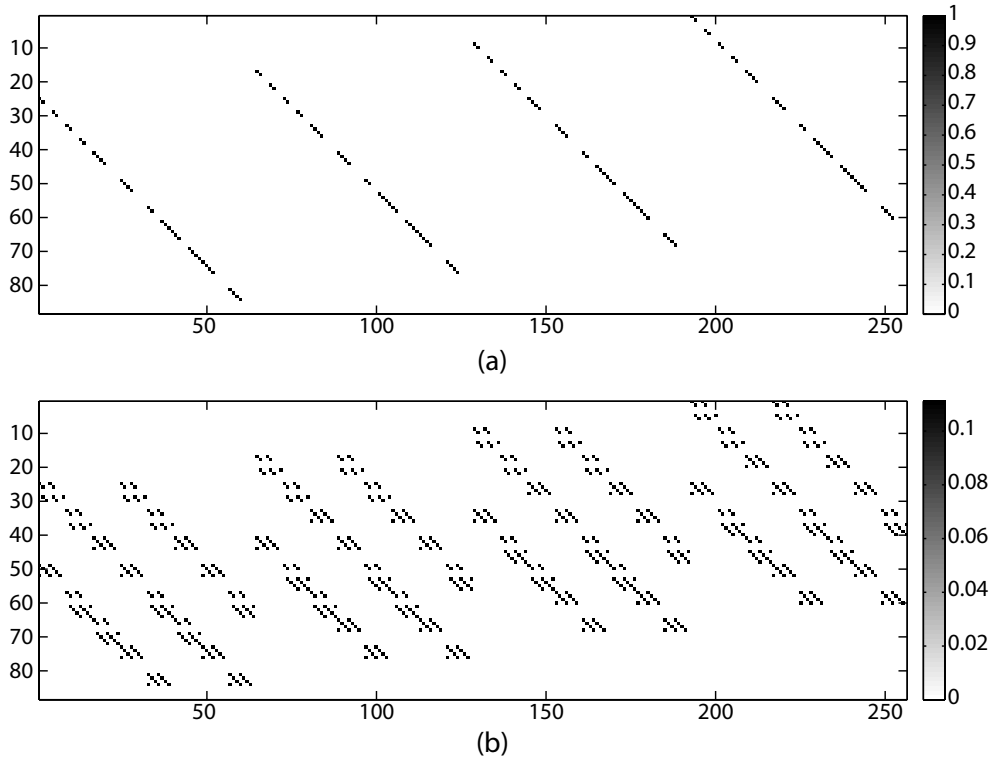

FIGURE 4.10: (a) Sensing matrix used to map an $8 \times 8 \times 4$ datacube on to the detector array of a conventional CASSI as shown in figure 4.9(a). (b) Sensing matrix used to map an $8 \times 8 \times 4$ datacube on to the detector array of a CASSI with a lenslet array consisting of 9 lenslets as shown in figure 4.9(b).

portional to the number of lenslets as shown in figure 4.10(b). As a consequence, the angles between the columns of the matrix also increase and the probability of any overlap between the columns decreases. The probability of having any zero columns in the matrix also decreases. Thus, the restricted isometry constant of a system matrix generated using Φ_{new} is likely to be smaller than the restricted isometry constant of a system matrix generated using the sensing matrix of the conventional CASSI, Φ, making the performance of the alternative design less scene dependent.

The performance of a lenslet array CASSI system consisting of a 7×7 lenslet array and a 50% transmission aperture code was studied with each of the three datacubes described in section 4.3. In a snapshot, the detector array would measure a superpo-

sition of 49 shifted and coded datacubes corresponding to the images produced by the lenslet array. Figures 4.11(b), (d) and (f) show the corresponding detector measurements. In all three cases, the number of detector measurements is the same as that generated by the conventional CASSI. Since the detector measurements multiplex the shifted copies of the datacube, the measurement does not bear any resemblance to the spatial distribution of the scene. In addition, while each lenslet has a lower light collection efficiency than the objective lens in the conventional CASSI, the overall amount of light collected by the instrument remains the same. As demonstrated by figure 4.11(d), the 49 shifted images produced by the lenslets have the largest amount of overlap around the center of the detector measurement and progressively lesser overlap towards the boundaries of the detector. As will be discussed in greater detail in the next section, this may cause the degradation of the conditioning of Φ_{new} and make the reconstructions from such a system less robust to noise.

Figure 4.11(a), (c) and (e) shows best reconstructions of the 'DU', 'phantom' and 'peppers' datacubes using the same aperture codes that produce the best reconstructions with the conventional CASSI. The high PSNRs of the reconstructions indicate highly accurate reconstructions of all 3 datacubes when no noise is present. The insensitivity of the PSNR to the nature of the sparse scenes can be attributed to the greater mutual orthogonality of the columns of Φ_{new} compared to that of Φ, the sensing matrix for the conventional CASSI.

The improvement can be roughly quantified by evaluating the mutual coherence $\mu(\Phi, \Psi)$ given in equation (4.3). For CASSI using a conventional objective lens and an aperture code with 50% transmission, $\mu(\Phi, \Psi)$ evaluates to $\frac{1}{2}\sqrt{N} \sim 180$. In addition, the mutual coherence using its alternative definition as given by equation (4.4), evaluates to its maximum possible value of 1. This indicates the presence of two or more identical columns in the CASSI sensing matrix. In contrast, the mutual coherence as given by equation (4.3) of the lenslet array based CASSI evaluates to 57.24,

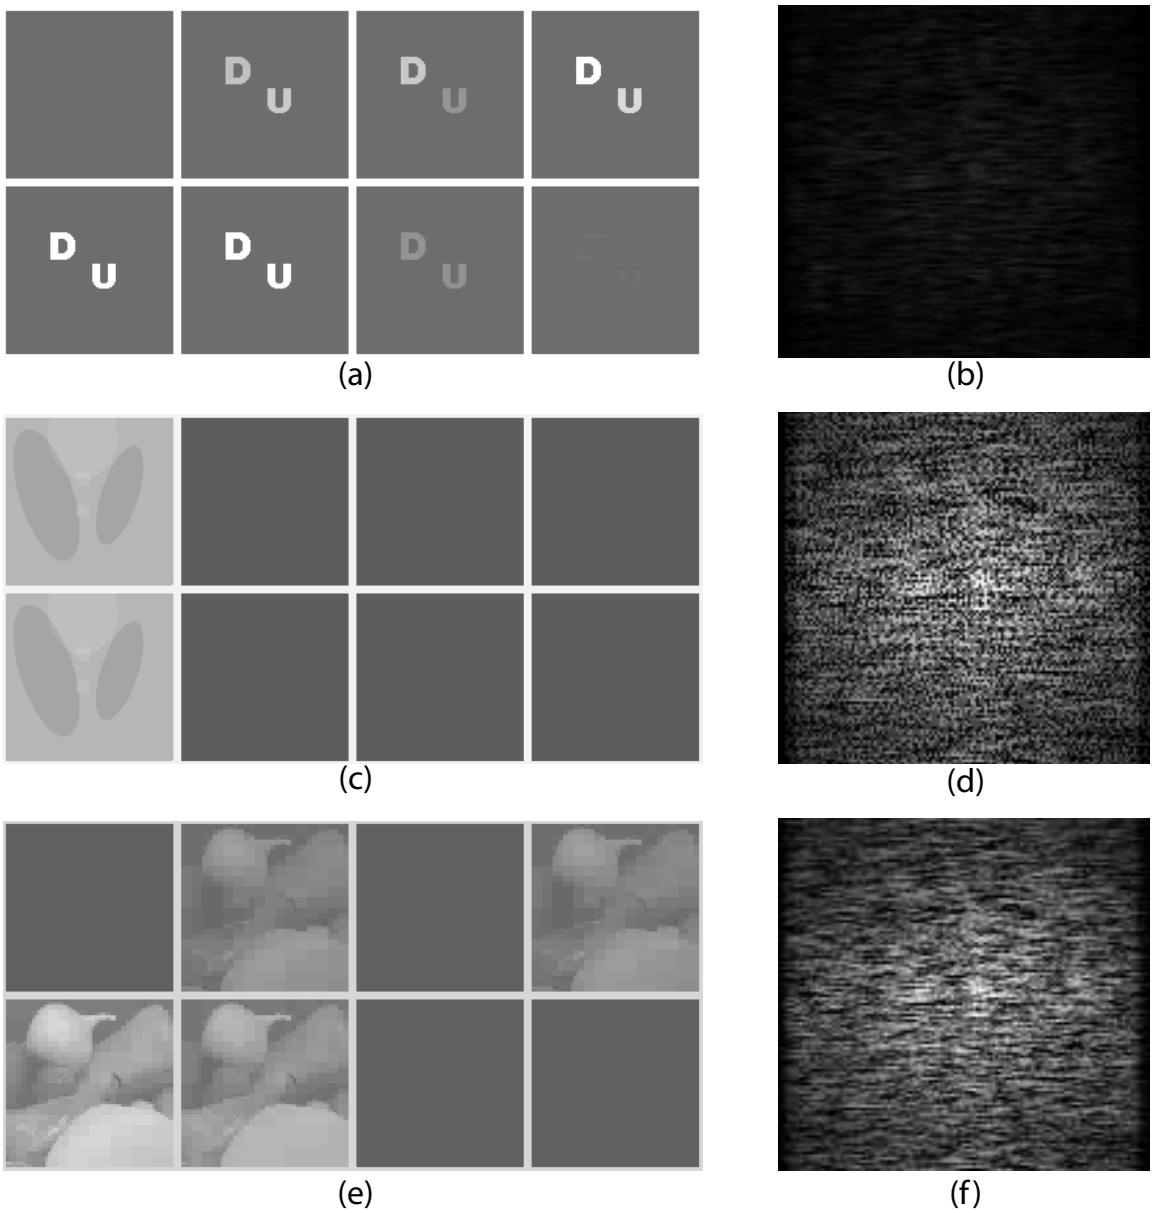

FIGURE 4.11: "Best" reconstructions of the 'DU', 'peppers' and 'phantom' datacubes. (a) PSNR=67.12, (c) PSNR=57.65 and (e) PSNR=56.72 show the best reconstructions of the three datacubes by a lenslet array CASSI, respectively. (b), (d) and (f) show their corresponding detector measurements.

while the alternative definition evaluates to 0.78. As the reconstruction examples show, the decrease in the mutual coherences leads to a significant improvement in the reconstruction fidelity.

4.5 Noise sensitivity of CASSI measurements

In this section, the performance of a conventional CASSI system and a lenslet array based CASSI is investigated in the presence of noise. Figure 4.12 and 4.13 show the changes in reconstruction PSNRs versus the changes in noise level for each system. Figures 4.12(a) and (b) show the results produced with conventional CASSI measurements corrupted by additive white Gaussian noise (AWGN) and Poisson noise, respectively. Similarly, figures 4.13(a) and (b) show the results produced with lenslet array based CASSI measurements. The AWGN is generated with the 'awgn' command in Matlab. The Poisson noise is generated with SNR defined by

$$\mathrm{SNR} = 10 \log_{10} \left[\frac{\sum_k g_k^2}{\sum_k (g_k^p - g_k)^2} \right], \tag{4.23}$$

where

$$g_k^p \sim \frac{1}{\alpha} \mathrm{Poisson} \left\{ \alpha g_k \right\}. \tag{4.24}$$

The constant factor α is chosen such that the measurement SNR equals the specified SNR in the plot. Poisson noise is added because real CASSI measurements are the intensities of optical fields which are subject to photon noise modeled by Poisson random processes. Both AWGN and Poisson noise are generated in 20 different levels to show the reconstruction PSNR changes. For each noise level, 20 different realizations of both noises are generated to produce the error bars.

The noise sensitivity curves for the conventional CASSI in figure 4.12 show two interesting features. First, the CASSI reconstructions are relatively robust to both Poisson and Gaussian noise. Although the reconstruction PSNR is scene dependent, it converges at a relatively low measurement SNR (around 40 dB). Examining the singular values of Φ reveals why this is the case. The singular values of Φ are simply the square roots of the number of 1's in each row, which means that the CASSI

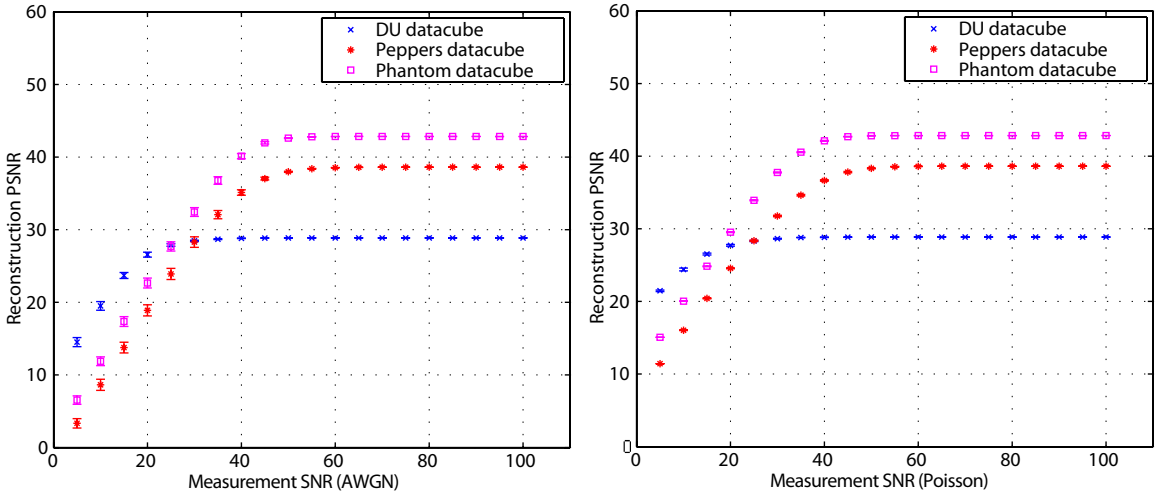

FIGURE 4.12: Plot of reconstruction PSNR from noisy CASSI measurements: (a) and (b) show reconstruction PSNRs produced with conventional CASSI measurements corrupted by additive white Gaussian noise and Poisson noise, respectively.

