Inverse Problems in Astronomical and General Imaging

by

Wun Ying Valerie Leung, B.E. (Hons I)

Department of Electrical and Computer Engineering

A thesis presented for the degree of
Doctor of Philosophy

University of Canterbury, August 2002

'Most people, if you describe a train of events to them, will tell you what the result would be. They can put those events together in their minds, and argue from them that something will come to pass. There are few people, however, who, if you told them a result, would be able to evolve from the own inner consciousness what the steps were which led up to that result. This power is what I mean when I talk of reasoning backwards, or analytically.' ... '... the grand thing is to be able to reason backwards.'

Sherlock Holmes in 'A study in Scarlet'

iii

Statement of Originality

This thesis contains no material which has been accepted for the award of any other degree or diploma in any tertiary institution. To the best of my knowledge and belief, the thesis contains no material previously published or written by another person, except when due reference is made in the text.

Valerie Leung

Abstract

The resolution and the quality of an imaged object are limited by four contributing factors. Firstly, the primary resolution limit of a system is imposed by the aperture of an instrument due to the effects of diffraction. Secondly, the finite sampling frequency, the finite measurement time and the mechanical limitations of the equipment also affect the resolution of the images captured. Thirdly, the images are corrupted by noise, a process inherent to all imaging systems. Finally, a turbulent imaging medium introduces random degradations to the signals before they are measured. In astronomical imaging, it is the atmosphere which distorts the wavefronts of the objects, severely limiting the resolution of the images captured by ground-based telescopes. These four factors affect all real imaging systems to varying degrees.

All the limitations imposed on an imaging system result in the need to deduce or reconstruct the underlying object distribution from the distorted measured data. This class of problems is called inverse problems. The key to the success of solving an inverse problem is the correct modelling of the physical processes which give rise to the corresponding forward problem. However, the physical processes have an infinite amount of information, but only a finite number of parameters can be used in the model. Information loss is therefore inevitable. As a result, the solution to many inverse problems requires additional information or prior knowledge. The application of prior information to inverse problems is a recurrent theme throughout this thesis.

An inverse problem that has been an active research area for many years is interpolation, and there exist numerous techniques for solving this problem. However, many of these techniques neither account for the sampling process of the instrument nor include prior information in the reconstruction. These factors are taken into account in the proposed optimal Bayesian interpolator. The process of interpolation is also examined from the point of view of superresolution, as these processes can be viewed as being complementary.

Since the principal effect of atmospheric turbulence on an incoming wavefront is a phase distortion, most of the inverse problem techniques devised for this seek to either estimate or compensate for this phase component. These techniques are classified into computer post-processing methods, adaptive optics (AO) and hybrid techniques.

Blind deconvolution is a post-processing technique which uses the speckle images to estimate both the object distribution and the point spread function (PSF), the latter of which is directly related to the phase. The most successful approaches are based on characterising the PSF as the aberrations over the aperture. Since the PSF is also dependent on the atmosphere, it is possible to constrain the solution using the statistics of the atmosphere. An investigation shows the feasibility of this approach. Bispectrum is also a post-processing method which reconstructs the spectrum of the object. The key component for phase preservation is the property of phase closure, and its application as prior information for blind deconvolution is examined.

Blind deconvolution techniques utilise only information in the image channel to estimate the phase which is difficult. An alternative method for phase estimation is from a Shack-Hartmann (SH) wavefront sensing channel. However, since phase information is present in both the wavefront sensing and the image channels simultaneously, both of these approaches suffer from the problem that phase information from only one channel is used. An improved estimate of the phase is achieved by a combination of these methods, ensuring that the phase estimation is made jointly from the data in both the image and the wavefront sensing measurements. This formulation, posed as a blind deconvolution framework, is investigated in this thesis. An additional advantage of this approach is that since speckle images are imaged in a narrowband, while wavefront sensing images are captured by a charge-coupled device (CCD) camera at all wavelengths, the splitting of the light does not compromise the light level for either channel. This provides a further incentive for using simultaneous data sets.

The effectiveness of using Shack-Hartmann wavefront sensing data for phase estimation relies on the accuracy of locating the data spots. The commonly used method which calculates the centre of gravity of the image is in fact prone to noise and is suboptimal. An improved method for spot location based on blind deconvolution is demonstrated.

Ground-based adaptive optics (AO) technologies aim to correct for atmospheric turbulence in real time. Although much success has been achieved, the space- and time-varying nature

of the atmosphere renders the accurate measurement of atmospheric properties difficult. It is therefore usual to perform additional post-processing on the AO data. As a result, some of the techniques developed in this thesis are applicable to adaptive optics.

One of the methods which utilise elements of both adaptive optics and post-processing is the hybrid technique of deconvolution from wavefront sensing (DWFS). Here, both the speckle images and the SH wavefront sensing data are used. The original proposal of DWFS is simple to implement but suffers from the problem where the magnitude of the object spectrum cannot be reconstructed accurately. The solution proposed for overcoming this is to use an additional set of reference star measurements. This however does not completely remove the original problem; in addition it introduces other difficulties associated with reference star measurements such as anisoplanatism and reduction of valuable observing time. In this thesis a parameterised solution is examined which removes the need for a reference star, as well as offering a potential to overcome the problem of estimating the magnitude of the object.

Acknowledgements

Contents

Preface

The research presented in this thesis is concerned with inverse problems in image processing, with a particular focus on astronomical imaging. This topic has been an active area of research for many years, and the aim is to obtain images with diffraction-limited (or even higher) resolution.

My interest in image processing emanated from my honours year project with Dr. Richard Lane. With this newly developed interest, and the lack of it for entering the work force, I started my Ph.D. studies with Dr. Lane in 1999.

As a lead in to my main research, I worked on an interpolation technique developed initially by Dr. Lane and Associate Professor Phil Bones, incorporating the sampling process and statistical information into the inverse problem. This led to a publication in Optical Engineering.

The principal area of my research is astronomical imaging. The main obstacle in astronomical imaging is the presence of the atmosphere between the object to be imaged and the ground-based telescope, which introduces a predominantly phase distortion to the incoming wavefront of the star. An obvious solution is to capture the images above the atmosphere as in the Hubble Space Telescope. Although this is a viable approach, it is expensive to operate and as a result access to it is regulated. Moreover, there is a limit to the size of the telescope that can be deployed in orbit. Ground-based methods are therefore by far the popular option.

There are three main approaches to overcoming the astronomical imaging problem: real-time adaptive optics, post-processing methods such as blind deconvolution, and hybrid methods which combine elements of the two. The main focus of my thesis was initially on blind deconvolution, which later widened to encompass hybrid methods and the processing of simultaneous wavefront sensing data.

Blind deconvolution is an ill-posed inverse problem and many existing techniques involve a maximum *a posteriori* solution to the problem, incorporating prior knowledge of the object and the point spread function. The progress of this research over the years has been the improvement in the quality of the prior information, as well as the parameterisation of the unknown quantities to constrain the solution. The prior information I examined includes the statistics of the atmosphere, and the phase closure property of the bispectrum method. Varying degrees of success have been achieved.

In 2001, Dr. Lane took sabbatical leave to perform research with the CRAL group at the Observatoire de Lyon in France, and I am very fortunate to have spent 5 months there. With the assistance of Dr. Michel Tallon and Dr. Eric Thiébaut, astronomical data was collected, some of which was used for my thesis. The availability of simultaneous speckle images and wavefront sensing data broadened the scope of my research to include hybrid methods.

In the process of using the wavefront sensing data, a deficiency in the existing centroid estimator is apparent, and a new improved approach based on blind deconvolution was devised. The application of the wavefront sensing data as a form of prior information in blind deconvolution is also investigated.

Although adaptive optics is not within the scope of this thesis, some of the intrinsic problems with this technique are similar to those in one of the most established hybrid methods, deconvolution from wavefront sensing. For instance, the use of the Shack-Hartmann wavefront sensor leads to the inability of these methods to accurately reconstruct the magnitude of the object spectrum. A solution has been proposed but uses additional reference star measurements. In view of this, a parameterised approach is investigated to remove the reliance on reference star measurements. This approach also offers the potential to overcome the problem with estimating the magnitude of the object.

One important recurring theme in this thesis is the concept of resolution. For a long time, it was thought that the highest attainable resolution is limited by the aperture of the optics, called the diffraction limit. In recent years, the phenomenon of reconstructing an object beyond the diffraction limit of the instrument, known as superresolution, is recognised to be attainable. Resolution is closely linked with sampling, however these two topics are rarely discussed together. In this thesis, I attempt to clarify the concept of resolution and its relation to sampling, as well as redefining some of the criteria for superresolution in relation

to the image-capturing instrument.

0.1 Thesis organisation

Chapter 1 presents an overview on image processing and the practical problems that are encountered. The specific problem of astronomical imaging is introduced, and a brief discussion of the possible solutions is given.

Chapter 2 introduces the notation and mathematics required to understand this thesis. Linear systems, transform theory and statistical processes are presented. This chapter concludes with an examination of real world data, linking the theory to its application in image processing.

Since all the problems examined in the thesis are inverse problems, an overview of the definition, feasibility and classes of solution for inverse problems is given in Chapter 3. Due to the ill-posed nature of many of these problems, regularisation, or the incorporation of prior information in the solution is essential. Different forms of prior information for image processing problems are also discussed in a general context in this chapter.

In Chapter 4, the theory of imaging through turbulence is presented. The principles of light propagation and the phenomenon of diffraction are examined. These are very important concepts because the design and analysis of the image processing algorithms in this thesis are built on these principles. The relationship between resolution and sampling, which is rarely found in literature, is clarified. The effect of the atmosphere on the wavefront of the star is defined, and properties of astronomical data are discussed. A review of the existing solutions for astronomical imaging is then given, and the associated pros and cons are pointed out.

Chapter 5 addresses the inverse problem of interpolation, and presents an optimal Bayesian interpolator which incorporates prior information in the form of the sampling process and statistical information in its design. Experimental results show that this algorithm has superior performance to existing algorithms. The framework of optimal interpolaters also encompasses the existing methods such as the bilinear and the sinc interpolators.

From Chapter 6 onwards, the focus is on astronomical imaging. Chapter 6 presents a review of existing blind deconvolution algorithms and investigates the application of two

new forms of prior information for solving this problem. Since the distortion introduced on the wavefront depends on the Earth's atmosphere, it is reasonable to constrain the solution with the statistics of the atmosphere. Simulation results indicate that this approach has the potential to out-perform existing algorithms. The phase closure property employed in the bispectrum technique is also considered as a form of prior information. The results of this feasibility study are also presented in this chapter.

One of the fundamental processes used on wavefront sensing data is the centroid estimator. The centroids of the spots need to be found in order to estimate the mean slope over the subapertures. However, the traditional centroid estimator suffers from many problems. An improved approach based on blind deconvolution is proposed in Chapter 7. Its superior performance is demonstrated on both simulated and experimental data.

With the availability of simultaneous wavefront sensing data and speckle images, the application of the former as an additional constraint in a blind deconvolution framework is examined in Chapter 8. The case of myopic deconvolution, where very good starting estimates of the quantities to be retrieved are available, is also addressed. The performance of this approach is demonstrated on experimental data.

In Chapter 9, the hybrid method of deconvolution from wavefront sensing (DWFS) is examined. The inability to accurately estimate the magnitude of the object spectrum in this method was pointed out by Roggemann *et al.* [149], who proposed a solution based on using a reference star. However, this does not completely remove the magnitude problem and requires reference star measurements which have other difficulties. An alternative solution based on parameterisation is proposed. This method has been derived from the NASRIF algorithm proposed by Kundur and Hatzinakos [102]. Comprehensive simulation results demonstrate the performance of this approach compared to DWFS.

Conclusions and recommendations for future work are presented in Chapter 10.

0.2 Supporting publications

Journal and conference papers published during the course of this thesis are listed below in the order of preparation.

1. W-Y.V. Leung, R.G. Lane and P.J. Bones, "Statistical interpolation of scanned data", *Proc. Image and Vision Computing New Zealand 1999*, D. Pairman and H. North

(Eds.), pp 103-108, August 1999.

2. W-Y.V. Leung and R.G. Lane, "Blind deconvolution of images blurred by atmospheric speckle", *Proc. SPIE Image Reconstruction from Incomplete Data*, Vol. 4123, pp 73-83, July/August 2000.

3. W-Y.V. Leung and R.G. Lane, "Joint blind deconvolution and registration of a sequence of undersampled images", *Proc. Image and Vision Computing New Zealand 2000*, M.J. Cree and A. Steyn-Ross (Eds.), pp 50-55, November 2000.

4. W-Y.V. Leung, P.J. Bones and R.G. Lane, "Statistical interpolation of sampled images", *Opt. Eng.*, Vol. 40, No. 4, pp 547-553, April 2001.

5. W-Y.V. Leung, M. Tallon and R.G. Lane, "Centroid estimation from undersampled wavefront sensing data", *Signal and Recovery Synthesis, OSA Technical Digest*, pp 102-104, November 2001.

6. W-Y.V. Leung and R.G. Lane, "Deconvolution from wavefront sensing via parameterisation", *Signal and Recovery Synthesis, OSA Technical Digest*, pp 124-126, November 2001.

7. W-Y.V. Leung and R.G. Lane, "Image reconstruction from undersampled images", *Proc. Image and Vision Computing New Zealand 2001*, B. McCane and K. Novins (Eds.), pp 79-83, November 2001.

8. W-Y.V. Leung, M. Tallon and R.G. Lane, "Centroid estimation by model-fitting from undersampled wavefront sensing images", *Opt. Comms.*, Vol. 201, pp 11-20, January 2002.

9. W-Y.V. Leung, R.M. Clare and R.G. Lane, "Blind deconvolution of speckle images constrained by wavefront sensing data", in *Proc. SPIE Image Reconstruction from Incomplete Data II, 2002.*

Papers in preparation include:

1. W-Y.V. Leung and R.G. Lane, "Post-processing of partially compensated imagery by recursive inverse filtering", submitted to *Opt. Comms.*

0.3   Notation

i	$\sqrt{-1}$
A^{-1}	inverse of the matrix A
$A^{\dagger}$	pseudo-inverse of the matrix A
A^{T}	transpose of the matrix A
$\odot$	convolution operation
$\otimes$	correlation operation
$\langle\rangle$	ensemble average
$\rightleftharpoons$	Fourier transform pair
(x,y)	rectangular coordinates in the spatial domain
(u,v)	rectangular coordinates in the Fourier or frequency domain
$\Re\{\cdot\}$	real part of
$\Im\{\cdot\}$	imaginary part of
∞	infinity
$E\{\cdot\}$	expectation operator
μ	mean
σ^2	variance
$P(\cdot)$	probability
$\sum$	summation
$\lvert\cdot\rvert$	magnitude
$\lVert\cdot\rVert$	2-norm
D	telescope aperture diameter
r_0	Fried's parameter, or the diameter of the seeing cell
k	wavenumber
λ	wavelength of light

$f(x,y)$ object

$h(x,y)$ point spread function

$g(x,y)$ noise free convolution of $f(x,y)$ and $h(x,y)$

$f^*(x,y)$ complex conjugate of $f(x,y)$

$n(x,y)$ measurement noise

$d(x,y)$ measured noisy data

$\delta(x,y)$ delta function

$\text{rect}(x,y)$ rectangle function

$\text{sinc}(x,y)$ sinc function

$\text{circ}(x,y)$ circle function

$J(r)$ Bessel function

$A(u,v)$ generalised pupil function

$\phi(u,v)$ phase of the generalised pupil function

$S(u,v)$ magnitude of the generalised pupil function

$H(u,v)$ optical transfer function

0.4   Abbreviations

ADC	analogue to digital converter
AO	adaptive optics
BD	blind deconvolution
BLUE	best linear unbiased estimator
CCD	charge-coupled device
CDF	cumulative distribution function
CG	conjugate gradient
DFT	discrete Fourier transform
DM	deformable mirror
DOF	degrees of freedom
DWFS	deconvolution from wavefront sensing
EM	expectation maximisation
FFT	fast Fourier transform
FFTW	fastest Fourier transform in the West
FOV	field of view
FT	Fourier transform
FWHM	full width at half maximum
HPF	highpass filter
IBD	iterative blind deconvolution
IFFT	inverse fast Fourier transform
KT	Knox-Thompson
LSI	linear shift-invariant
MAP	maximum *a posteriori*
ME	maximum entropy
ML	maximum-likelihood
MSE	mean-squared error
NASRIF	nonnegativity and support constraint recursive inverse filtering
OTF	optical transfer function
PDF	probability density function
PMAP	Poisson maximum *a posteriori*
POCS	projection onto convex sets
PSF	point spread function
RL	Richardson-Lucy
RV	random variable
SH	Shack-Hartmann

SNR   signal to noise ratio
SPID   SPeckle Imaging by Deconvolution
SPOT   Systeme Probatoire d'Observation de la Terre
STF   speckle transfer function
SVD   singular value decomposition
WFS   wavefront sensor

Chapter 1

Introduction

In many scientific and engineering applications, information is recorded digitally for later reference or further processing. In image processing, information is restricted to two dimensions, such as images captured on charge-coupled device (CCD) cameras, photographs converted into a digital format by scanners, and astronomical images captured at telescopes by photon-counting cameras. Although all the image-recording devices have different modes of operation and can be used in varying conditions, they all suffer from the effects of noise and distortions introduced by the imaging elements and the imaging media. As a result of this non-ideal imaging, much of the captured digital data must be post-processed before the information can be interpreted.

Image processing is divided into four main categories: rectification, enhancement, restoration and reconstruction [9]. Image rectification deals with the spatial transformation of images in order to align them with respect to the reference grid or each other, and image enhancement refers to the improvement of the presentation of an image. Image restoration involves the recovery of an object from measured data in the absence of information on the distortions, while image reconstruction is the recovery of detail in severely distorted images when the distortion is known *a priori*. This thesis deals with all the aspects of image processing, but focuses mainly on image restoration, in particular the restoration of astronomical images.

1
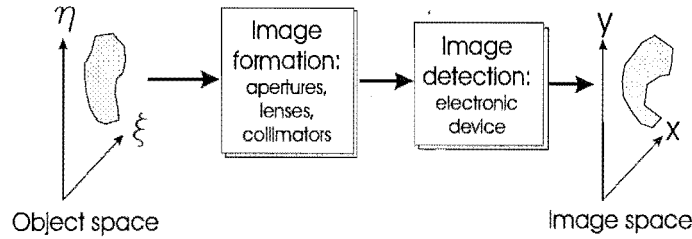
Figure 1.1: *The mapping of an object in object space into an image in image space.*

1.1   Image formation, detection and degradation

In order to solve the inverse problem of image processing, it is essential to understand the forward processes of image formation, detection and degradation. An object in object space (ξ, η) is imaged in the image space (x, y) when the light rays of the object pass through the imaging components and are detected by an image-capturing device (Fig. 1.1). There are several limitations imposed on this system. Firstly, diffraction occurs when the wavefront of the light is truncated, for example by an aperture. Such an image is called a diffraction-limited image. This phenomenon is common to all imaging systems, and the theory of diffraction is discussed in Chapter 4.

Secondly, the images are corrupted by noise which is inherent to the image formation and detection processes. There are two main sources of noise in a photoelectronic system: random thermal noise resulting from the resistive circuit elements, and the photoelectronic noise due to the statistical nature of light [19]. The former is adequately modelled by Gaussian statistics, and a common example of an image with Gaussian noise is a digital photograph. At low light levels, the small number of photons means that photoelectronic noise is better modelled as a Poisson process. In a speckle image captured at a telescope, both Poisson and Gaussian noises are present.

Finite sampling frequency, finite measurement time, and the mechanical limitations of the equipment also affect the resolution and the amount of data captured. An example is the collection of astronomical images. The exposure time of an image is dependent on the available shutter speed, while the resolution of the captured images is dependent on the density and the size of the CCD array. Finally, the imaging medium also contributes to the quality of the captured images. In astronomical imaging, a major limitation to the resolution of the

images is the Earth's turbulent atmosphere. For an 8-metre telescope, an angular resolution of 0.01-0.02 arcsec is theoretically achievable. However, the atmospheric turbulence limits the realisable resolution to approximately 1 arcsec even at the best observing sites [147].