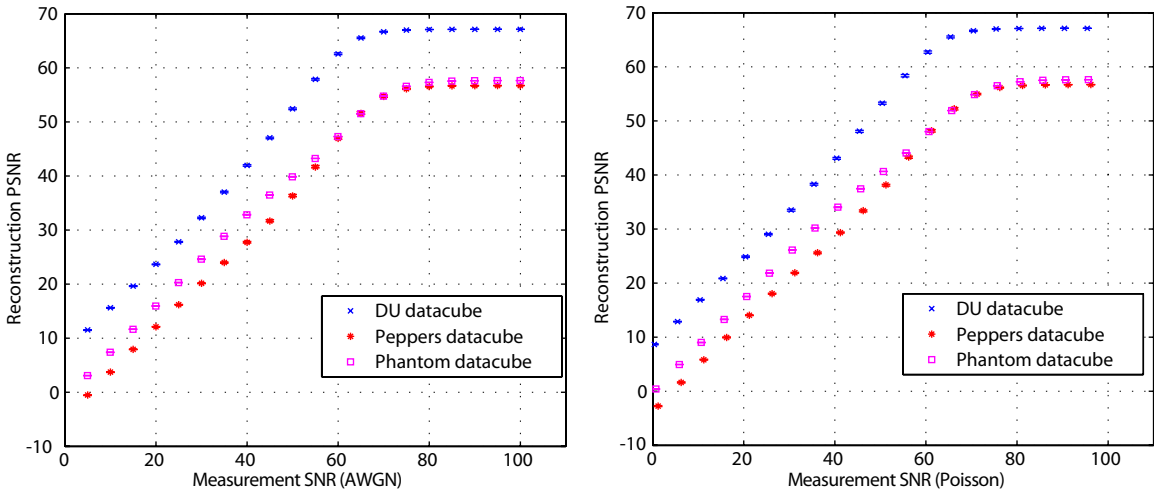

FIGURE 4.13: Plot of reconstruction PSNR from noisy lenslet array CASSI measurements: (a) and (b) show reconstruction PSNRs produced with measurements from a lenslet array CASSI corrupted by additive white Gaussian noise and Poisson noise, respectively.

sensing matrix is well-conditioned. For example, when there are 32 spectral slices, the largest possible singular value is $\sqrt{32}$ and the smallest possible singular value is $\sqrt{0}$, although the 0 singular values would rarely occur when N_λ is large enough. The rows corresponding to the 0 singular values can be easily ignored or removed without

losing information. Upon removing these rows, the effective minimum singular value is $\sqrt{1} = 1$, which produces a condition number $\sqrt{N_\lambda}$. Thus the stability of the performance of the conventional CASSI in the presence of noise can be attributed to the low condition number of Φ.

Second, the error bars are small in both the AWGN and Poisson noise cases. Interestingly, the Poisson noise error bars are smaller than the Gaussian noise error bars. Also, it is interesting that in low SNR levels, the CASSI measurements are more robust to the Poisson noise than the AWGN.

Note that the 'DU' datacube is less affected by noise because only a small set of pixels are significantly large such that they are affected by noise (especially by Poisson noise). Hence, for low SNR, the 'DU' datacube produces better reconstruction PSNRs compared to the other two datacubes.

The reconstruction results for the lenslet array based CASSI in figure 4.13 suggest that this system is not as robust to AWGN and Poisson noise as the conventional CASSI according to the slow convergence of the reconstruction PSNR. This can be attributed to the fact that the power distribution at the border of the detector measurement is much lower than at the center. This is because there is lower multiplexing of the different projections of the scene at the borders than at the center. A lower power distribution means a lower number of counts at each detector pixel and a corresponding increase in sensitivity to noise. Such non-uniform power distribution can be interpreted as leading to a condition number of Φ_{new} that is higher than that of Φ. The degradation in the condition number results in the alternative system being less robust to noise.

4.6 Implications for CASSI using a dual disperser architecture

It is important to note that many of the conclusions that were made in the previous sections for a CASSI using a single disperser are also valid for a CASSI system using

a dual disperser architecture [3]. The sensing matrix, Φ_{DD} for a dual disperser maps an $N_f = N_x \times N_y \times N_\lambda$ voxel datacube onto $M = N_x \times N_y$ detector pixels. Just like the single disperser CASSI, the sensing matrix for the dual disperser CASSI has the form, $\Phi_{DD} = [\Phi_1 \Phi_2 \cdots \Phi_{N_\lambda}] \in \mathbb{R}^{M \times N_f}$ as shown in figure 4.14(a). $\Phi_1 \in \mathbb{R}^{M \times (N_x \times N_y)}$ is a diagonal matrix that is constructed by vectorizing elements of the coded aperture column-wise. $\Phi_i \in \mathbb{R}^{M \times (N_x \times N_y)}$ is obtained by cycling the elements of the diagonal in Φ_1 a total of $(i-1) \times N_y$ rows. Compared to the single disperser CASSI sensing matrix shown in 4.14(b), Φ_{DD} is wider, which implies that the dual disperser maps a certain sized datacube to fewer pixels than the single disperser CASSI. Each pixel on the dual disperser detector is essentially programmed to apply a certain spectral filter on its corresponding pixel in the scene.

All of the algorithms described to process CASSI measurements in chapters 2 and 3 may be used for datacube recovery from dual disperser CASSI measurements. The algorithm described in [3] takes advantage of the fact that the detector measurement resembles the scene and utilizes spatio-spectral correlations in the scene to improve reconstruction quality.

Without coding introduced by the coded aperture and the dispersive elements, the detector on the dual disperser simply integrates the datacube along the wavelength axis and makes it impossible to recover the spatial content in different spectral channels. Given that the structure of the columns of Φ_{DD} is similar to that of Φ, the dual disperser is also expected to produce scene dependent datacube reconstructions. Just like the single disperser CASSI, the performance may become less scene dependent with the replacement of the front objective lens with a lenslet array, as discussed in section 4.4. Finally, since the rows of Φ_{DD} after removing all 0 rows are all orthogonal, Φ_{DD} will have a low condition number and will have stable performance in the presence of noise.

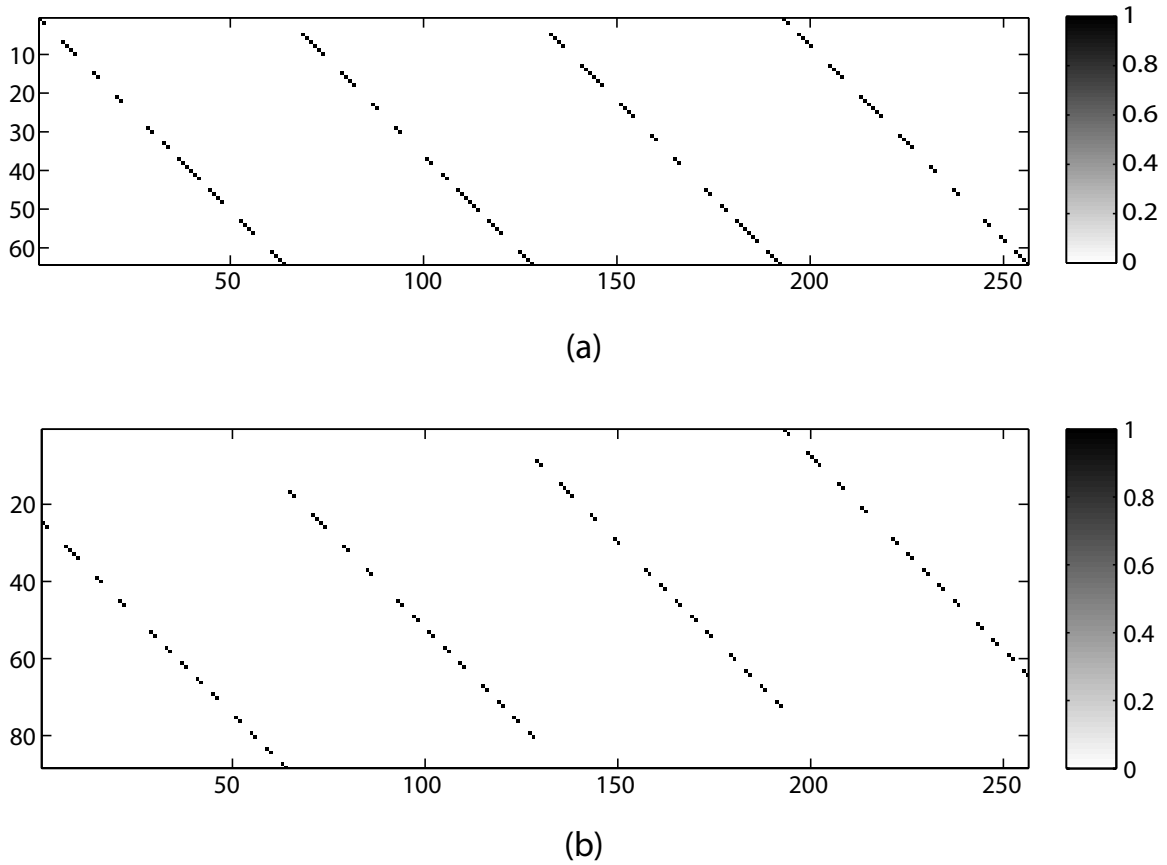

FIGURE 4.14: Comparison of CASSI sensing matrices for systems with a dual disperser and a single disperser architecture. (a) sensing matrix for a dual disperser CASSI, Φ_{DD}, that maps an $8 \times 8 \times 4 = 256$ voxel datacube to 64 detector pixels. (b) sensing matrix for a single disperser CASSI, Φ, that maps an $8 \times 8 \times 4 = 256$ voxel datacube to 88 detector pixels - identical to figure 3.12(a).

4.7 Conclusions

In this chapter, Matlab simulations were used to study the performance of the CASSI design presented in chapters 2 and 3 from the perspective of compressive sampling theory. Although optical distortions were ignored, the study provided key insights into properties of the CASSI sensing matrix, Φ and the system as a whole.

In particular, coding introduced by the coded aperture and the dispersive element was noted to be absolutely essential for any hope of datacube recovery from the significantly under-determined CASSI measurements. Furthermore, no optimal

aperture code was found; the open and closed features on the aperture code should simply be distributed as uniformly as possible in a random manner. Reconstructing 3 datacubes representing different scenes but having approximately the same number of sparse basis coefficients revealed the scene-dependent performance of CASSI. This was attributed to the presence of several identical columns, columns of zeros, and the lack of mutual orthogonality of the columns of Φ, making it the most significant insight provided by the Matlab simulations.

The sensing matrix, Φ_{new} for an alternative design that involved replacing the front objective lens of CASSI with a lenslet array was noted to have multiple 1's in each column, allowing the angles between the columns to be larger than those of Φ and allowing scene-independent reconstructions of datacubes. However, this also meant that, unlike the rows of Φ, the rows of Φ_{new} were not orthogonal and the condition number of Φ_{new} would likely be larger than that of Φ. Thus, while the alternative design did not exhibit scene-dependent performance in noiseless Matlab simulations, the performance was more sensitive in the presence of noise.

The results of this chapter have revealed that *snapshot* spectral imaging is a strong constraint for datacube reconstruction from a small number of conventional CASSI measurements. Nevertheless, it is important to note that there may be applications where the conventional CASSI might be useful due to the simplicity of the optical setup, particularly if the scene is sparse. A dual disperser CASSI system, which has a very similar sensing matrix to Φ, has been used for snapshot spectral imaging and identification of spatially uniform, but spectrally broadband, fluorescent microspheres that are often used in fluorescence microscopy [45, 46].

5

Imaging through turbulence using compressive coherence sensing

The Earth's atmosphere degrades the resolution of ground-based imaging systems attempting to image objects in space. The major contribution described in this chapter is an approach to image objects through turbulence by compressively sampling the mutual intensity, a coherence function of the optical field. The chapter begins with an introduction to the problem of imaging through turbulence in section 5.1. Section 5.2 reviews diverse methods for imaging through turbulence. The use of spectral imaging with CASSI is analyzed as a potential method to image through turbulence in section 5.3. Measurement of the mutual intensity, a 4D function, as a means of imaging and imaging through turbulence is introduced in section 5.4. A previous application of compressive sampling of the mutual intensity for imaging sparse astronomical objects is described in section 5.5. As the main focus of this chapter, an approach to image through turbulence using compressive sampling of the mutual intensity is presented in sections 5.6 and 5.7. The mutual intensity is sampled using a Rotational Shear Interferometer (RSI). Unlike a CASSI-based ap-

proach that would rely on the measurement of the power spectral density of the optical field, the RSI may be used to image through turbulence by measuring a subset of the mutual intensity of the field, assuming that the scene is spatially sparse. An experiment to test the ability of the RSI and the numerical estimation method to image remote point sources through isoplanatic turbulence is presented in section 5.8. Finally, a simulation testing the approach to image extended objects through turbulence is presented in section 5.9.

5.1 Introduction to imaging through turbulence

In the absence of atmospheric turbulence, an incoherent telescopic image of a space object can be modeled as a linear, shift-invariant process and represented as

$$g(x', y') = \iint \phi(x - x', y - y') f(x, y)\, dx\, dy, \tag{5.1}$$

where $f(x, y)$ is the object intensity distribution, $\phi(x, y)$ is the PSF of the imaging system and $g(x', y')$ is the image intensity distribution to be sampled by a detector array. For an incoherent imaging system, the PSF is given by

$$\phi(x, y) = \left| \iint P(u, v) e^{-i\frac{2\pi}{\lambda z}(ux+vy)}\, du\, dv \right|^2, \tag{5.2}$$

where $P(u, v)$ is a binary function defining the pupil of the imaging system, λ is the wavelength of light, z is the distance from the pupil to the image plane, and a scale factor is neglected. The resolution of an imaging system is limited by the finite size of the aperture, i.e. a $P(u, v)$ with a finite extent, which causes diffraction.

In part (a) of figure 5.1 (adapted from [47]), a diffraction limited image of a star obtained from a ground based imaging system in the absence of turbulence is simulated and displayed on a log scale to show the diffraction-induced Airy rings caused by the finite size of the telescope aperture. The figure is the convolution of the perfect image of the star with the PSF in equation (5.2).