The success in solving the inverse problem of image processing relies on an accurate model of the forward problem of image formation, detection and degradation. Section 1.2 describes a commonly used model for the forward problem in image processing.

1.2 Mathematical model

The degradation process is described using a continuous linear model, denoted mathematically by a Fredholm integral [4]:

$$d(x,y) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(\xi,\eta)h(x,y;\xi,\eta)d\xi d\eta + n(x,y) \qquad (1.1)$$

where $f(\xi,\eta)$ is the undegraded object, $h(x,y;\xi,\eta)$ is the space-variant point spread function (PSF), $n(x,y)$ is the noise on the system, and $d(x,y)$ is the degraded noisy data. The Cartesian coordinates are denoted by (x,y). If the PSF is assumed to be space-invariant, or isoplanatic, it can be written as $h(x-\xi, y-\eta)$. In this thesis, the PSF is restricted to the space-invariant case, which sufficiently describes the physical situations modelled.

Although the continuous model is representative of real-world objects, a discrete model is required if the processing is to be performed digitally on a computer. The discrete form of the degradation model assuming space-invariance is:

$$d(x,y) = \sum_{k=-\infty}^{\infty} \sum_{l=-\infty}^{\infty} f(k,l)h(x-k,y-l) + n(x,y). \qquad (1.2)$$

Taking into account the finite extent of the object and the PSF removes the need to perform a summation over all-space. This further simplifies Eq. (1.2) to:

$$d(x,y) = \sum_{k=-N/2}^{N/2-1} \sum_{l=-N/2}^{N/2-1} f(k,l)h(x-k,y-l) + n(x,y) \qquad (1.3)$$

for an object of size $N \times N$. Using a compact notation, the degradation model can be written as:

$$d(x,y) = f(x,y) \odot h(x,y) + n(x,y) \qquad (1.4)$$
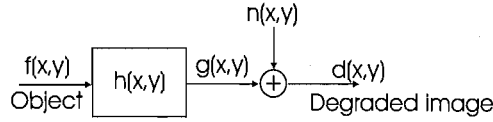
Figure 1.2: *An image degradation model.*

where $\odot$ denotes the two-dimensional convolution operation. A block diagram in Fig. 1.2 illustrates this degradation model. It is also common to represent this discrete model in matrix form [4]:

$$\mathbf{d} = \mathbf{Hf} + \mathbf{n} \tag{1.5}$$

where the two-dimensional functions $f(x,y)$, $n(x,y)$ and $d(x,y)$ are lexicographically ordered into vectors. If $\mathbf{f}$ is the object of size $N \times N$, then $\mathbf{f}$ is of size $N^2 \times 1$. $\mathbf{H}$ is a Toeplitz matrix of size $N^2 \times N^2$ which maps the object to the data.

1.3 Inverse problems

In many scientific problems, the direct measurement of the object is impossible or impractical. As a result, the object must be deduced from the degraded and distorted image by solving an inverse problem. In general, information loss is significant, and the correct modelling of the forward problem is vital in identifying the limited amount of information available for use in the inverse problem.

This thesis deals with two major areas of inverse problems: overcoming firstly the limitations imposed by image-capturing instruments, and secondly the distortions introduced by the imaging medium, particularly atmospheric turbulence. While in the former, the primary causes for such a need are the undersampling of many image-capturing devices and the presence of detector noise [64, 73, 83, 109], the latter arises due to the distortion of the astronomical object by the time-varying turbulent atmosphere [140, 147].

Although the situations from which the inverse problems arise vary, the solutions to these problems are best derived according to the following principle (Fig. 1.3). In most cases, the forward problem is well-modelled, and the object of interest is updated by comparing the measured data with the estimated data when the object is put through the model.
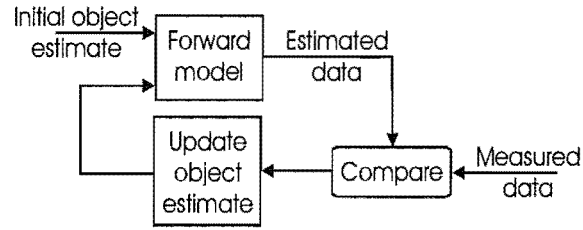
Figure 1.3: *The underlying principle for obtaining a solution for inverse problems.*

This methodology is adopted because the measured data is not altered, and as a result, no information within the data is inadvertently lost.

In general, the accuracy of a solution increases with the complexity of the model. However, the latter is not necessarily directly proportional to the former. A good example is fitting a model to a function, where the degrees of freedom (DOF) in the model depends on the realistic degrees of freedom in the function. Increasing the number of parameters in the model provides a better fit to the function, until the number of parameters equals the degrees of freedom in the function. After that, any increase in the model complexity does not improve but can potentially degrade the accuracy of the solution because the model is starting to fit the noise and not the data. The tradeoff between the potential accuracy of the solution and the ease of obtaining it must also be taken into account when solving the inverse problem.

The following subsections discuss the two areas of inverse problems treated in this thesis.

1.3.1 Overcoming instrumentation limitations

The acquisition of real-world continuous data by a digital instrument is called sampling. If the values of the samples correspond directly to the values of the original functions at the sample points, this process is called point sampling. Although theoretically possible, this process is not adopted in real instrumentation. In real devices such as CCD and Shack-Hartmann sensors, area sampling is performed where each sample value is obtained by integrating over each sensor and corresponds to the average of the light falling on it. This maximises the light captured, and is particularly relevant for instruments operating under

low light levels.

Although the sampling process does not record every point of the original function, provided that the function is band-limited and the samples are sufficiently close to each other, it is possible to reconstruct the function exactly. This is formalised in the Nyquist sampling criterion [9], which states that a band-limited signal must be sampled at at least twice the highest frequency component of the data for the original continuous function to be recovered. This is expressed mathematically as:

$$X_0 = \frac{1}{2f_H} \tag{1.6}$$

where X_0 is the sampling rate, and f_H is the highest frequency component of the data. The corresponding sampling frequency is $1/X_0$.

Consider a bandlimited one-dimensional continuous signal, whose spectrum is shown in Fig. 1.4(a). If this is critically sampled at the Nyquist frequency $2f_H$, the spectrum of the sampled signal is as shown in Fig. 1.4(b). However, if it was sampled at a frequency below that of $2f_H$, i.e. undersampled, the sampled spectrum is the sum of overlapping components, a phenomenon known as aliasing (Fig. 1.4(c)). Aliasing describes how a higher frequency component is shifted to a lower frequency sample in the spectrum. Since the shifted frequency is added to the existing lower frequency, the separation of these frequency components is difficult.
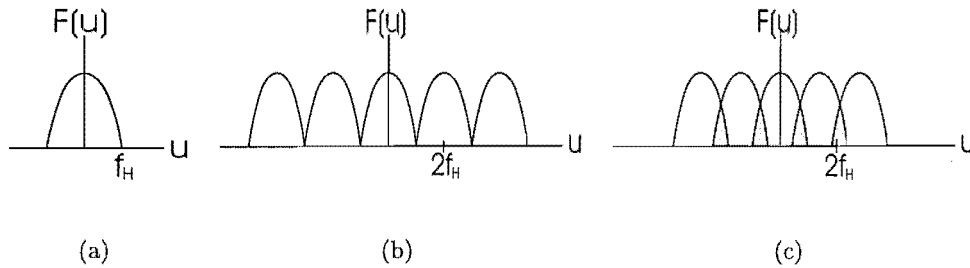
Figure 1.4: *Illustrating the problem of undersampling. (a) Spectrum of a one-dimensional band-limited continuous signal. (b) Spectrum of the critically sampled signal. (c) Spectrum of an undersampled signal. The grey areas indicate the overlapping of high and low frequency components.*

Undersampling is a common problem when capturing images. Often, the charge-coupled

device (CCD) cells in the cameras do not meet the Nyquist criterion. This problem is exacerbated in astronomical imaging, where the CCD cell size is optimised to minimise read noise; as a result, the images acquired can be sampled at as low as 60 % of the Nyquist frequency.

There are many approaches to overcoming the sampling limitations imposed by instrumentation, one of which is image interpolation. Image interpolation involves the deduction of the underlying high-resolution object from the noisy low-resolution measurements and constitutes an inverse problem. This topic has been an active area of research in image processing for many years, and numerous methods have been devised to solve this inverse problem. These include the common linear interpolator, the sinc and its windowed counterpart, polynomial interpolators and splines.

However, many of the existing interpolation algorithms neither deal with aliasing in the data directly, nor do they take into account the sampling process and the noise on the image capturing instrument. One of the topics examined in this thesis is an optimal interpolator which incorporates the sampling process and deals with aliasing via the use of prior information. This is presented in Chapter 5.

The process of interpolation is closely linked to superresolution as they can be viewed as being complementary to each other. Whilst the former describes the increase in sampling rate in the spatial domain, the latter depicts the extension of the meaningful frequency content of an image. In fact, in order to superresolve, the corresponding sampling rate must be increased [85].

Superresolution has generated much research interest, and there are many definitions of this process. One popular definition is the recovery of information beyond the frequency bandwidth of the system that produced the signal [85]. In the context of two-dimensional image processing, superresolution strives to overcome the diffraction limit imposed by the optical instrument. This is shown to be theoretically possible [61, 101], and several algorithms have had success in achieving superresolution in this sense. These include the Gerchberg algorithm [54] which projects between the spatial and Fourier domains to obtain the solution, and the reconstruction method for Poisson noise statistics, developed independently by Richardson [139], Lucy [114], and Shepp and Vardi [156]. Many algorithms proposed more recently [56, 157] are based on these original algorithms. In each case, information in

addition to the measured data is required to achieve superresolution.

In the above definition of superresolution, the diffraction limit of the optics is the principal form of resolution degradation in the imaging system. However, instead of the optics, it is more common for the CCD camera to be the major obstacle as it usually has a lower frequency cutoff compared to the optics. In this case, the process of superresolution must be redefined as the process of extending the frequency content of the recovered object beyond that imposed by the imaging element which has the lowest cutoff frequency. This is the definition adopted in this thesis. There are many published algorithms which address this problem of superresolution whereby the concept of subpixelly displaced images is used to improve the sampling rate [22,57,73,109]. In Chapter 7, the algorithm proposed by Hardie *et al.* [73] is extended and applied to severely undersampled astronomical data.

1.3.2 Overcoming the imaging medium: atmospheric turbulence

The wavefront from an astronmical object of interest such as a star propagates undistorted for lightyears, until it reaches the Earth's atmosphere (Fig. 1.5). The atmosphere can be viewed as a collection of lenslets with different and time-varying refractive indices. As a result, phase distortions are introduced to the wavefront. The atmospheric turbulence is sufficiently described by Kolmogorov statistics, characterised by the two-dimensional horizontal structure function [140]:

$$D(u_1, u_2, v_1, v_2) = 6.88 \left(\frac{\sqrt{(u_1 - u_2)^2 + (v_1 - v_2)^2}}{r_0} \right)^{5/3}, \qquad (1.7)$$

where r_0 is the effective diameter of the telescope and is in the order of about 5-15cm [147]. Another interpretation of r_0 is that it defines an aperture size over which the mean-square wavefront error is 1 rad^2 [77].

When the wavefront subsequently propagates to ground, its intensity is also distorted, and this distortion is called scintillation. Nevertheless, since the principal effect of the distortion introduced by the Earth's atmosphere on an image is a phase component, the scintillation is assumed negligible in many cases.

Due to the distorting effects of atmospheric turbulence, the images captured by a ground-based telescope have a 'speckly' appearance, hence the name speckle images. Figs. 1.6(a)

and 1.6(b) illustrate a simulated binary star before and after the distortion of the atmosphere in the absence of noise, and Fig. 1.6(c) shows the distortion when Poisson noise is present. It can be seen that the extent of the speckles is much larger than the extent of the object. It is therefore possible for some data information to be truncated by the image detector when capturing real astronomical data [92]. This causes errors in the second-order spectra in techniques such as speckle interferometry [103] and the Knox-Thompson [99] method.

The dominance of the phase distortion on the star's wavefront means that it is desirable to be able to estimate and compensate for this phase component. A common instrument for measuring this is a wavefront sensor, with the four main types being the Shack-Hartmann (SH) sensor, the curvature sensor, the shearing interferometer and the pyramid sensor. Fig. 1.7 shows a typical frame of wavefront sensing data captured by a SH sensor at the Observatoire de Lyon in France. The estimation of the phase distortion from wavefront sensing data is discussed in detail in Chapter 8.

Imaging through atmospheric turbulence is an active area of research, and there are three main approaches to this problem: real time adaptive optics (AO), computer post-processing, and hybrid techniques which are a combination of both. This thesis focuses on the latter two approaches.

Computer post-processing

In pure computer post-processing techniques, only the speckle images are used for recovering the intensity distribution of the object. The two main approaches in this category are blind deconvolution and bispectrum.

The problem of blind deconvolution can be modelled accurately using Eq. (1.4), where $f(x, y)$ is the object of interest, $h(x, y)$ is the PSF describing the combined effects of the atmosphere and the imaging optics, and $n(x, y)$ is the noise on the system. The captured image $d(x, y)$ is the speckle image (Fig. 1.6(c)).

Blind deconvolution is an ill-posed inverse problem where one wishes to find $f(x, y)$, $h(x, y)$ and $n(x, y)$ when only $d(x, y)$ is measured. In the absence of noise, blind deconvolution has been shown to be theoretically possible [107]. However, when noise is present, the problem
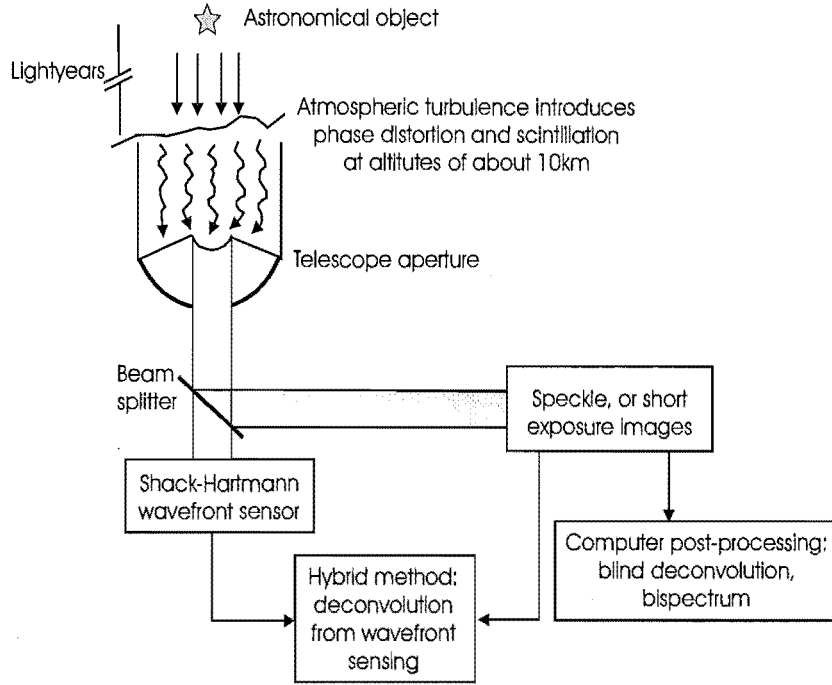
Figure 1.5: *Block diagram illustrating two classes of methods for solving the astronomical imaging problem.*

becomes ill-conditioned. As a result, the quality of the solution depends strongly on the amount of prior information available. When the only prior information available is the noise statistics, the problem is one of maximum likelihood (ML). This however leads to the trivial solution where the object is estimated to be a delta function, and the PSF the data (or vice versa). The ML method is therefore insufficient to produce a meaningful solution in this case. One type of commonly used prior information is in the form of multiple frames:

$$d_k(x,y) = f(x,y) \odot h_k(x,y) + n_k(x,y) \tag{1.8}$$

where k denotes the k^{th} realisation of the PSF, the noise, and the data. Other forms of prior information, such as the compact nature and the positivity of the object are also used.

Obviously, there are many forms of prior information that can be used to constrain the solution. The investigation of the feasibility and effectiveness of various forms of prior information is a major focus of this thesis. A strong constraint that can be imposed is the wavefront sensing data that is captured simultaneously with the speckle images, as this provides a good estimate for the phase distortion. This is discussed in Chapter 8. Another form of prior information that can be used for astronomical imaging is the Kolmogorov

(a)   (b)   (c)
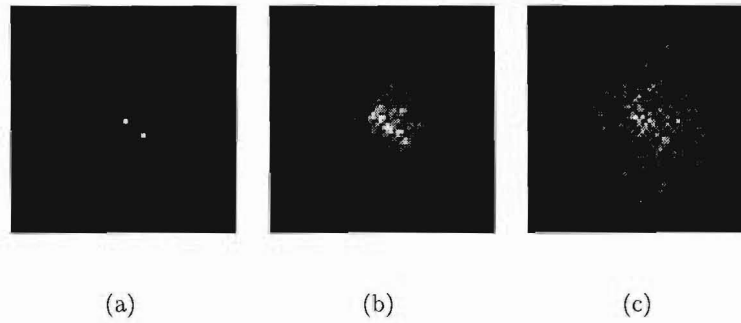
Figure 1.6: *(a) A simulated binary star. (b,c) An image of the star after the degradation caused by atmospheric turbulence in the absence and presence of Poisson noise. The images are displayed on a gray scale where the white pixels have a value of 1, and the black pixels have a value of 0.*
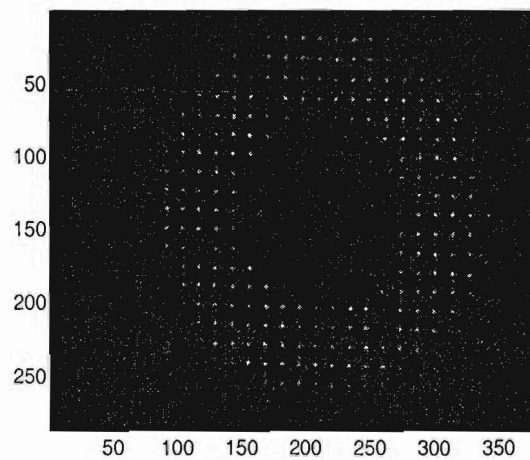
Figure 1.7: *An example of a wavefront sensing data frame.*

statistics of the turbulence. The investigation into its feasibility is presented in Chapter 6.

The bispectrum is an established technique for processing images formed with an incoherent imaging system [113]. The object is reconstructed in the Fourier domain, with the magnitude component obtained from speckle interferometry technique, and the phase component from triple correlation theory. This technique requires the use of multiple frames, and the survival of the phase information is attributed to the use of a four-dimensional ensemble quantity, called the bispectrum phase. Since most of the information is contained in several

dominant components in the bispectrum phase, it is common to use only these components and not the entire quantity. One such subset is the Knox-Thompson (KT) phase.

Although the technique of bispectrum in itself is a complete solution to the problem of image restoration, it can also be used as a form of prior information [58, 59, 72, 123–125]. In this thesis, a blind deconvolution algorithm with the KT phase as prior information is investigated. Since the KT phase of the ensemble PSF should equal to zero, the resultant KT phase must be attributed to the object. This property is used in a penalty term in the solution framework. The result of this investigation is presented in Chapter 6.

Hybrid techniques

The most frequently used hybrid technique for astronomical imaging is deconvolution from wavefront sensing (DWFS). It uses simultaneous speckle images and wavefront sensing data to recover the object distribution. This technique was first proposed by Fried [46] and later implemented by Primot *et al.* [135] in 1990. The retrieval of the object $\hat{f}(x, y)$ is an inverse problem, whose spectrum is recovered by:

$$\hat{F}(u, v) = \frac{\langle D(u, v)\hat{H}^*(u, v)\rangle}{\langle |\hat{H}(u, v)|^2\rangle + c} \tag{1.9}$$

where c is a constant, and the angular brackets denote an ensemble quantity. The optical transfer function (OTF) $H(u, v)$ is the autocorrelation of the generalised pupil function $A(u, v)$:

$$A(u, v) = S(u, v)\exp\{i\phi(u, v)\}. \tag{1.10}$$

$S(u, v)$ is the magnitude and $\phi(u, v)$ is the phase distortion. In its implementation, the phase $\phi(u, v)$ is reconstructed using the wavefront sensing data which is also an inverse problem. Thus the method of DWFS comprises two inverse problems.