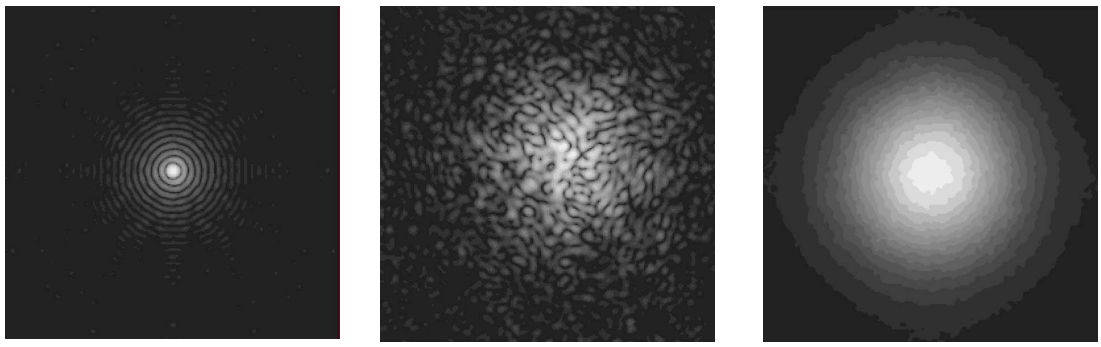

FIGURE 5.1: Simulated images of a star: (a) diffraction-limited image in the absence of turbulence; (b) short exposure image in turbulence; (c) long exposure image in turbulence. To produce these results, the authors simulated a telescope with a diameter of 1 m, average atmospheric turbulence conditions, and a mean wavelength of $\lambda = 550$ nm.

The resolution of a ground-based imaging system is further limited by phase aberrations induced by atmospheric turbulence. Turbulence arises due to temperature variations of the atmosphere over the course of the day, which leads to variations in the refractive index. The net effect is that the atmosphere acts as a collection of lenses of different refractive indices so that wavefronts from incoherent sources of light radiating in space undergo phase distortions [48]. These phase distortions are distributed across the pupil of the imaging system. Discussion of this problem throughout this chapter is limited to isoplanatic imaging, where the aberrations introduced by the turbulence are assumed to be uniform across the field of view of the imaging system.

Given an instantaneous turbulence-induced, non-zero optical path delay (OPD), $d(u, v)$, across the pupil, the expression for the PSF of an incoherent imaging system is equation (5.2) is modified to

$$\phi(x, y) = \left| \iint P(u, v) e^{j\, d(u,v)} e^{-i \frac{2\pi}{\lambda z}(ux+vy)}\, du\, dv \right|^2 . \tag{5.3}$$

Figure 5.1(b) above illustrates a short exposure image of the star in the presence of

turbulence. Here, short exposure means that the image exposure time is sufficiently short to freeze the atmosphere during the exposure. A short exposure typically has an integration time of less than 10 ms [49]. This image is a convolution of the perfect image of the star with a short-exposure PSF of the form shown in equation (5.3). As the figure shows, the effect of the turbulence is to broaden the system PSF compared to the diffraction limited PSF and to be highly speckled in its appearance.

Figure 5.1(c) above illustrates a long exposure image of the star in the presence of turbulence. In such an image, the exposure time is long enough to integrate over multiple realizations of the turbulence-induced aberrations. Thus, the long exposure PSF is broader than the diffraction limited PSF and smoother than the short exposure PSF shown in figure 5.1(b). A long exposure typically has an exposure time of greater than 0.1 s [49].

The effect of turbulence can be quantified in terms of the loss of angular resolution of the telescope. The angular resolution of a diffraction limited telescope of diameter D operating at a wavelength λ is

$$\theta \propto \frac{\lambda}{D}. \tag{5.4}$$

In the presence of turbulence,

$$\theta \propto \frac{\lambda}{r_0}, \tag{5.5}$$

where r_0 is the Fried parameter. This parameter is given by

$$r_0 = \left[0.42 \left(\frac{2\pi}{\lambda}\right)^2 \sec(\zeta) \int C_N^2(h) dh\right]^{-\frac{3}{5}}, \tag{5.6}$$

where ζ is the zenith angle and $C_N^2(h)$ is the refractive index structure constant as a function of height which provides a measure of the vertical turbulence profile [48]. The smaller r_0 is, the more severe the turbulence and the poorer the angular

resolution. Thus, r_0 may be thought of as the effective telescope diameter imposed by the atmosphere.

Removing the effect of atmospheric turbulence requires a method for characterization of the wavefront distortions as well as a method for the subsequent correction to recover an undistorted image of the scene. In the next section, some techniques designed to eventually recover undistorted images of objects through atmospheric turbulence are reviewed.

5.2 Approaches to imaging through turbulence

Speckle imaging

Speckle imaging is regarded as one of the earliest approaches to partially overcome atmospheric turbulence in astronomical imaging [50]. It relies on the fact that each speckle in a short exposure image is the size of the diffraction limited PSF and therefore contains high frequency information of the object [51]. Several short exposure images of the objects and a nearby reference star are captured and postprocessed to estimate the Fourier transform of the object, which is later inverted to recover an undistorted image of the object [50]. This technique has previously been used to image binary stars, where the speckle patterns are simple.

A time sequence of short exposures images of a common scene can also be postprocessed by using multiframe blind deconvolution, a technique where the undistorted image of the scene is estimated from turbulence degraded images and the time-varying PSFs are unknown [52]. This approach has the benefit of not requiring additional imaging of a reference star.

Lucky imaging

A significantly simpler technique for processing the short exposure images to recover a high resolution image of the object through turbulence is called lucky imaging [53].

Several thousand short exposure images of the object of interest are captured, but only $1-5\%$ of these images with the highest Strehl ratio of a reference star in the field are selected to be registered and added. This subset of images corresponds to exposures that are least affected by the time-varying effects of turbulence. The technique is limited by the large number of images that must be collected and the need for expensive, low noise, electron multiplying CCD detectors for faint object detection.

Adaptive optics

Possibly the most well known turbulence correction systems that are capable of real time turbulence correction are those based on adaptive optics technologies routinely in use on large telescopes like the Keck II telescope in Hawaii [54]. In these systems a deformable mirror placed in the optical path is controlled so that the path length errors in the light reflected from the mirror cancel out the errors induced by the turbulence. The actuators of the deformable mirror are controlled based on inputs received from wavefront sensing systems used to characterize the wavefront distortions.

Wavefront sensing

A commonly used wavefront sensor (WFS) is the Hartmann-Shack sensor [55]. This sensor consists of a 2D lenslet array located in a plane conjugate to the pupil plane of the telescope which produces a 2D array of spots on a detector array at the back focal plane of the lenslets. The positions of these spots are compared with the positions of reference spots and the displacements between the two are proportional to the local slope of the wavefront [56]. Since the light from science targets under observation is typically faint, a natural or laser-generated artificial guide star [57] is used for measuring the wavefront distortions.

The wavefront estimated from a Hartmann-Shack WFS may also be used to calculate the instantaneous PSF which can be used to post-process the instantaneous short exposure image of the object. Over time, different spatial frequencies of the object can be recovered from different short exposure images because of the random time variation of the turbulence [50]. This technique offers a compromise between an adaptive optics system, which relies on the mechanical correction of the wavefront, and speckle interferometry, which requires a larger number of short exposure images for correcting turbulence effects [56].

Another approach to wavefront sensing that can provide feedback to an adaptive optics system relies on phase diversity [58]. Two images of the object are simultaneously captured with different (known) phase aberrations optically added to each image. The known aberration is typically defocus [59]. Recovery of the wavefront from the recorded images is achieved through nonlinear optimization of some metric that measures the difference between the recorded data and a model of the data obtained with a given estimate of the wavefront [60]. The phase diversity sensor does not require a reference guide star, making it independent of the scene content.

Imaging through turbulence at non-optical wavelengths using radio interferometry

The universe is full of sources of electromagnetic radiation that are not observable by telescopes operating at optical frequencies. Imaging incoherent sources of radio waves like stars, galaxies and pulsars that may not be visible in optical frequencies requires a radio interferometric telescope. This section briefly introduces the van Cittert-Zernike (VCZ) theorem, which is used to image such sources using radio telescopes. The use of the VCZ theorem for imaging will be revisited in greater detail in sections 5.5 and 5.6.

From equation (5.4), it is easy to see that as the emission wavelength of the source increases from optical to radio wavelengths, the diameter of a telescope needed to

image the source with high angular resolution has to be considerably larger - too large to build as a continuous aperture. A radio telescope, like the Very Large Array in New Mexico, USA, consists of widely separated small radio telescopes that are connected using interferometry.

In the absence of turbulence, these small telescopes can be combined to synthesize a telescope whose effective aperture diameter is the size of the two telescopes furthest apart in the array - a technique known as aperture synthesis. The signal from each pair of telescopes is interfered to allow estimation of the amplitude and phase of the complex mutual intensity, a coherence function of the optical field. To recover an image of the object with an angular resolution dependent on the effective aperture diameter, the van Cittert-Zernike (VCZ) theorem [61] is employed, which states that the measured samples of the mutual intensity are in fact samples of the 2D Fourier transform of the spatial distribution of the incoherent radio wave source in space. By changing the locations and separations (baselines) of the small telescopes, different Fourier components can be measured [62]. To overcome the limitation of atmospheric turbulence on the effective telescope aperture, a technique known as phase closure may be used to eliminate the errors in the phases of the measured Fourier components [63, 64].

5.3 Spectral imaging to image through turbulence

In contrast to multiframe blind deconvolution, where an undistorted image of an object is estimated from a time sequence of frames degraded by turbulence, multi-channel blind deconvolution (MCBD) is a technique where images of an object are collected through two or more imaging channels with known inter-channel differences. The technique does not require the collection of a time sequence of frames.

Defocus-based phase diversity has previously been used as a MCBD scheme for simultaneous estimation of the object as well as the wavefront aberrations [60, 65].

Noting the requirement for multiple detector arrays, which can be a considerable expense for a phase diversity based MCBD system, it is interesting to consider alternative schemes to introduce inter-channel diversity.

Wavelength diversity has previously been considered as a MCBD scheme [66], although not specifically for the astronomical imaging case. To this end, spectral imaging with CASSI, as described in the previous 3 chapters, was investigated as a means of imaging through turbulence in a snapshot by exploiting any chromatic dispersion of the atmosphere for wavelength diversity of the PSF.

Expected wavelength diversity introduced by the atmosphere

Given that light incident on a spectral imager is divided into multiple spectral channels, the integration time would have to be considerably longer than $10 - 100$ ms, i.e. the typical integration time for a short exposure image. As mentioned earlier, a longer integration time results in a smoother and less structured PSF than for a short exposure. The full width half maximum (FWHM) of a long exposure image in the presence of atmospheric turbulence is called seeing [67], which is quantified as

$$seeing = 0.98\frac{\lambda}{r_0}. \tag{5.7}$$

Referring to equation (5.6), $r_0 \propto \lambda^{\frac{6}{5}}$, so that

$$seeing \approx \frac{\lambda}{r_0} \propto \lambda^{\frac{-1}{5}}. \tag{5.8}$$

Thus, seeing has a very weak dependence on wavelength. Considering that the spectral range of CASSI is only between 450 and 650 nm, the FWHM of the atmosphere-induced long exposure PSF at 650 nm only shrinks to 93% of its value at 450 nm - certainly not a very big change.

Experimental examination of turbulence-induced wavelength diversity

The analysis above suggested that the atmosphere alone would not provide significant wavelength diversity. To experimentally examine the change in the structure and size of the turbulence-induced PSF as a function of wavelength in the lab, 3 points generated using lasers at 473, 532 and 632.8 nm were observed one at a time by an imaging system. Adjustable neutral density filters were placed in front of each laser to ensure that the intensity of the three laser spots were the same.

The imaging system consisted of a grayscale CCD array with a camera lens in front of it, along with a phase screen placed immediately in front of the lens aperture to simulate the effect of the atmosphere. Different static phase screens were tested including the end of a petri dish, the end of a glass bottle, and a warped acrylic plate. A dynamically changing phase screen using the vapors from a propane flame was also tested.

None of the phase screens tested in the lab produced any significant variation in the PSFs as a function of wavelength. As equation (5.8) suggests, relying on the atmosphere to provide wavelength diversity requires a drastic change in wavelength. As it turns out, researchers in the wavefront sensing community have previously noted that a Hartmann-Shack WFS works well with polychromatic radiation [56], so that sensing the wavefront and aligning an adaptive optics system at an optical wavelength between $400 - 700$ nm is enough to adequately align the system even at IR wavelengths.

The quantitative analysis and experimental evidence described above suggested that the atmosphere does not introduce significant chromatic aberrations over the optical wavelengths. Without sufficient inter-channel diversity in the spectral image of the object degraded by turbulence, it would not be possible to use a MCBD scheme to recover an undistorted image of the object. For this reason, CASSI was not given

any further consideration as a means to image through turbulence.

5.4 Coherence imaging to image through turbulence

Besides radio interferometry, all of the approaches to image through turbulence described in section 5.2 and 5.3 treat the formed image as incoherent and do not attempt to infer coherence properties of the source from the image. Partial coherence is a seldom exploited, yet rich source of information that could be used not only for imaging, but also for imaging through turbulence.

As light radiates from incoherent sources in the far field, it propagates through free space or through a turbulent atmosphere and as it does so, it gains coherence. The mutual intensity is a measure of optical coherence. In addition to its use in imaging of incoherent sources by radio interferometry, the mutual intensity may also be used to image the turbulence-induced wavefront aberrations.

In a 2D aperture, the mutual intensity is a 4D function, $J(\mathbf{r_1}, \mathbf{r_2})$, capturing the correlations between every pair of points $\mathbf{r_1} = (x_1, y_1)$ and $\mathbf{r_2} = (x_2, y_2)$ in the aperture. The most direct way to measure J is to measure the interference of the light incident on two pinholes that are spatially scanned across all pairs of coordinates in the aperture. This approach is tedious and time consuming.