There are two disadvantages in the conventional technique of DWFS. Firstly, the phase is estimated solely from wavefront sensing data, and phase information in the speckles is not utilised. An improved method is to estimate the phase from data in both the speckle and wavefront sensing channels in an iterative blind deconvolution framework. This is the topic of discussion in Chapter 8.

Secondly, the correlation properties of the true phase $\phi(u,v)$ and the reconstructed phase $\hat{\phi}(u,v)$ are different [149]. This means that the magnitude of the object spectrum is not reconstructed correctly. The solution to this problem is the use of a reference star to obtain measurements for the denominator of Eq. (1.9) [149]. However, this does not completely remove the problem of inaccurate magnitude estimation; moreover reference star measurements are in practice very difficult to obtain. These problems can be overcome by a parameterised solution derived from the NASRIF algorithm [102], the formulation of which is presented in Chapter 9.

Chapter 2

Mathematical preliminaries

This chapter introduces the general mathematical definitions and notations required for this thesis. However, theories that are specific to each subsequent chapter are introduced as required. Section 2.1 defines the notation used, including vectors, complex numbers, functions and coordinate spaces. It is convenient to assign special notation to commonly used functions; these are discussed in Section 2.2. Section 2.3 presents the theory of linear systems, and common operations on the variables such as convolution and correlation are defined. There are situations where the variables can be manipulated more easily in a different coordinate space. This is part of transform theory and is discussed in Section 2.4. Many phenomena in the real world such as atmospheric turbulence and noise are not deterministic; therefore the theory of random processes is required and it is presented in Section 2.5. Section 2.6 deals with the properties of real world data.

2.1 Notation

2.1.1 Scalars, vectors and matrices

A scalar is a real number denoted by x. A vector, denoted by the boldface letter $\mathbf{x}$, has n elements $(x_1, x_2, \cdots, x_n)$. Physically, it can be interpreted as a point in an n-dimensional space. A matrix is an array of numbers which is of two or more dimensions and is usually denoted by a boldface uppercase letter. An example of a two-dimensional matrix of size

$m \times n$ is denoted by:

$$\mathbf{X} = \begin{pmatrix} x_{1,1} & \cdots & x_{1,n} \\ & \ddots & \\ x_{m,1} & \cdots & x_{m,n} \end{pmatrix}.$$

The row vectors in this matrix are the vectors in $\mathbf{R}^n$, where $\mathbf{R}$ denotes a real space and the superscript n defines the dimension of this real space. The row vectors are:

$$\begin{aligned} \mathbf{r}_1 &= (x_{1,1}, x_{1,2}, \ldots, x_{1,n}) \\ &\vdots \\ \mathbf{r}_m &= (x_{m,1}, x_{m,2}, \ldots, x_{m,n}) \end{aligned} \tag{2.1}$$

while the column vectors are the vectors in $\mathbf{R}^m$ and can be written in a similar manner.

A vector $\mathbf{x}$ is a linear combination of vectors $\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n$ if there exist scalars $c_1, c_2, \ldots, c_n$ such that $\mathbf{x} = c_1\mathbf{x}_1 + c_2\mathbf{x}_2 + \ldots + c_n\mathbf{x}_n$. A set of vectors $\{\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n\}$ is linearly independent if the equation:

$$c_1\mathbf{x}_1 + c_2\mathbf{x}_2 + \ldots + c_n\mathbf{x}_n = 0 \tag{2.2}$$

has only the trivial solution $c_1 = c_2 = \ldots = c_n = 0$.

For an $m \times n$ matrix $\mathbf{X}$, the range of $\mathbf{X}$ is:

$$\text{range}(\mathbf{X}) = \{\mathbf{b} \in \mathbf{R}^m : \mathbf{b} = \mathbf{X}\mathbf{a} \text{ for some } \mathbf{a} \in \mathbf{R}^n\}; \tag{2.3}$$

the null space of $\mathbf{X}$ is:

$$\text{null}(\mathbf{X}) = \{\mathbf{a} \in \mathbf{R}^n : \mathbf{X}\mathbf{a} = 0\}; \tag{2.4}$$

and the rank of $\mathbf{X}$ is:

$$\text{rank}(\mathbf{X}) = \dim(\text{range}(\mathbf{X})) \tag{2.5}$$

where $\dim(\cdot)$ indicates the dimension of a matrix. The trace of a square matrix is the sum of its diagonals. Mathematically, it is defined as:

$$\text{trace}(\mathbf{X}) = \sum_{j=1}^{n} x_{jj} \tag{2.6}$$

for an $n \times n$ matrix.

The transpose of a matrix $\mathbf{X}^T$ transforms the row vectors into column vectors and vice versa. Therefore the transpose of an $m \times n$ matrix has size $n \times m$. In the case where the original matrix is identical to its transpose, i.e. $\mathbf{X} = \mathbf{X}^T$, the matrix is symmetric.

The inverse $\mathbf{X}^{-1}$ of a square matrix of size $n \times n$ is such that $\mathbf{X}\mathbf{X}^{-1} = \mathbf{X}^{-1}\mathbf{X} = \mathbf{I}$ where $\mathbf{I}$ is the identity matrix defined at the end of this section [101]. For a non-square matrix, a pseudo-inverse can be computed using the technique of singular value decomposition (SVD) [60]. For a matrix of size $m \times n$, it can be decomposed into:

$$\mathbf{X} = \mathbf{U}\Lambda\mathbf{V}^T \tag{2.7}$$

where $\mathbf{U}$ has the eigenvectors of $\mathbf{X}\mathbf{X}^T$ and $\mathbf{V}$ has the eigenvectors of $\mathbf{X}^T\mathbf{X}$. Λ is a diagonal matrix containing the singular values of $\mathbf{X}$, whose squares are the eigenvalues of $\mathbf{X}^T\mathbf{X}$ [162]. The pseudo-inverse of $\mathbf{X}$, $\mathbf{X}^\dagger$, is:

$$\mathbf{X}^\dagger = \mathbf{V}\Lambda^{-1}\mathbf{U}^T. \tag{2.8}$$

For any matrix, if an inverse exists, the matrix is nonsingular. Otherwise, the matrix is singular. The inverse and the pseudo-inverse are one way of solving an inverse problem. Consider a system denoted by $\mathbf{X}\mathbf{a} = \mathbf{b}$, where $\mathbf{a}$ and $\mathbf{b}$ are the unknown and measured vectors respectively, and $\mathbf{X}$ is the matrix which maps $\mathbf{a}$ onto $\mathbf{b}$. The estimate of $\mathbf{a}$ can be found by calculating $\mathbf{X}^{-1}\mathbf{b}$ or $\mathbf{X}^\dagger\mathbf{b}$. The inverse or the pseudo-inverse thus provides one of the simplest solutions to such an inverse problem.

The large singular values in Λ correspond to significant amounts of information being encoded in them; conversely small singular values signify that they only contain very little information. A problem arises when Λ is inverted, leading to the amplification of small singular values.

The eigenvalue is related to the eigenvector. For a square matrix $\mathbf{X}$ of size $n \times n$, the scalar λ is an eigenvalue of $\mathbf{X}$ if there exists a non-zero column vector $\mathbf{b}$ in $\mathbf{R}^n$ such that $\mathbf{X}\mathbf{b} = \lambda\mathbf{b}$ [40]. When all the eigenvalues of a square matrix are positive (and not equal to zero), this matrix is positive definite and its invertibility is guaranteed. If all the eigenvalues are greater than or equal to zero, the matrix is positive semidefinite.

A Toeplitz matrix $\mathbf{T}$ is one whose entries are constant along each diagonal and arises in many applications [60]. Formally, $\mathbf{T} \in \mathbf{R}^{n \times n}$ is Toeplitz if there exist scalars $t_{-n+1}, \ldots, t_0, \ldots, t_{n-1}$

such that $x_{i,j} = t_{j-i}$ for all i and j, with i and j being the row and column indices:

$$\mathbf{T} = \begin{pmatrix} t_0 & t_1 & t_2 \\ t_{-1} & t_0 & t_1 \\ t_{-2} & t_{-1} & t_0 \end{pmatrix}.$$

An example of a Toeplitz matrix is the point spread function matrix $\mathbf{H}$ in Eq. (1.5).

A commonly used special matrix is the identity matrix $\mathbf{I}$. The identity matrix is square and has the value 1 along the main diagonal, and 0 for all other entries. The identity matrix is symmetric and Toeplitz.

Basis functions and orthogonality

A set of vectors $S = \{\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n\}$ in a vector space $\mathbf{X}$ is a basis for $\mathbf{X}$ if

1. S is linearly independent and

2. S spans $\mathbf{X}$.

Basis functions are used commonly in image processing. The Fourier transform is a ubiquitous example, while the Zernike basis functions [131] are used to represent atmospheric turbulence. Two vectors $\mathbf{x}_1$ and $\mathbf{x}_2$ are said to be orthogonal if $\mathbf{x}_1^T \mathbf{x}_2 = 0$.

2.1.2   Complex numbers

A complex number is one which has a real and an imaginary part. Let us denote a complex number by $z = x + iy$, where $i = \sqrt{-1}$. Then the real part of z is $\Re\{z\} = x$ and its imaginary part is $\Im\{z\} = y$. It is convenient to represent any associated value of z on the x–y plane, also called the z plane (Fig. 2.1). It is also possible to denote a complex number in polar form where:

$$\begin{aligned} r &= \sqrt{x^2 + y^2} \\ \theta &= \arctan\left(\frac{y}{x}\right). \end{aligned} \tag{2.9}$$
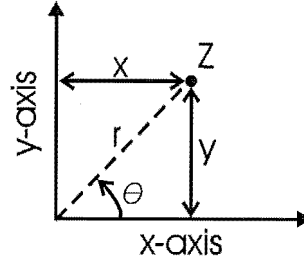
Figure 2.1: *The z-plane.*

Then

$$z \ = \ x + iy$$
$$= \ r(\cos\theta + i\sin\theta)$$
$$= \ r\exp\{i\theta\}. \tag{2.10}$$

2.1.3 Functions

The first usage of the term 'function' was by Leibniz in 1673 to describe the dependence of one quantity on another. However, Leonard Euler was the first person to denote functions by letters of the alphabet [5]. If the letter f is used to denote a function, then the expression:

$$z = f(x) \tag{2.11}$$

means that z is dependent on x. When extended to two-dimensional (2D) spaces, the notation:

$$z = f(x,y) \tag{2.12}$$

can be used. This notation is valid for real numbers and complex numbers. For example, the above equation can be used to describe a real function $z = 2x + y$, or a complex function $z = x + iy$. The most common usage of functions in imaging is the coordinate space which is the topic of discussion in the next section.
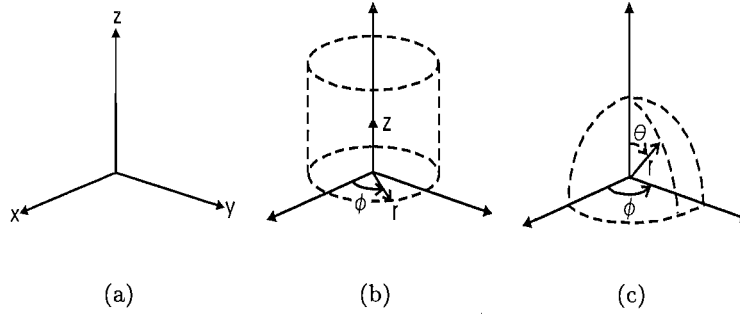
Figure 2.2: *(a) Cartesian or rectangular coordinate system. (b) Cylindri-cal coordinate system. (c) Spherical coordinate system.*

2.1.4 Coordinate space

The three commonly used coordinate systems are the Cartesian or rectangular system, the cylindrical system, and the spherical system. Since a point in three-dimensional space can be specified by three parameters, all these systems are three-parameter systems, depicted in Fig. 2.2. The rectangular coordinate system consists of three orthogonal planes intersecting at the origin. This is denoted by (x, y, z) and is the most commonly used system in image processing; the image is represented in the two-dimensional x-y plane, while z is often used to describe distances such as the focal length of a lens. Although the light waves that radiate from a point source are spherical, when they reach the imaging plane after a large propagation distance, they are essentially planar. It is therefore justified to employ the Cartesian coordinates in image processing.

In the cylindrical system [137], the coordinate surfaces are firstly a set of circular cylinders r, secondly a set of planes passing through the axis ϕ, and thirdly the distance from the origin along the axis z. In spherical coordinates [137], the three parameters are the distance from the origin r, the angle about the azimuthal axis ϕ, and the polar angle θ. The cylindrical and spherical systems are not used in this thesis.

2.2 Special functions

Dirac delta function

The Dirac delta function, or the unit impulse, is denoted by $\delta(x,y)$. It is defined as having zero amplitude everywhere except at the origin $(x,y) = (0,0)$, where it is infinitely large and contains unit volume under its curve [78]:

$$1 = \int_{0^-}^{0^+} \int_{0^-}^{0^+} \delta(x,y)dxdy. \tag{2.13}$$

A two-dimensional (2D) delta function is defined as:

$$\delta(x,y) = \begin{cases} \infty & x = 0, y = 0 \\ 0 & \text{otherwise} \end{cases}.$$

Intuitively, the delta function can be viewed as the limiting form of a pulse of unit volume as the duration of the pulse approaches zero. Fig. 2.3(a) shows a representation of a delta function.

Sampling function

The sampling function, or the Dirac comb, consists of an infinite sequence of uniformly spaced delta functions. This two-dimensional function is defined as:

$$\text{comb}(x,y) = \sum_{m=-\infty}^{\infty} \sum_{n=-\infty}^{\infty} \delta(x - mX_0, y - nY_0) \tag{2.14}$$

where X_0 and Y_0 are the sampling intervals in the $x-$ and $y-$ directions. It is shown in Fig. 2.3(b). This function is usually called the ideal sampling function and is ideal for the case of point sampling. However in real instruments, this ideal sampling is physically unrealisable. Moreover as discussed in Section 1.3.1, area sampling is the process adopted by real instruments, not point sampling. Nevertheless, this sampling function is widely used in signal processing literature for analysing systems. In Section 2.6.1, the area sampling of a function is discussed.

Rectangle function

The rectangle function, also known as the rect function, is defined as:

$$\text{rect}(x,y) = \begin{cases} 1 & |x| \leq \frac{1}{2}, |y| \leq \frac{1}{2} \\ 0 & \text{otherwise}. \end{cases}$$

(a)   (b)
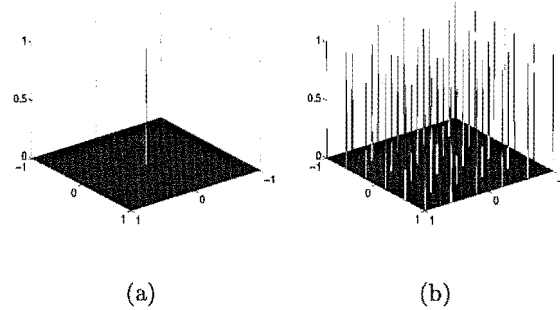
Figure 2.3: *(a) Two-dimensional Dirac delta function. (b) Two-dimensional Dirac comb or sampling function.*

The 2D rect function is depicted in Fig. 2.4(a) and can be separated into two 1D functions. This function is useful for describing a square aperture in an imaging system (see Chapter 4).

Sinc function

The 2D sinc function (Fig. 2.4(b)) is defined as:

$$\text{sinc}(x,y) \quad = \quad \frac{\sin(\pi x)}{\pi x} \frac{\sin(\pi y)}{\pi y}$$
$$= \quad \text{sinc}(x)\text{sinc}(y)$$

and can be separated into two 1D functions. There exists a Fourier transform relationship between the sinc and the rect functions. This Fourier transform pair is therefore useful for investigating diffraction effects on an imaging system with a square aperture such as a Shack-Hartmann lenslet array.

Sign function

The sign function is defined as:

$$\text{sgn}(x) = \begin{cases} 1 & x > 0 \\ 0 & x = 0 \\ -1 & x < 0 \end{cases}$$

and is shown in Fig. 2.5.
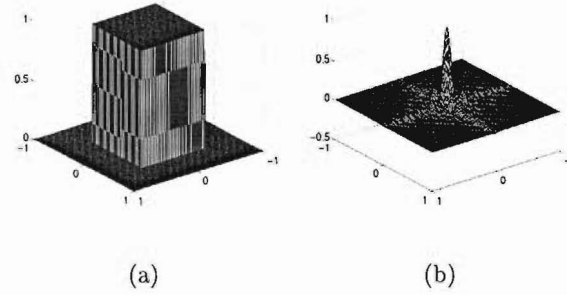
Figure 2.4: *(a) Two-dimensional rect function. (b) Two-dimensional sinc function.*
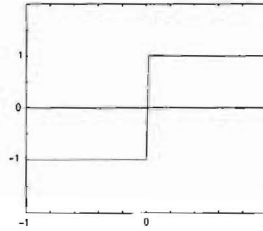
Figure 2.5: *Sign function*

Circle function

The circle function shown in Fig. 2.6 is defined as:

$$\text{circ}(x, y) = \begin{cases} 1 & \sqrt{x^2 + y^2} \leq 1 \\ 0 & \text{otherwise.} \end{cases}$$

Alternatively, due to the circular symmetry of this function, it can be defined with one variable r alone:

$$\text{circ}(r) = \begin{cases} 1 & r \leq 1 \\ 0 & \text{otherwise.} \end{cases}$$

This function is useful for describing a circular aperture of an imaging system. The circle function has a Fourier transform relationship with the Bessel function which is discussed in Section 2.4.3. As a result of this relationship, this Fourier transform pair is useful for the investigation of diffraction effects at a circular aperture which is discussed in Chapter 4.

Gaussian function

The Gaussian function is a commonly encountered function in many physical systems. A
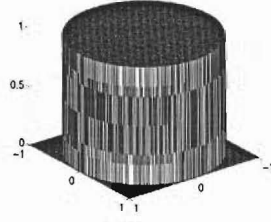
Figure 2.6: *Circle function.*

(a)   (b)
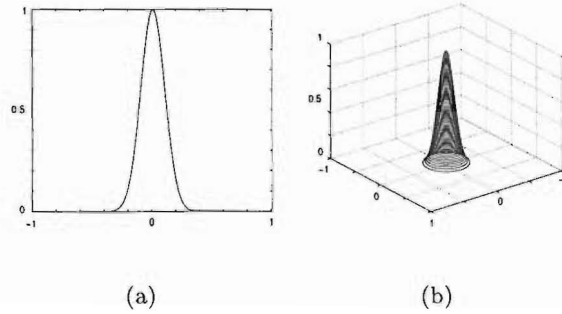
Figure 2.7: *(a,b) One and two-dimensional Gaussian functions.*

Gaussian function is a smooth function specified with two parameters: the mean μ, and the standard deviation σ. In one dimension, it is defined as:

$$G(x) = \frac{1}{\sqrt{2\pi}\sigma_X}\exp\left\{ -\frac{(x-\mu_X)^2}{2\sigma_X^2} \right\} \tag{2.15}$$

for all x. In two dimensions, it is a separable function defined as:

$$G(x,y) = \frac{1}{2\pi\sigma_X\sigma_Y}\exp\left\{ -\frac{(x-\mu_X)^2}{2\sigma_X^2} \right\}\exp\left\{ -\frac{(y-\mu_Y)^2}{2\sigma_Y^2} \right\}. \tag{2.16}$$

Examples of a 1D and 2D function are shown in Figs. 2.7(a) and 2.7(b).