An alternative approach is to use an astigmatic coherence sensor shown in figure 5.2 (adapted from [68]), that essentially captures 2D planes of correlations at once. The ACS consists of a CCD array that is scanned axially and three cylindrical lenses which are rotated to adjust the horizontal and vertical focal lengths of this imaging system. Over time, a 4D array of intensity measurements can be captured which have a 4D Fourier transform relationship with J. Thus the ACS may be viewed as a generalized phase diversity sensor. The ACS has previously been used to image remote incoherent sources through an isoplanatic, refractive distortion in its pupil by using a technique known as the coherent mode decomposition [69]. This technique

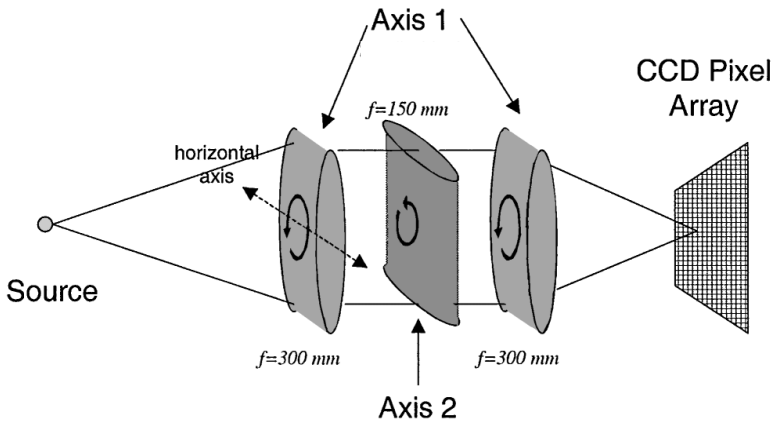

FIGURE 5.2: The astigmatic coherence sensor can sample the 4D mutual intensity function, J, in its aperture. It consists of a CCD array that is scanned axially and three cylindrical lenses which are rotated to adjust the horizontal and vertical focal lengths of this imaging system. Over time, a 4D array of intensity measurements can be captured which have a 4D Fourier transform relationship with J.

utilizes the idea that a common distortion occurs to the coherent modes from each source of light in the scene and this idea is used to separate the distortion from the intensity distribution of sources behind it. Unfortunately, the use of the ACS for imaging through turbulence is impractical and slow for undistorted image formation, as it requires the complete sampling of the 4D J, which can take several hours.

5.5 Imaging by compressively sampling the mutual intensity

Compressive sampling can be used to remove the need for exhaustive sampling of J given the prior information that the object being imaged is sparse. In radio interferometry, imaging of sparse objects can be cast into an inverse problem based on the VCZ theorem. According to the theorem, in a 2D aperture located in the far field of an incoherent source (a space object), the mutual intensity of the optical field is $J(\mathbf{r_1}, \mathbf{r_2}) = J(\mathbf{r_1} - \mathbf{r_2}) = J(\mathbf{\Delta r})$ and is related by a 2D Fourier transform to the source intensity distribution of the object [50]. Image reconstruction of the sparse objects must be performed using noisy Fourier measurements (samples of the mutual

intensity, J) that are captured by interfering signals from pairs of radio telescopes. These measurements represent an incomplete sampling of the effective aperture of the radio telescope because they are captured at a finite number of locations and separations within the aperture.

Incomplete Fourier measurements of scenes that are spatially sparse (e.g. star fields) are particularly attractive in the CS sense because the incoherence between the Fourier basis as the measurement basis, and real space as the basis where the scene is sparse, is maximum. Upon discretization of the inverse problem, this is equivalent to stating that the mutual coherence, as defined in equation (4.3), between Φ, the discrete Fourier transform matrix, and Ψ, the identity matrix that discretizes real space, is $\mu = 1$. Assuming that the ratio of the number of incomplete (random) Fourier measurements to the number of sparse objects in the discretized scene is between $3 - 5$ [70], the system matrix $H = \Phi\Psi$ satisfies the RIP in equation (4.1) and thus assures the accuracy of the signal recovered by solving the problem in equation (4.7).

Simulations have shown the feasibility of this idea for radio interferometric imaging of sparse astrophysical objects, as well as cosmic strings in the temperature field of the cosmic microwave background radiation [71], which may be considered sparse in the TV sense. The focus of the remainder of this chapter is the extension of compressively sampling of J to not only image sparse objects, but to image them through isoplanatic turbulence.

5.6 Imaging through turbulence by compressively sampling the mutual intensity

In this section, a rotational shear interferometer (RSI) is described as a means for imaging through turbulence by measuring a limited subset of samples of J. As mentioned earlier, the mutual intensity is a 4D function describing the coherence of

the optical field and can be represented as

$$J\left(\bar{\mathbf{r}}-\frac{\Delta\mathbf{r}}{2},\bar{\mathbf{r}}+\frac{\Delta\mathbf{r}}{2}\right)=\left\langle E\left(\bar{\mathbf{r}}-\frac{\Delta\mathbf{r}}{2}\right)^{*}E\left(\bar{\mathbf{r}}+\frac{\Delta\mathbf{r}}{2}\right)\right\rangle, \tag{5.9}$$

where $E(\mathbf{r})$ is a quasimonochromatic, stochastic electromagnetic field of wavenumber $k=\frac{\omega}{c}$ and $\bar{\mathbf{r}}$ is the mean position of a pair of coordinates separated by $\Delta\mathbf{r}$. Thus, $\bar{\mathbf{r}}=\frac{\mathbf{r_1}+\mathbf{r_2}}{2}$ and $\Delta\mathbf{r}=\mathbf{r_2}-\mathbf{r_1}$. Here the quasimonochromatic assumption ensures that the bandwidth of the optical field is finite, but small enough that the temporal coherence effects can be ignored.

The van Cittert-Zernike (VCZ) theorem states that J in a 2D aperture in the far field of a spatially incoherent source with planar irradiance $S(\mathbf{r})$ depends only on $\Delta\mathbf{r}$ as

$$J(\Delta\mathbf{r})=\int\frac{S(\mathbf{r})}{z^2}e^{\frac{-2\pi i}{\lambda z}(\Delta\mathbf{r}\cdot\mathbf{r})}d^2\mathbf{r}. \tag{5.10}$$

Thus, J is the 2D Fourier transform of the far field source distribution. As mentioned in the previous section, this relationship forms the basis for radio interferometric imaging. To image a remote incoherent object using the same relationship at optical frequencies, J may be sampled as a function of $\Delta\mathbf{r}$ about a central point using an interferometer with a 2D detector array and then inverse Fourier transformed.

If isoplanatic turbulence with an optical path delay, $d(\mathbf{r}')$, is present between the scene and the aperture,

$$J\left(\bar{\mathbf{r}}-\tfrac{\Delta\mathbf{r}}{2},\bar{\mathbf{r}}+\tfrac{\Delta\mathbf{r}}{2}\right)=\left[\int\frac{S(\mathbf{r})}{z^2}e^{\frac{-2\pi i}{\lambda z}[(\Delta\mathbf{r})\cdot\mathbf{r}]}d^2\mathbf{r}\right]\cdot e^{\frac{2\pi i}{\lambda}\left[d\left(\bar{\mathbf{r}}-\frac{\Delta\mathbf{r}}{2}\right)-d\left(\bar{\mathbf{r}}+\frac{\Delta\mathbf{r}}{2}\right)\right]}. \tag{5.11}$$

Note that J is now a function of both $\bar{\mathbf{r}}$ and $\Delta\mathbf{r}$. Thus, when turbulence is present, directly applying the VCZ theorem described in equation (5.10) for inversion produces a distorted image of $S(\mathbf{r})$. Accurate recovery of $S(\mathbf{r})$ would normally require knowledge of the complete 4D distribution of J, which may be measured using an instrument like the ACS. However, an RSI is uniquely suited to measure J as a

product of the Fourier transform of the source distribution and a phase screen that depends on the OPD. The use of an RSI to estimate both the source distribution and the phase screen from just a 2D subset of the 4D J is experimentally demonstrated in section 5.8.

An RSI [61], as shown in figure 5.3, consists of two right-angle fold mirrors, a beam splitter and a CCD array. The field incident on the input aperture is first split into two copies by the beam splitter and directed towards the two fold mirrors. Light reflected from each fold mirror is recombined by the beam splitter to form an interferogram on the CCD array. One of the two fold mirrors is placed on a piezo translation stage that must be translated to allow isolation of the complex J from the intensity measurements made on the CCD.

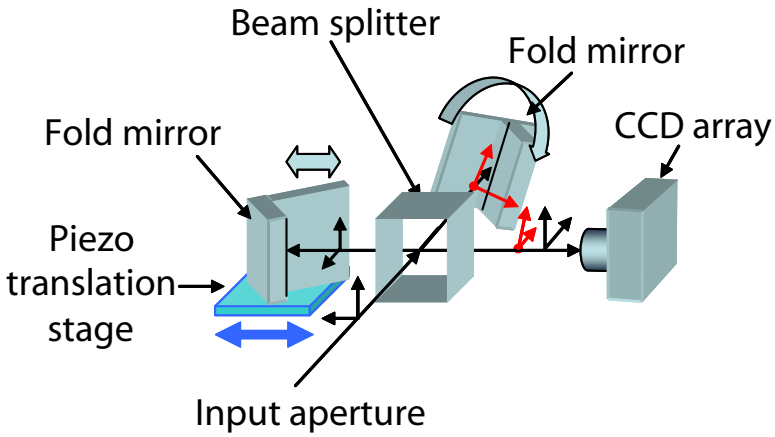

FIGURE 5.3: The rotational shear interferometer can sample the mutual intensity in its aperture. The field incident on the aperture is split into two copies by the beam splitter and directed towards the two fold mirrors. Light reflected from each fold mirror is recombined by the beam splitter to form an interferogram on the CCD array. One of the two fold mirrors is placed on a piezo translation stage. By capturing intensity measurements on the CCD at multiple translations, a 2D subset of samples of the complex 4D J can be isolated.

Each fold mirror consists of two mirrored surfaces joined at right angles. Unlike an ordinary mirror, a fold mirror with its fold axis positioned vertically flips the image horizontally. Thus, any text visible in the reflection from a fold mirror will

still be readable from left to right. Another interesting property is that a reflection from a fold mirror when its fold axis is oriented at an angle θ (the shear angle) will undergo a rotation of 2θ. Figure 5.4(b) (reproduced from [72]) displays the image of the scene in figure 5.4(a) as recorded by the RSI detector. The effect of rotation on the reflections from each mirror is easily observed.

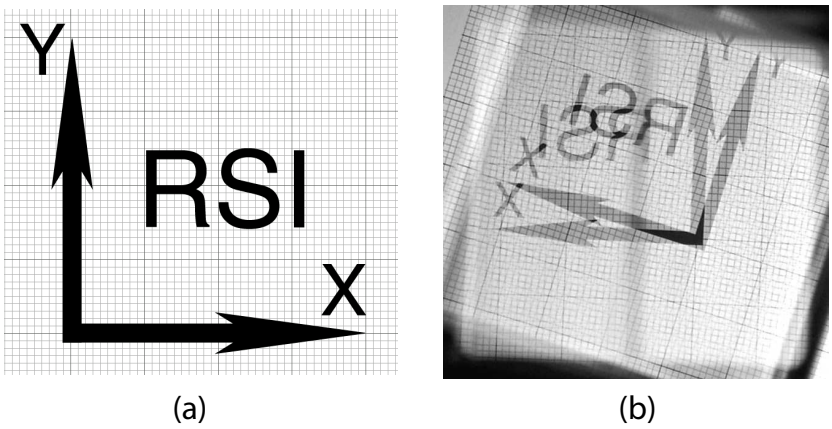

FIGURE 5.4: (a) An example of a scene to be imaged through the RSI. (b) Image of the scene recorded on the RSI detector array. One of the copies of the image is rotated due to the rotation of one of the two fold mirrors.

The continuous intensity distribution sampled at the RSI detector plane is a function of the mutual intensity between points in the two rotated copies of the input aperture that are being superimposed by the beam splitter after reflection from each fold mirror. Denoting the points in each copy as $\mathbf{r_1}$ and $\mathbf{r_2}$ respectively, the intensity measured at the detector is

$$\begin{aligned} I &= \left\langle \left| E(\mathbf{r_1}) + e^{-2\pi i \frac{l}{\lambda}} E(\mathbf{r_2}) \right|^2 \right\rangle \\ &= \left\langle \left(E(\mathbf{r_1}) + e^{-2\pi i \frac{l}{\lambda}} E(\mathbf{r_2}) \right) \left(E(\mathbf{r_1}) + e^{-2\pi i \frac{l}{\lambda}} E(\mathbf{r_2}) \right)^* \right\rangle \\ &= \left\langle \left(E(\mathbf{r_1}) + e^{-2\pi i \frac{l}{\lambda}} E(\mathbf{r_2}) \right) \left(E(\mathbf{r_1})^* + e^{2\pi i \frac{l}{\lambda}} E(\mathbf{r_2})^* \right) \right\rangle \\ &= I(\mathbf{r_1}) + I(\mathbf{r_2}) + 2\mathbb{R}\left\{ J(\mathbf{r_1}, \mathbf{r_2})\, e^{-2\pi i \frac{l}{\lambda}} \right\}, \end{aligned} \tag{5.12}$$

where l is the path length difference between the two optical paths in the RSI. Isolation of the complex J from the intensity measurements on the CCD requires the capture of four interferograms captured while translating one of the fold mirrors using the piezo stage to modify the path length difference between the two arms, a technique known as phase shifting. By capturing interferograms I_1, I_2, I_3 and I_4 at instances where $2\pi\frac{l}{\lambda} = 0$, $2\pi\frac{l}{\lambda} = \frac{\pi}{2}$, $2\pi\frac{l}{\lambda} = \pi$ and $2\pi\frac{l}{\lambda} = \frac{3\pi}{2}$ respectively, the complex J can be isolated by computing

$$\begin{aligned} J(\mathbf{r_1}, \mathbf{r_2}) &= \frac{1}{4}\left(I_1 e^{0i} + I_2 e^{\frac{\pi}{2}i} + I_3 e^{\pi i} + I_4 e^{\frac{3\pi}{2}i}\right) \\ &= \frac{1}{4}\left[(I_1 - I_3) + i\,(I_2 - I_4)\right]. \end{aligned} \tag{5.13}$$

Now consider the specific case where the fold axes of the two fold mirrors are rotated by $-\theta$ and θ respectively and let the transverse coordinates at the detector plane of the RSI be (x, y). The fields being superimposed at (x, y) are

1. the field that would appear at $\mathbf{r_1} = (-x\cos 2\theta + y\sin 2\theta\,,\; x\sin 2\theta + y\cos 2\theta)$ if the fold mirror rotated by $-\theta$ is replaced by a plane mirror, and

2. the field that would appear at $\mathbf{r_2} = (-x\cos 2\theta - y\sin 2\theta\,,\; -x\sin 2\theta + y\cos 2\theta)$ if the fold mirror rotated by θ is replaced by a plane mirror [72].