The Gaussian function is versatile and has wide applications in image processing. It can be used as a model of an object, as well as a probability density functon for random noise, the latter of which is described in further detail in Section 2.5.2. It should be pointed out that because the Fourier transform of a Gaussian function is also a Gaussian function, it is compact in both the spatial and frequency domains.

2.3 Linear systems

The theory of linear systems provides models for many physical processes. Although the property of linearity is an assumption, it is a good approximation for most applications. It has the advantage of being mathematically tractible and provides a means for analysing complex processes.

The process of image degradation is modelled by a linear system. Let the input signal be $f(x, y)$, and the output signal be $g(x, y)$. Then the transformation which maps f onto g is H:

$$g(x, y) = H\{f(x, y)\}. \tag{2.17}$$

One of the most important properties of a linear system is superposition. This states that if the input to a system is a sum of a number of individual signals, then the output is the sum of the outputs of the respective input signals as if these signals were acting alone. This can be illustrated mathematically as:

$$\begin{aligned} H\{c_1 f_1(x, y) + c_2 f_2(x, y)\} &= c_1 H\{f_1(x, y)\} + c_2 H\{f_2(x, y)\} \\ &= c_1 g_1(x, y) + c_2 g_2(x, y) \end{aligned} \tag{2.18}$$

where c_1 and c_2 are arbitrary constants.

The response of a linear system can be characterised by examining its impulse response, i.e. its response to a Dirac delta function. In a 2D system, it is called the point spread function (PSF). In a shift-variant system, the PSF can be defined as:

$$h(x, y; \xi, \eta) = H\{\delta(x, y; \xi, \eta)\} \tag{2.19}$$

where the PSF has a different effect on the input depending on its position. Nevertheless, in most image processing applications, it is sufficient to assume that the system is shift-invariant, and the PSF is written accordingly as:

$$h(x, y; \xi, \eta) = h(x - \xi, y - \eta). \tag{2.20}$$

A linear system with a shift-invariant property is called a Linear Shift-Invariant (LSI) system, and is the characteristic of all the systems examined and modelled in this thesis. In astronomical imaging this implies that the field of view (FOV) of the telescope is sufficiently narrow that the portion of the atmosphere through which the object is imaged is shift-invariant.

2.3.1   Convolution

The transformation of a signal by a LSI system can be denoted using the convolution operation. For a continuous signal, it is defined as [4]:

$$g(x,y) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(\xi,\eta)h(x-\xi, y-\eta)d\xi d\eta. \tag{2.21}$$

For a discrete signal (with finite extent), the convolution is represented as [8,87]:

$$\begin{aligned} g(x,y) &= \sum_{k=-M/2}^{M/2-1} \sum_{l=-N/2}^{N/2-1} f(k,l)h(x-k, y-l) \\ &= f(x,y) \odot h(x,y) \end{aligned} \tag{2.22}$$

for an input signal f of size $M \times N$, and $\odot$ denotes a two-dimensional convolution operation. The convolution operation has many useful properties, including the commutative, distributive and associative properties. The result of the convolution operation, $g(x,y)$, is in general smoother than either of the two inputs. This smoothing or lowpass filtering effect is characteristic of all real imaging systems.

2.3.2   Correlation

The cross-correlation operation is defined as [8,87]:

$$\begin{aligned} g(x,y) &= f(x,y) \odot h(-x,-y) \\ &= f(x,y) \otimes h(x,y) \end{aligned} \tag{2.23}$$

where $\otimes$ denotes the correlation operation. It has the same properties as the convolution operation except for the commutative property. If the correlation is performed on a quantity with itself, then it is known as autocorrelation.

The cross-correlation of two signals is useful in determining the similarity between them [181]. A cross-correlation value of unity implies that the two signals are identical. A cross-correlation value of almost unity indicates a close match between the signals. If the cross-correlation is of a very small value or zero, then there is little or no similarity between the two signals. This operation forms a fundamental component of iterative (blind) deconvolution algorithms.

2.4 Transform theory

Transform theory describes the process of taking a function and manipulating it in such a way that operations are more easily performed on it. Consider a simple example, where a page of text is upside-down and is difficult to read. The transformation of rotation (by 180°) makes it much easier to read without altering the contents of the page.

A more formal definition of transform theory is the process of taking a function which is expandable on a set of basis functions, and expressing it using a different set of basis functions without altering the information content.

The need for performing transforms is two-fold. Firstly, the manipulation of a function may be simplified when expressed using a different set of basis functions. For example, a convolution operation in the spatial domain is computationally intensive. By transforming the quantities into the Fourier domain, the convolution operation becomes a multiplication which is easier to perform. Secondly, certain features in the signal are more easily detected after transformation.

In this section, transforms that are relevant to this thesis are discussed. These include the Fourier transform, the Fourier-Bessel transform, and the Hankel transform.

2.4.1 Fourier transform

The Fourier transform (FT) is very important in image processing, especially image restoration, because it describes many physical processes in an imaging system. For instance, the wavefront at the aperture and that at the focal plane are related by the Fourier transform. The Fourier transform is also an ideal tool for analysing linear systems because its basis functions, which are a set of complex exponentials, are also the eigenvectors of LSI systems.

In this section, the Fourier transform is defined and some of its properties are discussed. For a more detailed treatment, the reader should consult the books by Papoulis [132] and Bracewell [14].

The two-dimensional Fourier transform of a complex variable $f(x, y)$ is defined as [78]:

$$\mathfrak{F}\{f(x,y)\} = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x,y)\exp\{-i2\pi(ux + vy)\}dxdy$$
$$= F(u, v) \qquad (2.24)$$

where (u, v) are appropriately normalised coordinates in the frequency domain [61], and $\mathfrak{F}\{\cdot\}$ is the Fourier transform operator. The Fourier transform of f, F, is also called the spectrum of f. The inverse Fourier transform (IFT) is defined as [78]:

$$\mathfrak{F}^{-1}\{F(u,v)\} = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} F(u,v)\exp\{i2\pi(ux + vy)\}dudv$$
$$= f(x, y). \qquad (2.25)$$

The notation

$$f(x, y) \rightleftharpoons F(u, v) \qquad (2.26)$$

indicates that $f(x, y)$ and $F(u, v)$ are a Fourier transform pair and this relationship is unique.

The Fourier transform has many useful properties, some of which are listed below.

Separability

A two-dimensional Fourier transform of $f(x, y)$

$$\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x,y)\exp\{-i2\pi(ux + vy)\}dxdy \qquad (2.27)$$

can be separated into two one-dimensional FTs in the $x-$ and $y-$ directions:

$$\int_{-\infty}^{\infty} \left(\int_{-\infty}^{\infty} f(x,y)\exp\{-i2\pi ux\}dx \right) \exp\{-i2\pi vy\}dy. \qquad (2.28)$$

Superposition

If $f_1(x, y) \rightleftharpoons F_1(u, v)$ and $f_2(x, y) \rightleftharpoons F_2(u, v)$, then

$$c_1 f_1(x, y) + c_2 f_2(x, y) \rightleftharpoons c_1 F_1(u, v) + c_2 F_2(u, v) \qquad (2.29)$$

where c_1 and c_2 are arbitrary constants.

Time scaling

If $f(x, y) \rightleftharpoons F(u, v)$, then

$$f(ax, ay) \rightleftharpoons \frac{1}{|a|^2} F\left(\frac{u}{a}, \frac{v}{a}\right) \qquad (2.30)$$

where a is a scaling constant. This property states that if the function in the spatial domain is expanded by a factor of a, then its Fourier transform is compressed correspondingly by a factor of a. This property is often associated with the sampling of a signal. For example, if a is reduced from 1 to 0.5, then the sampling in the spatial domain is twice as dense in each dimension, and the bandwidth of the spectrum is expanded by a factor of two.

Duality

If $f(x, y) \rightleftharpoons F(u, v)$, then

$$F(x, y) \rightleftharpoons f(-u, -v). \tag{2.31}$$

Time and frequency shifting

If $f(x, y) \rightleftharpoons F(u, v)$, then

$$f(x - x_0, y - y_0) \rightleftharpoons F(u, v)\exp\{-i2\pi(ux_0 + vy_0)\} \tag{2.32}$$

where x_0 and y_0 are shifts in the $x-$ and $y-$directions. This means that shifting an image in the spatial domain does not alter the magnitude of the Fourier transform, but introduces a linear phase shift.

Alternatively, if the shift is introduced in the frequency domain, then the exponential factor is present in the spatial domain:

$$\exp\{i2\pi(u_0x + v_0y)\}f(x, y) \rightleftharpoons F(u - u_0, v - v_0). \tag{2.33}$$

Conjugation

If $f(x, y) \rightleftharpoons F(u, v)$, then

$$f^*(x, y) \rightleftharpoons F^*(-u, -v) \tag{2.34}$$

where the conjugation operation is denoted by $*$.

Convolution and correlation

If $f_1(x, y) \rightleftharpoons F_1(u, v)$ and $f_2(x, y) \rightleftharpoons F_2(u, v)$, then

$$f_1(x, y) \odot f_2(x, y) \rightleftharpoons F_1(u, v)F_2(u, v). \tag{2.35}$$

The convolution of two signals in the spatial domain transforms into the multiplication of the Fourier transforms of the individual signals in the frequency domain, an operation which

is easier and faster to manipulate. This property is used frequently in this thesis. As for the correlation of two signals:

$$f_1(x,y) \otimes f_2(x,y) \rightleftharpoons F_1(u,v)F_2^*(u,v). \tag{2.36}$$

2.4.2   Discrete Fourier transform, the fast Fourier transform, and the fastest Fourier transform in the West

The discrete Fourier transform (DFT) is required to implement the Fourier transform digitally, where discretisation must be performed. The two-dimensional $N-$point DFT of $f(x,y)$ is defined as [78]

$$F(u,v) = \frac{1}{N} \sum_{x=0}^{N-1} \sum_{y=0}^{N-1} f(x,y)\exp\left\{ \frac{-i2\pi(ux+vy)}{N} \right\} \tag{2.37}$$

while the corresponding inverse DFT is

$$f(x,y) = \frac{1}{N} \sum_{u=0}^{N-1} \sum_{v=0}^{N-1} F(u,v)\exp\left\{ \frac{i2\pi(ux+vy)}{N} \right\} \tag{2.38}$$

where N is the number of samples in one dimension. The assumption made when performing a DFT is that the underlying signal is periodic, which can introduce the problem of spectral leakage. One useful application of the DFT is its ability to diagonalise circulant matrices, which are approximations to Toeplitz matrices such as the convolution matrix $\mathbf{H}$ in Eq. (1.5) [30]. This enables the fast implementation of multiplications such as $\mathbf{Hf} = \mathbf{d}$.