Thus $\bar{\mathbf{r}} = \frac{\mathbf{r_2}+\mathbf{r_1}}{2} = (-x\cos 2\theta, y\cos 2\theta)$ and $\mathbf{\Delta r} = \mathbf{r_2} - \mathbf{r_1} = (-2y\sin 2\theta, -2x\sin 2\theta)$. Substituting these expressions into equation (5.11) and defining the coordinates in the plane of the far field, incoherent source as $\mathbf{r} = (x', y')$, the continuous distribution of the complex mutual intensity sampled by the detector array is given by

$$J_{RSI}(x, y) = \left[\iint \frac{S(x', y')}{z^2} e^{\frac{4\pi i}{\lambda z}(yx' + xy')\sin(2\theta)} dx' dy'\right] \cdot e^{i\gamma(x,y)}, \tag{5.14}$$

where $\gamma(x, y) = \frac{2\pi}{\lambda}\left[d\left(\bar{\mathbf{r}} - \frac{\Delta\mathbf{r}}{2}\right) - d\left(\bar{\mathbf{r}} + \frac{\Delta\mathbf{r}}{2}\right)\right]$. This measurement represents a rotated and scaled version of the 2D J in equation (5.10) that is modulated by a 2D

turbulence-induced phase screen, $\gamma(x, y)$, in the RSI aperture. In the absence of turbulence, the measured samples have Hermitian symmetry, that is $J_{RSI}(x, y) = J_{RSI}(-x, -y)^*$, since they represent a Fourier transformation of a real function, the irradiance of the scene, $S(x', y')$. Unless the shear angle $\theta = \frac{\pi}{4}$, the samples do not have Hermitian symmetry in the presence of turbulence. Because of this disparity, the turbulence phase can be distinguished from the phase due to the scene itself.

5.7 Numerical estimation method to image through turbulence

Let g represent a vector that lists discrete samples of $J_{RSI}(x, y)$ in equation (5.14), the 2D subset of the 4D mutual intensity measured by the RSI. This measurement can be represented as a matrix-vector product

$$\begin{aligned} g &= PFx + n \\ &= \Phi x + n, \end{aligned} \tag{5.15}$$

where n is noise in the measurement, x is a vector whose elements are a discrete representation of the scaled and rotated unaberrated image of the scene consisting of incoherent sources radiating in the far field, F is the 2D discrete Fourier transform (DFT) matrix, P is a discrete representation of the turbulence-induced phase screen that is denoted by $e^{i\gamma(x,y)}$ in equation (5.14), and Φ is the sensing matrix for a given instance of the phase screen.

At this point, it is important to distinguish the RSI measurement from a conventional CS measurement like the one implemented by CASSI. A 2D CASSI measurement is digitally processed to recover a 3D datacube *given* the under-determined sensing matrix that models the mapping from the 3D datacube to the 2D measurement. In contrast, the problem of recovering an undistorted image of the scene from just a 2D subset of the 4D J measured using the RSI is complicated by the fact that the phase screen P is unknown, which implies that the sensing matrix, Φ in equation

(5.15) is unknown. If there is no turbulence present, then an RSI measurement of the scene is *not* compressive because the 2D measurement is related to the image of the scene by a 2D Fourier transform according to the VCZ theorem in equation (5.10). However, in the presence of turbulence, an RSI measurement *is* compressive in the sense that both the 2D scene and the 2D phase screen must be estimated from a single 2D subset of J.

Recovering an undistorted image of the scene involves digitally processing the vector of measurements from the RSI to *jointly* estimate x and P. Mathematically, solving this problem can be posed as

$$[x_e, P_e] = \arg\min_{x,P} ||g - PFx||_2^2, \tag{5.16}$$

where x_e is a discrete estimate of the 2D scene $S(x', y')$, P_e is a discrete estimate of the turbulence-induced phase screen $e^{i\gamma(x,y)}$, g is the vector of the 2D subset of samples of J measured on the RSI, and F is the discrete Fourier transform (DFT) matrix. Since P and x are both unknown, this is a non-convex problem, which implies that a globally optimal solution cannot be found efficiently.

Part of a potential solution to this joint recovery problem involves borrowing the CS assumption of sparsity. Just as the datacube representing a scene was assumed to be sparse when imaging with CASSI, a sparsity assumption may be placed on the scene being imaged through turbulence using the RSI. Consider the case where the scene to be imaged through turbulence is sparse and consists of point sources, the astronomical equivalent of a star field. Since the sensing matrix in equation (5.15) is not known, an approach to solving the problem posed in equation (5.16) is an iterative two-step alternating minimization (AM) technique [42]. In the first step, the TwIST algorithm described in section 3.3, is used to estimate an undistorted and sparse distribution of sources in the scene that fits the measurements of J as closely as possible, given an estimate of the phase screen. In the other step, the

phase screen estimate is refined using the measurements of J from the RSI and a progressively undistorted version of J that is modeled using the current estimate of the sparse scene. These two steps are repeated until there is minimal change in the estimate of the sparse scene. Mathematically, the i^{th} iteration of the AM technique is represented as

$$x_e^i = \arg\min_x \frac{1}{2}||g - \Phi x||_2^2 + \lambda||x||_1$$

$$= \arg\min_x \frac{1}{2}||g - P_e^{i-1} F x||_2^2 + \lambda||x||_1 \tag{5.17a}$$

$$P_e^i = \text{sgn}\left[\left(F x_e^i\right)^* \cdot g\right], \tag{5.17b}$$

where x_e^i and P_e^i are the estimates of the scene and the phase screen at the i^{th} iteration to be recovered from g, a vector of samples of the measured J. Equation (5.17b) is the least squares solution to $g = Xp$, where the matrix X is a diagonal matrix listing the DFT coefficients of x_e^i on the diagonal and p is a vector of the elements of P_e^i with the constraint that $|P_e^i| = 1$.

5.8 Experimental demonstration of imaging a sparse scene through turbulence by compressively sampling the mutual intensity

To experimentally verify the utility of this approach, a scene consisting of light emitting diodes (LEDs) was arranged and was observed through an aberrating phase distortion using the RSI, as detailed in figure 5.5.

The LEDs emitted light over 590 ± 30 nm (coherence length = 5.8 μm) and had their plastic lenses sanded off to produce mutually incoherent radiation over 2π sr. Note that the spectral bandwidth of the LEDs was wide enough to ensure that the radiation was spectrally incoherent, but narrow enough so that interference fringes could still be observed on the interferometer. These LEDs were collimated by a lens to place their image at infinity as would occur for stellar imaging. An iris was used

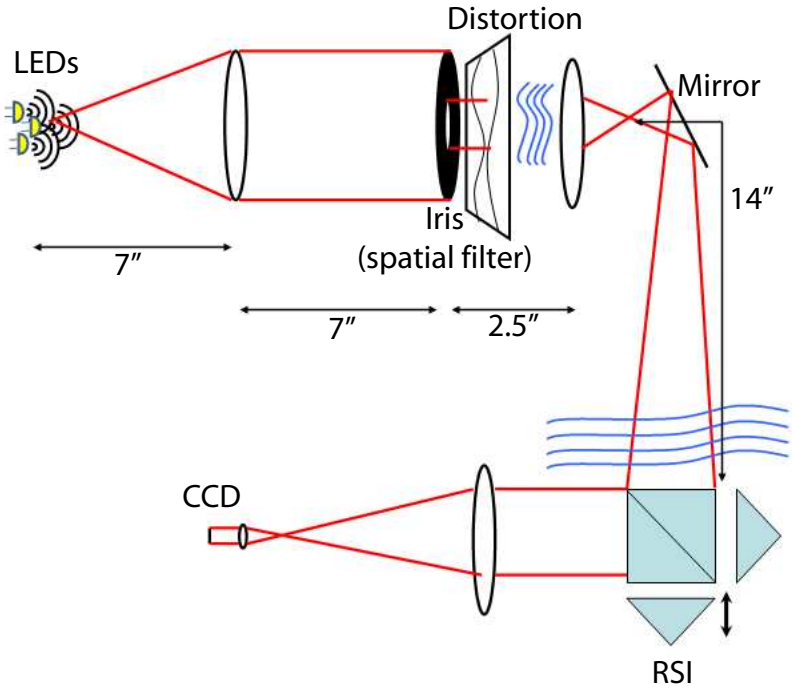

FIGURE 5.5: Experimental setup to measure the mutual intensity at the RSI aperture from LEDs radiating through a phase distortion. Light from the LEDs is collimated and passes through a phase distortion plate placed behind the iris to produce an instance of the phase distortion caused by turbulence. The distorted wavefronts then propagate to the RSI aperture. One of the two fold mirrors in the RSI is translated using a piezo stage to adjust the path length difference between the two arms of the RSI. 4 intensity measurements are captured on the CCD at set path length differences. A 2D subset of the complex, 4D mutual intensity can be isolated from these intensity measurements.

to eliminate vignetting effects from the collimation lens. A phase distortion plate was placed behind the iris to produce an instance of the phase distortion caused by turbulence. This plate was constructed by dripping polydimethylsiloxane (Sylgard 184) silicone onto a microscope slide and then heating it with a heat gun to cure it before the surface became uniformly flat. A second lens relayed the partially coherent aberrated field to the entrance of the RSI. The RSI was used to measure interferograms as the delay between the two arms of the interferometer was varied. The interferograms at four delays with phases separated by $\pi/2$ were used to find a

2D subset of J according to equation (5.13).

The measurements were then processed using a slightly modified version of the AM algorithm described in equation (5.17) to accommodate an experimental non-ideality. Since the LEDs are not perfect point sources, the interference fringes are filtered by a Gaussian-shaped window, as shown in figure 5.6(a), that essentially places more weight on the lower spatial frequencies. This non-ideality can be accounted for by modifying the problem to be solved in equation (5.16) to

$$[x_e, P_e] = \arg\min_{x,P} ||g - APFx||_2^2, \tag{5.18}$$

where A is a diagonal matrix whose elements vectorize a 2D Gaussian. The iterative, two-step AM algorithm is modified from equation (5.17) to

$$\begin{aligned} x_e^i &= \arg\min_x \frac{1}{2}||g - \Phi x||_2^2 + \lambda||x||_1 \\ &= \arg\min_x \frac{1}{2}||g - AP_e^{i-1}Fx||_2^2 + \lambda||x||_1 \end{aligned} \tag{5.19a}$$

$$P_e^i = \text{sgn}\left[\left(AFx_e^i\right)^* \cdot g\right]. \tag{5.19b}$$

Figure 5.6(a) shows the absolute value of the 2D subset of samples of J measured using the RSI. Figure 5.6(b) shows a distorted image of the scene recovered through an inverse Fourier transformation based on the VCZ theorem in equation (5.10), which does not account for turbulence. Figures 5.6(c) and (d) show discrete estimates of the undistorted scene, x_e, and the phase screen, $arg\{P_e\}$, recovered using the algorithm in equation (5.19). The recovered image of the 3 LEDs is a clear improvement over the aberrated image in figure 5.6(b).

Figure 5.7 shows the results of a control experiment where the phase distortion between the LEDs and the RSI was removed. Figure 5.7(a) shows the absolute value of the 2D subset of samples of J measured using the RSI. Figure 5.7(b) shows an image of the scene recovered using the VCZ theorem. Figures 5.7(c) and (d) show

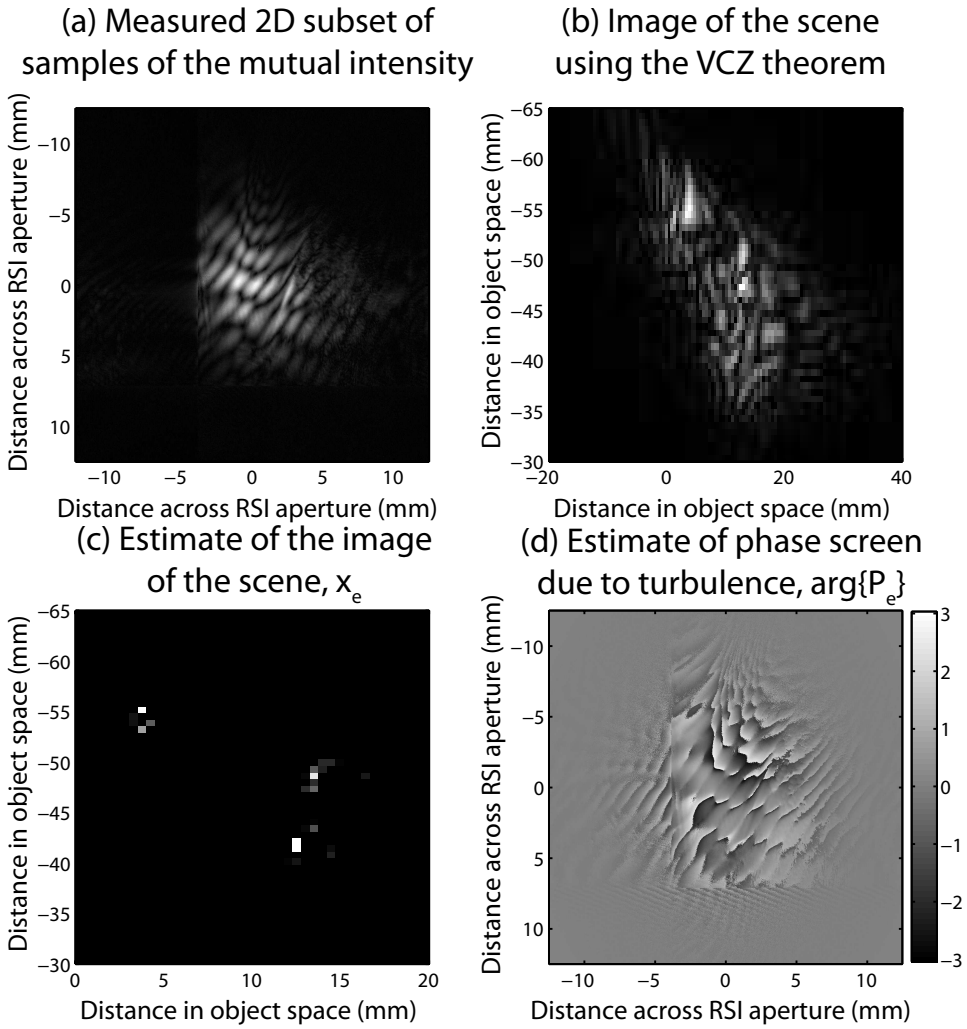

FIGURE 5.6: Results of using the two step AM algorithm described in equation (5.19) to image three LEDs aberrated by turbulence. (a) shows the absolute value of the 2D subset of samples of J measured using the RSI. (b) shows a distorted image of the scene recovered through direct application of the VCZ theorem, which does not account for turbulence. (c) and (d) show discrete estimates of the undistorted scene, x_e, and the phase screen, $arg\{P_e\}$, recovered using the AM algorithm.

discrete estimates of the scene, x_e, and the phase screen, $arg\{P_e\}$, recovered using the algorithm in equation (5.19). With the sparsity constraint on the scene, the algorithm sharpens the images of the 3 LEDs in the scene. As expected, the phase screen distribution is uniform due to the lack of turbulence. The control experiment suggests that the algorithm in equation (5.19) does not merely deconvolve the image,

but actually estimates the phase screen and applies it to the sampled J to produce an improved image.

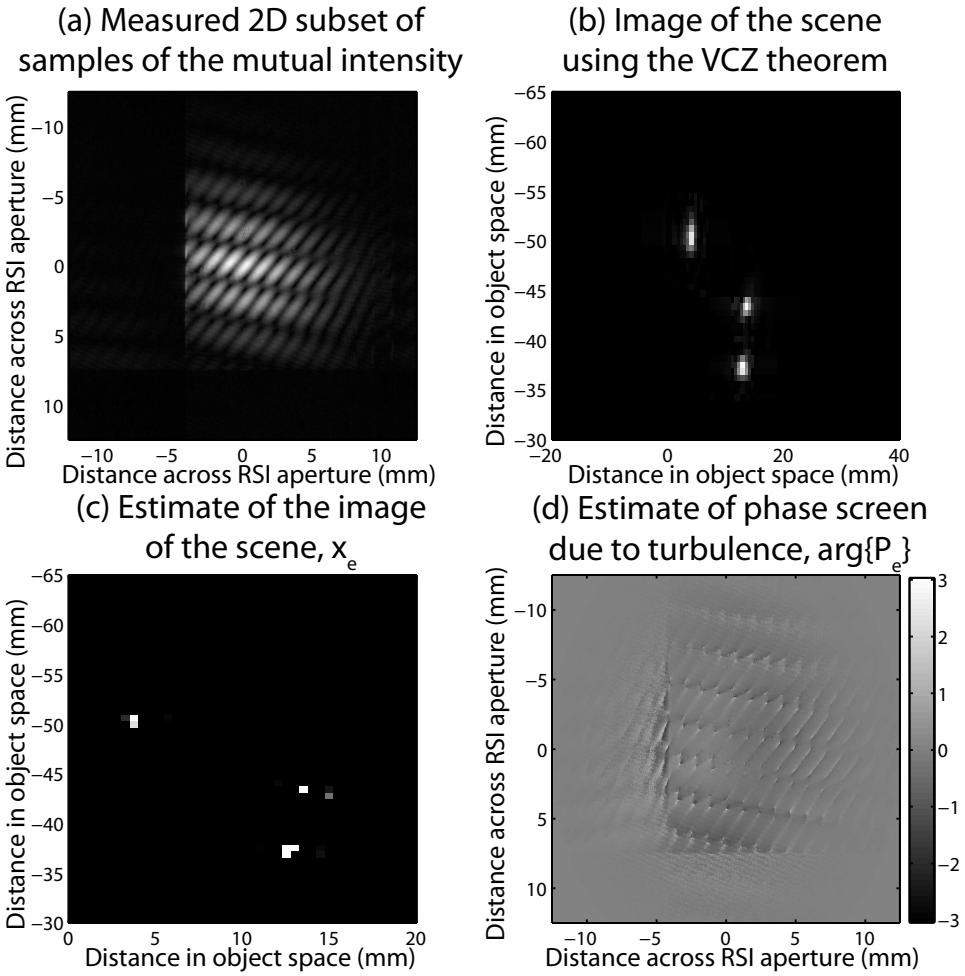

FIGURE 5.7: Results of a control experiment for RSI imaging of three LEDs not aberrated by turbulence. (a) shows the absolute value of the 2D subset of samples of J measured using the RSI. (b) shown an image of the scene recovered using the VCZ theorem. (c) and (d) show discrete estimates of the undistorted scene, x_e, and the phase screen $arg\{P_e\}$ recovered using the AM algorithm.

Comparing figures 5.6(c) and 5.7(c) illustrates that the absolute positions of the LEDs is not recovered. This is because the algorithm is unable to recover the tilt component of the phase screen that leads to a relative shift in the position of the objects with respect to their true positions.

5.9 Imaging extended objects through turbulence by compressively sampling the mutual intensity

In the previous section, experimental verification was provided for the RSI's ability to image point sources through turbulence with the assumption of spatial sparsity of the scene. Subsequently, a noiseless Matlab simulation was performed to investigate the potential to image an extended object through turbulence with an assumption of sparsity in the spatial gradient of the scene. For this case, the iterative, two-step AM algorithm is modified from equation (5.17) to

$$\begin{aligned} x_e^i &= \arg\min_x \frac{1}{2}||g - \Phi x||_2^2 + \lambda||x||_{TV} \\ &= \arg\min_x \frac{1}{2}||g - AP_e^{i-1}Fx||_2^2 + \lambda||x||_{TV} \end{aligned} \tag{5.20a}$$

$$P_e^i = \text{sgn}\left[\left(AFx_e^i\right)^* \cdot g\right]. \tag{5.20b}$$

In the simulation, the object shown in figure 5.8(a) was imaged by the RSI through the phase screen shown in 5.8(b). The object consists of a face with a small box located in the top right corner with a different intensity. The phase screen was generated based on McGlamery's algorithm [73], which defines the phase screen statistically by means of the Kolmogorov spectrum [50].

Figure 5.9(a) shows the absolute value of the 2D subset of samples of J measured using the RSI (scaled for contrast). Figure 5.9(b) shows a distorted image of the scene recovered through an inverse Fourier transformation based on the VCZ theorem in equation (5.10), which does not account for turbulence. While the smile and eyes on the object are somewhat visible, the small box in the top right corner of the object is completely buried in the turbulence. Figures 5.9(c) and (d) show discrete estimates of the scene, x_e, and the phase screen, $arg\{P_e\}$, recovered using the algorithm in equation (5.20). The recovered image of the face and the small box

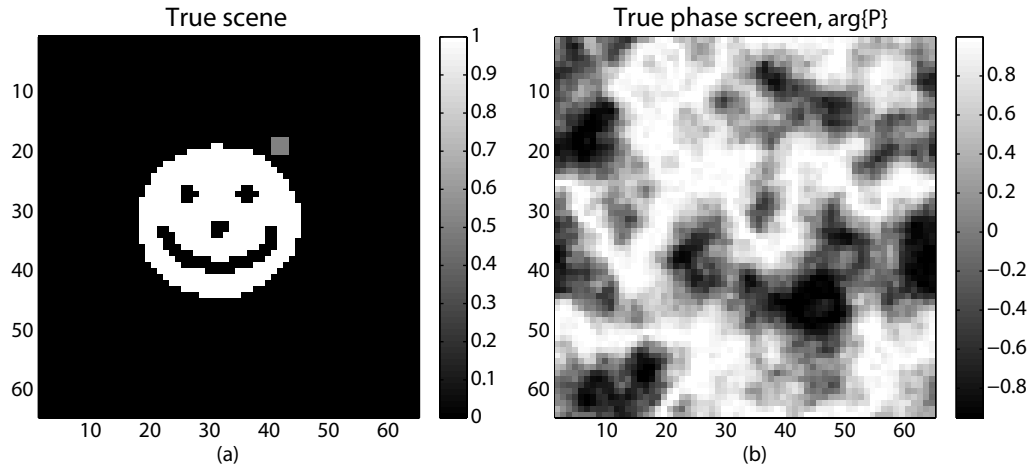

FIGURE 5.8: (a) An extended object to be imaged by the RSI in simulation. (b) Distribution of the simulated phase screen across the RSI aperture. The phase screen was generated based on McGlamery's algorithm, which defines the phase screen statistically by means of the Kolmogorov spectrum.

is a clear improvement over the aberrated image in figure 5.9(b).

Figure 5.10(a) shows the true phase screen, while figure 5.10(b) shows the phase screen estimated using the algorithm described in equation (5.20). Figure 5.10(c) shows the difference between the true and estimated phase screens. The difference map demonstrates that the algorithm is able to recover the true phase screen up to a tilt component from the bottom left to the top right. The lack of tilt correction is compensated by a shift in the position of the reconstructed extended object up and to the right from its true position shown in figure 5.8(a). The estimated object in figure 5.9(c) and the estimated phase screen in figure 5.10(b) are a valid solution, as they combine to produce almost the same measurements as the combination of the true object and phase screen shown in figures 5.8(a) and 5.8(b).

This simulation suggests that the RSI may be able to image extended objects through turbulence. However, the SNR of the RSI is known to degrade as a function of the number of independent sources in the scene. This is because, unlike in a focal imaging design, every source contributes light to every pixel on the detector and

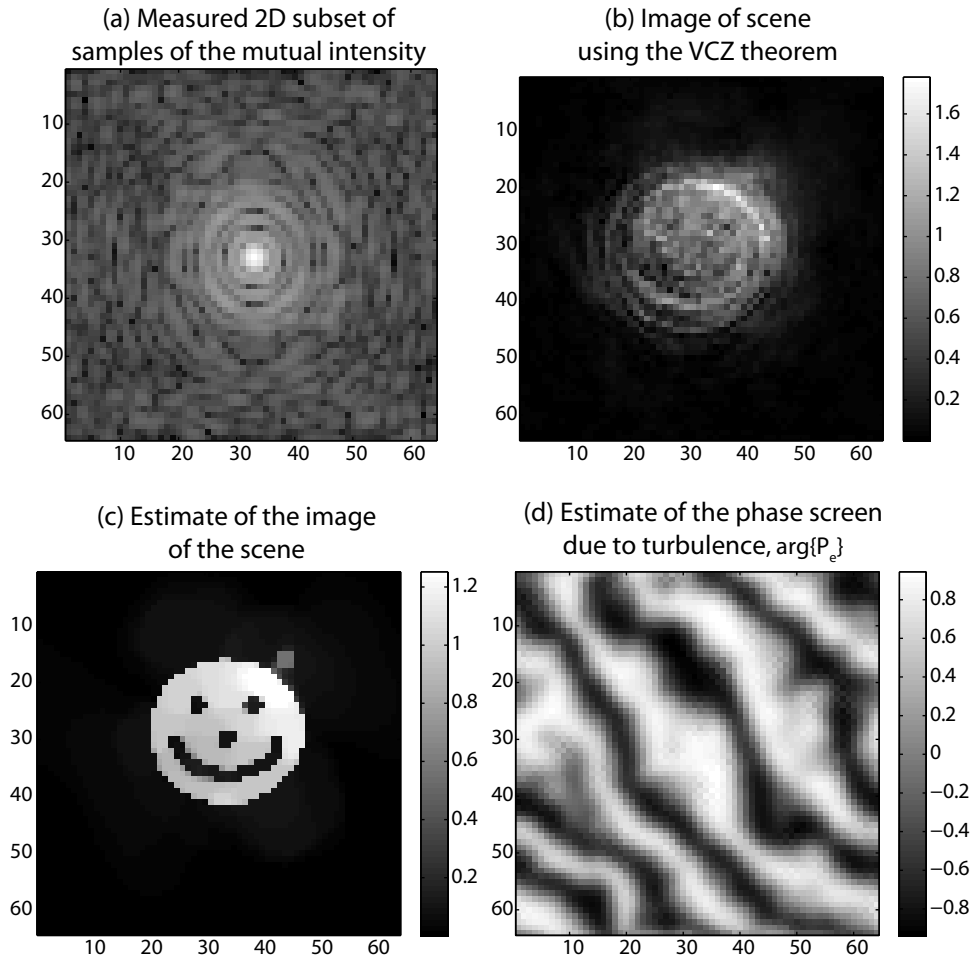

FIGURE 5.9: Noiseless simulation to demonstrate the imaging of an extended object through turbulence using the RSI and the AM algorithm described by equation (5.20). (a) shows the absolute value of the 2D subset of samples of J measured using the RSI. (b) shown an image of the scene recovered using the VCZ theorem. The face is barely recognizable and the small box in the top right corner can no longer be discerned. (c) and (d) show discrete estimates of the undistorted scene, x_e, and the phase screen $arg\{P_e\}$ recovered using the AM algorithm.

the actual signal is a weak modulation of the intensity on a high background [74]. This may make imaging of extended sources by measuring samples of J on the RSI practically challenging.

However, the proposed approach can be a practical means of imaging remote astronomical objects that are close together through turbulence if the complex 2D subset of 4D J can be sampled in a snapshot. This can potentially be achieved

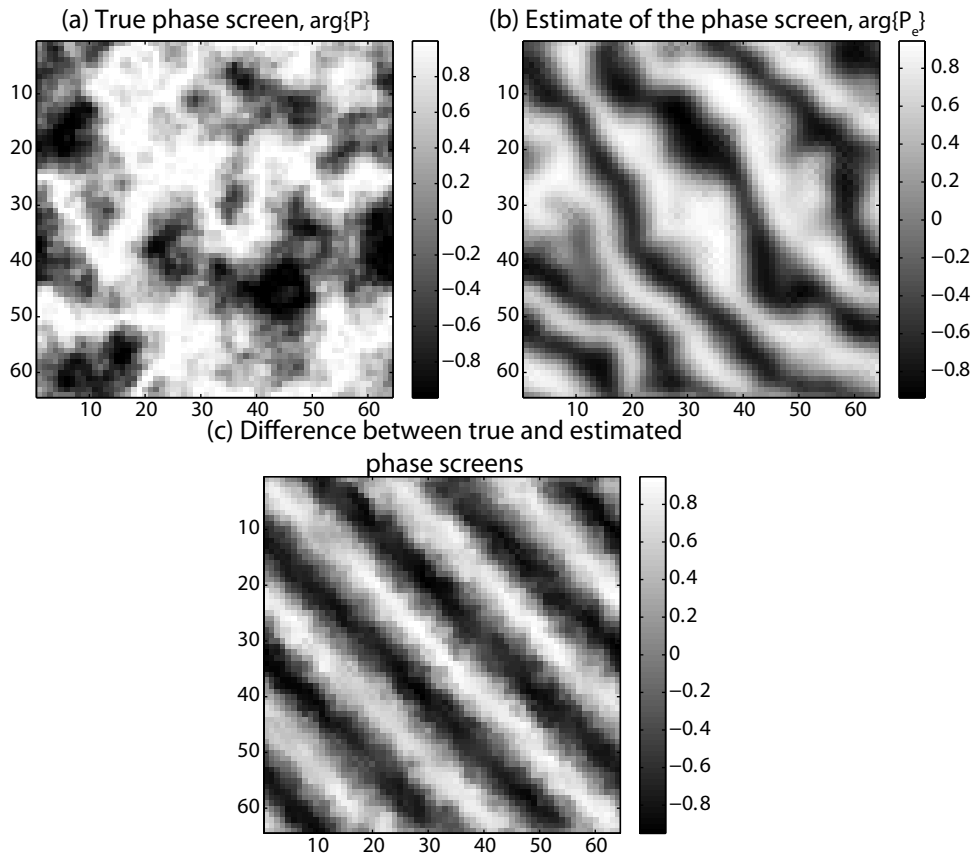

FIGURE 5.10: Comparison of the estimated phase screen with the true phase screen. The difference map demonstrates that the algorithm is able to recover the true phase screen up to a tilt component. The lack of tilt correction is compensated by a shift in the absolute position of the reconstructed extended object.

using the quadrature phase interferometer [75], which measures four phase-shifted interferograms in quadrature across the entire pupil in a single exposure, rather than over four separate exposures.

5.10 Conclusions

In this chapter, CASSI was briefly considered as a means of imaging through turbulence by exploiting the chromatic dispersion of the turbulent atmosphere to produce wavelength diversity of the PSF. Unfortunately, the atmosphere induces negligible wavelength diversity of the PSF.

Thus, an approach to image objects through turbulence by measurement of the power spectral density of the optical field was no longer considered. Instead, a rotational shear interferometer was used to image a sparse scene through turbulence by measuring a 2D subset of the complex, 4D mutual intensity of the optical field. This measurement was processed using an alternating minimization algorithm to jointly estimate the sparse scene as well the turbulence-induced phase screen. The sparsity of the scene was enforced with sparsity constraints similar to those discussed in the previous chapters.

Experimental results demonstrated the feasibility of the algorithm for imaging a sparse set of point sources through an isoplanatic phase distortion. Noiseless simulations also demonstrated the feasibility of the algorithm to image extended objects sparse on some basis through turbulence.

6

Summary and open questions

In this dissertation, the coded aperture snapshot spectral imager (CASSI) and the rotational shear interferometer (RSI) were demonstrated as physical examples of systems that perform generalized sampling of the optical field. Both sensors were used to recover particular signals of interest through an integrated approach to optical and algorithm design.

In a snapshot, CASSI measured a linear, 2D coded projection of the power spectral density of the optical field, a 3D datacube. The coding was physically implemented in optical hardware. In particular, a coded aperture was used to introduce spatial coding and a prism was used to introduce spectral coding. The measurements were digitally processed to recover the datacube. In general, inversion of the snapshot projection for recovery of the datacube is an ill-posed problem because the number of measurements made is significantly smaller than the number of voxels in the datacube. Thus, any numerical estimation algorithm used for datacube recovery had to exploit properties or structures in the datacube known *a priori* that helped to constrain the estimated datacube to a certain solution space.

CASSI's potential for datacube recovery in a snapshot was experimentally demonstrated in chapters 2 and 3. The first CASSI prototype, presented in chapter 2, was built using off-the-shelf components. The second prototype was custom-designed in Zemax ray tracing software to achieve significantly better imaging quality. Three different iterative numerical estimation algorithms were investigated for datacube recovery. The GPSR and TwIST methods solved the linear inversion problem by reformulating it into a non-linear convex unconstrained optimization problem with a regularizing penalty term in the objective function that encouraged sparsity of the datacube in a sparsifying basis. In contrast to these methods, the NeAREst did not require the specification of a regularization parameter and used the Richardson-Lucy algorithm to ensure that the estimated solution would be non-negative. NeAREst iterations were also computationally cheaper than GPSR and TwIST iterations.

In chapter 4, CASSI was examined from the perspective of compressive sampling theory. Noiseless Matlab simulations illustrated that even with a constraint on the number of sparse basis coefficients in the scene to be imaged, the performance of CASSI was scene dependent, an undesirable property for CASSI's application for mainstream spectral imaging. The scene dependence was attributed to lack of orthogonality of the columns of the CASSI sensing matrix. An alternative design was proposed with better column properties that was less scene dependent but was more sensitive to noise in the measurements.

As expected from any new body of research, there are some open questions and immediate extensions to the work on CASSI. CASSI was an example of a system that attempted to physically implement a compressively sampled measurement of a signal in a snapshot. In terms of sensing hardware, the discussion about properties of the columns of the CASSI sensing matrix presented in chapter 4 may stimulate various new designs of physical compressive measurement systems. In terms of numerical estimation methods used to process CASSI data, it would be interesting to study

the scene dependence of CASSI when the NeAREst algorithm is used. Alternative methods should consider the non-negativity of the datacube to be reconstructed as well as the fact that incoherent imaging is often impacted by Poisson noise. Recent work on sparse Poisson intensity reconstruction algorithms [76] may find application here.

CASSI was briefly considered as a means of imaging through turbulence by exploiting the chromatic dispersion of the aberrating medium to produce wavelength diversity of the PSF. However, a medium like the Earth's atmosphere produces negligible wavelength diversity in the PSF. Furthermore, a datacube of an object to be imaged through turbulence and illuminated with broadband light would not be sparse. Thus, CASSI was no longer considered as a means of imaging through turbulence. However, a related idea that is worthy of further investigation involves imaging through turbulence by using an RGB camera (a 3-channel spectral imager) with a chromatic lens to increase the wavelength diversity of the PSF.

As an alternative means of imaging through turbulence, compressive measurement of the 4D mutual intensity, a coherence function, was performed using an RSI, as described in chapter 5. This was not a conventional compressive measurement in the sense that it could not be analyzed from the perspective of compressive sampling theory. The effective sensing matrix was unknown in the presence of turbulence, making recovery of an undistorted image of the scene along with the distribution of the phase screen a non-convex problem.

Prior to its use in the work described in this dissertation, the RSI had been used for extended depth of field microscopic imaging [77] as well as tomographic imaging [78] of incoherent sources by using the van Cittert Zernike theorem on samples of the mutual intensity as measured by the RSI. In the presence of a turbulence-induced phase screen between the incoherent sources and the RSI, direct application of the van Cittert Zernike theorem to the RSI measurements would produce a distorted

image of the sources. However, with the constraint that the incoherent sources in the scene are sparse in some basis, an alternating minimization method was presented in chapter 5 to successfully recover an undistorted image of a sparse scene as well as the distribution of the turbulence-induced phase screen. In principle, this approach does not require capture of several images and post-processing of the measured data in batch mode. Thus, this approach may find utility in real-time astronomical imaging of distant stars through the turbulence of Earth's atmosphere.

Just like CASSI, there are several open questions that may be asked as extensions to the RSI approach for imaging through turbulence. An immediate extension of the experimental results described in chapter 5 would be an attempt to image astronomical point sources through turbulence. A non-negativity constraint should be placed on the scene. The potential of the alternating minimization approach for imaging extended sources from RSI measurements in the presence of noise should be investigated.

Bibliography

[1] D. J. Brady, *Optical Imaging and Spectroscopy*. New York: Wiley, 2009.

[2] J. Mait, R. Athale, and J. van der Gracht, "Evolutionary paths in imaging and recent trends," *Opt. Express*, vol. 11, no. 18, pp. 2093–2101, 2003.

[3] M. E. Gehm, R. John, D. J. Brady, R. M. Willett, and T. J. Schulz, "Single-shot compressive spectral imaging with a dual-disperser architecture," *Opt. Express*, vol. 15, no. 21, pp. 14013–14027, 2007.

[4] A. Wagadarikar, R. John, R. Willett, and D. Brady, "Single disperser design for compressive, single-snapshot spectral imaging," vol. 6714 of *Proc. SPIE*, 2007.

[5] A. Wagadarikar, R. John, R. Willett, and D. J. Brady, "Single disperser design for coded aperture snapshot spectral imaging," *Applied Optics*, vol. 47, no. 10, pp. B44–B51, 2008.

[6] A. Wagadarikar, N. Pitsianis, X. Sun, and D. Brady, "Spectral image estimation for coded aperture snapshot spectral imagers," vol. 7076 of *Proc. SPIE*, 2008.

[7] A. Wagadarikar, N. Pitsianis, X. Sun, and D. Brady, "Video rate spectral imaging using a coded aperture snapshot spectral imager," *Optics Express*, vol. 17, no. 8, pp. 6368–6388, 2009.

[8] A. A. Wagadarikar, D. Marks, K. Choi, and D. J. Brady, "Compressive coherence sensing," in *Computational Optical Sensing and Imaging*, p. CTuA6, Optical Society of America, 2009.

[9] A. A. Wagadarikar, D. Marks, K. Choi, R. Horisaki, and D. J. Brady, "Imaging through turbulence using compressive coherence sensing (to appear)," *Optics Letters*, 2010.

[10] R. Levenson and C. C. Hoyt, "Spectral imaging and microscopy," *American Laboratory*, vol. 32, no. 22, pp. 26–33, 2000.

[11] R. Lukac and K. Plataniotis, *Color Image Processing: Methods and Applications*. Boca Raton: CRC Press, 2006.

[12] R. Green, M. L. Eastwood, C. M. Sarturea, T. G. Chriena, M. Aronssona, B. J. Chippendalea, J. A. Fausta, B. E. Pavria, C. J. Chovita, M. Solisa, M. R. Olaha, and O. Williams, "Imaging Spectroscopy and the Airborne Visible/Infrared Imaging Spectrometer (AVIRIS)," *Remote Sensing of Environment*, vol. 65, no. 3, pp. 227–248, 1998.

[13] M. R. Descour and E. L. Dereniak, "Nonscanning no-moving-parts imaging spectrometer," in *Imaging Spectrometry*, vol. 2480, pp. 48–64, Proc. SPIE, 1995.

[14] M. E. Gehm, S. T. McCain, N. P. Pitsianis, D. J. Brady, P. Potuluri, and M. E. Sullivan, "Static two-dimensional aperture coding for multimodal multiplex spectroscopy," *Applied Optics*, vol. 45(13), pp. 2857–3183, 2006.

[15] A. A. Wagadarikar, M. E. Gehm, and D. J. Brady, "Performance comparison of aperture codes for multimodal, multiplex spectroscopy," *Applied Optics*, vol. 46, no. 22, pp. 4932–4942, 2007.

[16] M. Gehm, M. Kim, C. Fernandez, and D. Brady, "High-throughput, multiplexed pushbroom hyperspectral microscopy," *Optics Express*, vol. 16, no. 15, pp. 11032–11043, 2008.

[17] J. M. Mooney, V. E. Vickers, M. An, and A. K. Brodzik, "High-throughput hyperspectral infrared camera," *Journal of the Optical Society of America A*, vol. 14, no. 11, pp. 2951–2961, 1997.

[18] D. Brady and M. Gehm, "Compressive imaging spectrometers using coded apertures," vol. 6246, Proc. SPIE, 2006.

[19] A. Harvey and D. Fletcher-Holmes, "Spectral imaging in a snapshot," vol. 5694, Proc. SPIE, 2005.

[20] H. Morris, C. Hoyt, and P. Treado, "Imaging spectrometers for fluorescence and Raman microscopy: acousto-optic and liquid-crystal-tunable filters," *Applied Spectroscopy*, vol. 48, pp. 857–866, 1994.

[21] C. Volin, B. Ford, M. Descour, J. Garcia, D. Wilson, P. Maker, and G. Bearman, "High-speed spectral imager for imaging transient fluorescent phenomena," *Applied Optics*, vol. 37, no. 34, pp. 8112–8119, 1998.

[22] K. Hege, D. O'Connell, W. Johnson, S. Basty, and E. Dereniak, "Hyperspectral imaging for astronomy and space surveillance," vol. 5159 of *Proc. SPIE*, pp. 380–391, 2003.

[23] W. Johnson, D. Wilson, W. Fink, M. Humayun, and G. Bearman, "Snapshot hyperspectral imaging in ophthalmology," *Journal of Biomedical Optics*, vol. 12, no. 1, 2007.

[24] N. Gat, G. Scriven, J. Garman, M. D. Li, and J. Zhang, "Development of four-dimensional imaging spectrometers (4D-IS)," vol. 6302 of *Proc. SPIE*, 2006.

[25] L. Gao, R. T. Kester, and T. S. Tkaczyk, "Compact image slicing spectrometer (ISS) for hyperspectral fluorescence microscopy," *Opt. Express*, vol. 17, no. 15, pp. 12293–12308, 2009.

[26] M. Descour, C. Volin, E. Dereniak, K. Thorne, A. Schumacher, D. Wilson, and P. Maker, "Demonstration of a high-speed nonscanning imaging spectrometer," *Optics Letters*, vol. 22, no. 16, pp. 1271–1273, 1997.

[27] W. R. Johnson, D. W. Wilson, and G. Bearman, "Spatial-spectral modulating snapshot hyperspectral imager," *Applied Optics*, vol. 45, no. 9, pp. 1898–1908, 2006.

[28] M. Descour and E. Dereniak, "Computed tomography imaging spectrometer: experimental calibration and reconstruction results," *Appl. Opt.*, vol. 34, pp. 4817–4826, 1995.

[29] A. D. Portnoy, *Coded measurement for imaging and spectroscopy.* PhD thesis, Duke University.

[30] M. Harwit and N. Sloane, *Hadamard Transform Optics.* New York, USA: Academic Press, 1979.

[31] J. M. Lerner, "Imaging spectrometer fundamentals for researchers in the biosciences - a tutorial," *Cytometry Part A*, vol. 69, no. 8, pp. 712–734, 2006.

[32] M. A. Figueiredo, R. D. Nowak, and S. J. Wright, "Gradient projection for sparse reconstruction: application to compressed sensing and other inverse problems," *IEEE Journal of Selected Topics in Signal Processing: Special Issue on Convex Optimization Methods for Signal Processing*, vol. 1, no. 4, pp. 586–597, 2007.

[33] J. Bioucas-Dias and M. Figueiredo, "A new TwIST: two-step iterative shrinkage/thresholding for image restoration," *IEEE Transactions on Image Processing*, vol. 16, no. 12, pp. 2992–3004, 2007.

[34] X. Sun and N. P. Pitsianis, "Solving non-negative linear inverse problems with the nearest method," vol. 7074, p. 707402, SPIE, 2008.

[35] R. Blahut, *Theory of Remote Image Formation*. Cambridge, UK: Cambridge University Press, 2004.

[36] E. J. Candès, "Compressive sampling," Proceedings of the International Congress of Mathematicians, 2006.

[37] R. Baraniuk, M. Davenport, R. DeVore, and M. Wakin, "A simple proof of the restricted isometry property for random matrices," *Constructive Approximation*, vol. 2008, 2007.

[38] R. G. Baraniuk, "Compressive sensing," *IEEE Signal Processing Magazine*, pp. 118–124, July 2007.

[39] E. Candès, J. Romberg, and T. Tao, "Stable signal recovery from incomplete and inaccurate measurements," *Communications on Pure and Applied Mathematics*, vol. 59, pp. 1207–1223, 2006.

[40] R. Gribonval and M. Nielsen, "Sparse representations in unions of bases," *IEEE Transactions on Information Theory*, vol. 49, no. 12, pp. 3320–3325, 2003.

[41] E. J. Candès and T. Tao, "Near-optimal signal recovery from random projections: Universal encoding strategies?," *IEEE Trans. Inform. Theory*, vol. 52, no. 12, pp. 5406–5425, 2006.

[42] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, March 2004.

[43] E. Candès, "The restricted isometry property and its implications for compressed sensing," *Comptes Rendus Mathematique*, vol. 346, no. 9-10, pp. 589–592, 2008.

[44] E. J. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Inform. Theory*, vol. 52, no. 2, pp. 489–509, 2006.

[45] C. Fernandez, A. Wagadarikar, D. Brady, S. McCain, and T. Oliver, "Fluorescence microscopy with a coded aperture snapshot spectral imager," vol. 7184, p. 71840Z, SPIE, 2009.

[46] C. Fernandez-Cull, K. Choi, D. Brady, and T. Oliver, "Identification of fluorescent beads using a coded aperture snapshot spectral imager," *Applied Optics (to appear)*, April 2010.

[47] M. Roggemann, B. Welsh, and R. Fugate, "Improving the resolution of ground-based telescopes," *Review of Modern Physics*, vol. 69, no. 2, pp. 437–505, 1997.

[48] R. Johnston, C. Worley, J. Mohr, R. Lane, and P. Cottrell, "Turbulence monitoring at Mount John," Proceedings of Image and Vision Computing New Zealand, pp. 497–492, 2005.

[49] F. Brunner, "The effects of atmospheric turbulence on telescopic observations," *Bulletin Geodesique*, vol. 56, no. 4, pp. 341–355, 1982.

[50] M. Roggemann and R. Welsh, *Color Image Processing: Methods and Applications*. Boca Raton, FL: CRC Press, 1996.

[51] A. Labeyrie, "Attainment of diffraction limited resolution in large telescopes by fourier analysing speckle patterns in star images," *Astronomy and Astrophyics*, vol. 6, pp. 85–87, 1970.

[52] T. Schulz, "Multiframe blind deconvolution of astronomical images," *Journal of the Optical Society of America A*, vol. 10, no. 5, pp. 1064–1073, 1993.

[53] J. Baldwin, R. Tubbs, G. Cox, C. Mackay, R. Wilson, and M. Anderson, "Diffraction-limited 800 nm imaging with the 2.56 m nordic optical telescope," *Astronomy and Astrophyics*, vol. 368, pp. L1–L4, 2001.

[54] P. L. Wizinowich, D. Le Mignant, A. H. Bouchez, R. D. Campbell, J. C. Y. Chin, A. R. Contos, M. A. Van Dam, S. K. Hartman, E. M. Johansson, R. E. Lafon, H. Lewis, P. J. Stomski, D. M. Summers, C. G. Brown, P. M. Danforth, C. E. Max, and D. M. Pennington, "The W.M. Keck observatory laser guide star adaptive optics system: Overview," *Publications of the Astronomical Society of the Pacific*, vol. 118, pp. 297–309, 2006.

[55] R. Lane and M. Tallon, "Wave-front reconstruction using a Shack-Hartmann sensor," *Applied Optics*, vol. 31, no. 32, pp. 6902–6908, 1992.

[56] J. Primot, G. Rousset, and J. Fontanella, "Deconvolution from wave-front sensing: a new technique for compensating turbulence-degraded images," *J. Opt. Soc. Am. A*, vol. 7, no. 9, pp. 1598–1608, 1990.

[57] R. Fugate, D. Fried, G. Ameer, B. Boeke, S. Browne, P. Roberts, R. Ruane, G. Tyler, and L. Wopat, "Measurement of atmospheric wavefront distortion using scattered light from a laser guide-star," *Nature*, vol. 353, pp. 144–146, 1991.

[58] R. Kendrick, D. Acton, and A. Duncan, "Phase-diversity wave-front sensor for imaging systems," *Applied Optics*, vol. 33, no. 27, pp. 6533–6546, 1994.

[59] J. Dolne and H. Schall, "Cramer-Rao bound and phase-diversity blind deconvolution performance versus diversity polynomials," *Applied Optics*, vol. 44, no. 29, pp. 6220–6227, 2005.

[60] S. Jefferies, M. Lloyd-Hart, E. Hege, and J. Georges, "Sensing wave-front amplitude and phase with phase diversity," *Applied Optics*.

[61] D. Marks, R. Stack, and D. Brady, "Three dimensional coherence imaging in the Fresnel domain," *Applied Optics*, vol. 38, no. 8, pp. 1332–1342, 1999.

[62] T. Pearson and A. Readhead, "Image formation by self-calibration in radio astronomy," *Annual Reviews of Astronomy and Astrophysics*, vol. 22, pp. 97–130, 1984.

[63] V. Uvarov, "Astigmatic coherence sensor for digital imaging," *Radiophysics and Quantum Electronics*, vol. 31, no. 1, pp. 13–19, 1988.

[64] D. Rogstad, "A technique for measuring visibility phase with an optical interferometer in the presence of atmospheric seeing," *Applied Optics*, vol. 7, no. 4, pp. 585–588, 1968.

[65] R. G. Paxman, T. J. Schulz, and J. R. Fienup, "Joint estimation of object and aberrations using phase diversity," *J. Opt. Soc. Am. A*, vol. 9, pp. 1072–1085, 1992.

[66] H. Ingleby and D. McGaughey, "Real data results with wavelength-diverse blind deconvolution," *Optics Letters*, vol. 30, no. 5, pp. 489–491, 2004.

[67] A. Glindemann, S. Hippler, T. Berkefeld, and W. Hackenberg, "Adaptive optics on large telescopes," *Experimental Astronomy*, vol. 20, no. 1, pp. 5–48, 1999.

[68] D. Marks, R. Stack, and D. Brady, "Astigmatic coherence sensor for digital imaging," *Optics Letters*, vol. 25, no. 23, pp. 1726–1728, 2000.

[69] D. Marks, R. Stack, and D. Brady, "Digital refraction distortion correction with an astigmatic coherence sensor," *Applied Optics*, vol. 41, no. 29, pp. 6050–6054, 2002.

[70] M. Lustig, D. Donoho, and J. Pauly, "Sparse MRI: The application of compressed sensing for rapid MR imaging," *Magnetic Resonance in Medicine*, vol. 58, no. 6, pp. 1182–1195, 2007.

[71] Y. Wiaux, L. Jacques, G. Puy, A. Scaife, and P. Vandergheynst, "Compressed sensing imaging techniques for radio interferometry," *Monthly Notices of the Royal Astronomical Society*, vol. 395, pp. 1733–1742, 2009.

[72] J. Gallicchio, "Spatio-spectral triangulation of visible and infrared point sources using a portable rotational shear interferometer," Master's thesis, University of Illinois at Urbana-Champaign, Illinois.

[73] B. McGlamery, "Computer simulation studies of compensation of turbulence degraded images," vol. 76 of *Proc. SPIE*, pp. 225–233, 1976.

[74] D. Marks, *Four-dimensional coherence sensing*. PhD thesis, University of Illinois at Urbana-Champaign, Illinois.

[75] B. Kern, P. Dimotakis, C. Martin, D. Lang, and R. Thessin, "Imaging through turbulence with a quadrature-phase optical interferometer," *Applied Optics*, vol. 44, pp. 7424–7438, 2005.

[76] Z. Harmany, R. Marcia, and R. Willett, "Spiral out of convexity: sparsity-regularized algorithms for photon limited imaging," vol. 7533 of *Proc. SPIE*, p. 75330R, 2010.

[77] P. Potuluri, M. Fetterman, and D. Brady, "High depth of field microscopic imaging using an interferometric camera," *Optics Express*, vol. 8, no. 11, pp. 624–630, 2001.

[78] D. Marks, R. Stack, D. Brady, D. M. Jr., and R. Brady, "Visible cone beam tomography with a lensless interferometric camera," *Science*, vol. 284, pp. 2164–2166, 2001.

Biography

Ashwin Wagadarikar was born in Kalyan, India on January 19, 1984. He received a Bachelor of Applied Science in Engineering Science (Biomedical Option) at the University of Toronto and graduated with honours in 2005. In the summers of 2004 and 2005, he investigated the biomechanics of the optics nerve head under the guidance of Dr. Ross Ethier and Dr. Ian Sigal. His work was supported by the Natural Sciences and Engineering Research Council (NSERC) of Canada. Ashwin began pursuing his PhD in Electrical & Computer Engineering at Duke University in August 2005. There he became a member of the Duke Imaging and Spectroscopy Program (DISP) led by Dr. David Brady. He received a Master of Science in Electrical & Computer Engineering in 2007 after defending his thesis entitled "Performance comparison of aperture codes for multimodal, multiplex spectroscopy". He continued his research in computational imaging for the duration of his PhD. While at Duke, he was awarded a prestigious 3 year doctoral fellowship from NSERC in 2008, as well as an SPIE educational scholarship in 2009. Since 2008, he has served as a reviewer for the journals *Applied Optics* and *Optics Letters*.

Ashwin was actively involved in the Optical Society of America (OSA) / SPIE student chapter at Duke University, serving as Secretary ($2006 - 2007$), Outreach Coordinator ($2007 - 2008$) and President ($2008 - 2009$) of the chapter. During his term as President, he was interviewed about the group's activities by Photonics Spectra magazine in February 2009.

Publications

1. **Ashwin Wagadarikar**, Daniel Marks, Kerkil Choi, Ryoichi Horisaki, and David Brady, *"Imaging through turbulence using compressive coherence sensing,"* Optics Letters 35 (6), 838 – 840 (2010)

2. Kerkil Choi, **Ashwin Wagadarikar**, Ryoichi Horisaki, and David Brady, *"Compressive coded aperture computed tomography - Part I: Coded Aperture Snapshot Spectral Imager (CASSI),"* Submitted to IEEE Transactions on Image Processing (July 2009)

3. Kerkil Choi, Ryoichi Horisaki, **Ashwin Wagadarikar**, and David Brady, *"Compressive coded aperture computed tomography - Part II: Compressive X-ray Tomography,"* Submitted to IEEE Transactions on Image Processing (July 2009)

4. **Ashwin Wagadarikar**, Nikos Pitsianis, Xiaobai Sun, and David Brady, *"Video rate spectral imaging using a coded aperture snapshot spectral imager,"* Optics Express 17 (8), 6368 – 6388 (2009)

5. **Ashwin Wagadarikar**, Renu John, Rebecca Willett, and David Brady, *"Single disperser design for coded aperture snapshot spectral imaging,"* feature issue on Computational Optical Sensing and Imaging, Applied Optics 47 (10), B44 – 51 (2008). **(A figure from the paper appeared on the cover of this issue.)**

6. **Ashwin Wagadarikar**, Michael Gehm, and David Brady, *"Performance comparison of aperture codes for multimodal, multiplex spectroscopy,"* Applied Optics 46 (22), 4932 – 4942 (2007)

Conference Presentations

1. **Ashwin Wagadarikar**, Daniel Marks, Kerkil Choi, and David Brady, *"Compressive coherence sensing,"* at OSA topical meeting on Computational Optical

Sensing and Imaging (COSI) in San Jose, CA (2009)

2. Christy Fernandez, **Ashwin Wagadarikar**, David Brady, and Scott McCain, *"Fluorescence microscopy with a coded aperture snapshot spectral imager,"* in Synthetic Illumination and Synthetic Apertures I, Proc. SPIE 7184 (2009)

3. **Ashwin Wagadarikar**, Nikos Pitsianis, Xiaobai Sun, and David Brady, *"Spectral image reconstruction techniques for a coded aperture snapshot spectral imager,"* in Image Reconstruction from Incomplete Data V, Proc. SPIE 7076 (2008) **(Invited Talk)**

4. **Ashwin Wagadarikar**, Renu John, Rebecca Willett, and David Brady, *"Single disperser design for compressive, single-snapshot spectral imaging,"* in Adaptive Coded Aperture Imaging and Non-imaging Sensors, Proc. SPIE 6714 (2007)

5. **Ashwin Wagadarikar**, Michael Gehm, and David Brady, *"Evaluation of apertures codes for high throughput spectroscopy,"* at OSA Annual Meeting in Rochester, NY (2006)