The fast Fourier transform (FFT) is the efficient implementation of the DFT and has been studied extensively [33]. The fastest algorithms are those that employ the 'divide and conquer' strategy [78] and break down the computation into many small DFT steps. This is achieved by diagonalising the transformation matrix and exploiting its symmetry, enabling a fast implementation of matrix multiplication. Whilst the DFT operation on a complex signal is of order N^2, the FFT is of order $N\ln(N)$. The Cooley-Tukey and the split-radix algorithms are examples of FFT implementations.

Recently, an adaptive software architecture for the FFT has been developed [48]. This is called the fastest Fourier transform in the West (FFTW) with a C subroutine library for computing the DFT [49]. This has been shown to have superior performance to many

existing FFT implementations on various platforms. The major difference between the FFTW and existing implementations is that the former is adaptively optimised and does not require the array size to be a power of 2. FFTW has been incorporated in Matlab 6 [117], the language in which most of the algorithms in this thesis are developed.

2.4.3   Fourier-Bessel transform and Hankel transform

A special case of the Fourier transform of a function with circular symmetry is called the Fourier-Bessel transform, or the Hankel transform of zero order. The Fourier-Bessel transform of a function $f(r)$ is defined as [61]:

$$F(\rho) = 2\pi \int_0^\infty r f(r) J_0(2\pi r \rho) dr \qquad (2.39)$$

where $r = \sqrt{x^2 + y^2}$ and $\rho = \sqrt{u^2 + v^2}$ are the radii in the spatial and frequency domains respectively. Similarly, the inverse Fourier transform, or the inverse Fourier-Bessel transform of a circularly symmetric function is:

$$f(r) = 2\pi \int_0^\infty \rho F(\rho) J_0(2\pi r \rho) d\rho. \qquad (2.40)$$

Thus, for circularly symmetric functions, there is no difference between the forward and inverse transform kernels.

The function $J_0(\cdot)$ is the zeroth order Bessel function, and the standard definition of a Bessel function for all real orders p is [35]:

$$J_p(x) = \sum_{k=0}^\infty \frac{(-1)^p x^{p+2k}}{2^{p+2k} k! \Gamma[p+k+1]}, \quad p \neq -1, -2, -3, \ldots \qquad (2.41)$$

$$J_{-n} = (-1)^n J_n(x), \quad n = 1, 2, 3, \ldots . \qquad (2.42)$$

The Gamma function $\Gamma[n]$ is defined as [35]:

$$\Gamma[n] = \int_0^\infty t^{n-1} \exp\{-t\} dt, \quad n > 0 \qquad (2.43)$$

for positive values of n. For the negative values of n, the Gamma function is defined via the equation:

$$\Gamma[n+1] = n\Gamma[n]. \qquad (2.44)$$

(a)   (b)
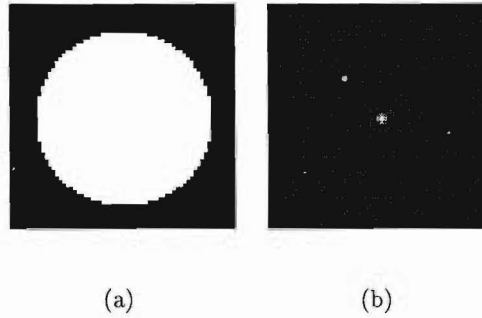
Figure 2.8: *An example of (a) the circle function and (b) its Fourier-Bessel transform.*

An example of the application of the Fourier-Bessel transform is the Fourier transform of the circle function (defined in Section 2.2). The Fourier-Bessel transform pair is:

$$\text{circ}(r) \rightleftharpoons \frac{J_1(2\pi\rho)}{\rho} \tag{2.45}$$

where J_1 is the Bessel function of the first kind, order one [61]. In a physical sense, while the circle function describes a circular aperture, its Bessel transform describes the complex wavefront at the focal plane of the system, and is useful for investigating the PSF of a circular aperture. An example of the circle function and its Fourier-Bessel transform are shown in Figs. 2.8(a) and 2.8(b).

2.5   Statistics and random processes

Deterministic signals are those that can be modelled precisely as specified functions. However, many real world signals are not deterministic but random in nature. They cannot be predicted precisely in advance, but can be described by certain statistical properties. A ubiquitous example of a random signal is thermal noise in the circuits of an image-acquiring instrument [4], which can be sufficiently described by a zero-mean Gaussian process with a uniform power spectrum.

The statistical characterisation of random signals is known as probability theory [78] and is discussed in this section. Other aspects of random signals relevant to this thesis are also presented.

2.5.1   Probability theory

Probability theory in its simplest form can be explained by a random experiment. Such an experiment has an outcome that is subject to chance and cannot be repeated. However, if a large number of trials are performed, then an average pattern of the outcome is observed. The sample space S contains all the possible outcomes of an experiment, and a particular outcome of an experiment is called a sample point s.

Mathematically, probability theory can be defined by a set of axioms [78]:

1. The probability of the sample space is unity: $P(S) = 1$.

2. The probability of an event A is between zero and unity: $0 \leq P(A) \leq 1$.

3. If the events A and B are two mutually exclusive events, then $P(A + B) = P(A) + P(B)$.

Conditional probability

The conditional probability of B given A is the probability of event B occuring given that event A has occured already. Provided that A has nonzero probability, then the conditional probability $P(B|A)$ is:

$$P(B|A) = \frac{P(AB)}{P(A)}. \qquad (2.46)$$

In many image processing applications, the conditional probability is usually unavailable. However, applying Bayes' rule allows this to be written as [4, 160]:

$$P(B|A) = \frac{P(A|B)P(B)}{P(A)} \qquad (2.47)$$

where each of the three terms on the right-hand side can be calculated more easily. In the special case where the events A and B are statistically independent i.e. $P(AB) = P(A)P(B)$, then $P(B|A) = P(B)$. Bayes' rule is of considerable importance in solving inverse problems in image processing. The term $P(A|B)$ corresponds to the consistency between the data and the reconstruction, and $P(B)$ is the prior probability of the object to be reconstructed. $P(A)$ is the prior probability of the data, which in the case of image

restoration is equal to 1. The application of Bayes' rule in image restoration is discussed further in Section 3.3.2.

2.5.2   Random variables

A random variable (RV) is a possible outcome of an experiment whose range is a set of real numbers. If the outcome of an experiment is s, then the associated RV is $X(s)$. A more general definition is to consider a random variable X and the probability of the event $X \leq x$. The cumulative distribution function (CDF), or simply the distribution function, is defined as:

$$F_X(x) = P(X \leq x). \tag{2.48}$$

There are many situations in image processing involving more than one random variable. In this case, the joint distribution function for the RVs X and Y is:

$$F_{X,Y}(x,y) = P(X \leq x, Y \leq y). \tag{2.49}$$

An important property of the CDF is that its derivative gives the probability density function (PDF) of the random variable(s). In two dimensions, the joint PDF is:

$$\frac{\partial^2 F_{X,Y}(x,y)}{\partial x \partial y} = f_{X,Y}(x,y). \tag{2.50}$$

Common forms of PDFs in image processing include the Gaussian PDF [4], the Poisson PDF [101,158], the Gibbs PDF [73] and the Markov random field [53]. The Gaussian PDF has the same form as Eq. (2.15) and is repeated here for completeness. For a variable x it is defined by [4]:

$$p(x) = \frac{1}{\sigma_X \sqrt{2\pi}} \exp \left\{ -\frac{1}{2} \left(\frac{x - \mu_X}{\sigma_X} \right)^2 \right\} \tag{2.51}$$

with the mean μ and the variance σ_X^2. σ_X is the standard deviation. The Poisson PDF is defined by [101,158]:

$$p(x) = \frac{\lambda^x}{x!} \exp\{-\lambda\} \tag{2.52}$$

where the mean and the variance are both equal to λ.

2.5.3   Statistical quantities

This section defines the statistical quantities that are relevant to this thesis.

Mean

The mean, or expected value, of the random variable X is [78]:

$$\mu_X = E\{X\} = \int_{-\infty}^{\infty} x f_X(x) dx \tag{2.53}$$

where $E\{\cdot\}$ is the expectation operator. In more general terms, let us define a function of X, $g(X)$, whose output Y is also a random variable. Then the PDF of Y can be written as [78]:

$$E\{g(X)\} = \int_{-\infty}^{\infty} g(x) f_X(x) dx. \tag{2.54}$$

Moments

The n–th moment of the probability distribution of the random variable X is defined as [78]:

$$E\{X^n\} = \int_{-\infty}^{\infty} x^n f_X(x) dx. \tag{2.55}$$

The most commonly used moments are the first and second moments, also known as the mean and the mean-square value of X. The central moment, on the other hand, takes into account the mean of the process μ_X:

$$E\{(X - \mu_X)^n\} = \int_{-\infty}^{\infty} (x - \mu_X)^n f_X(x) dx. \tag{2.56}$$

In the case where $n = 2$, the central moment is also called the variance σ_X^2. The variance essentially describes the variation of the samples about its mean. In a zero-mean process, the variance and the mean-square value are identical.

Joint moments

For two random variables X and Y, the joint moment is defined by [78]:

$$E\{X^j Y^k\} = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} x^j y^k f_{X,Y}(x,y) dx dy \tag{2.57}$$

with j and k being positive integers. The most important joint moment is the correlation where $j = k = 1$. This can be calculated using the Fourier transform as discussed in Section 2.4.1. The correlation between two RVs is called the cross-correlation, while the

correlation of a RV with itself is termed autocorrelation. The covariance corresponds to the mean-removed correlation where:

$$\text{cov}\{XY\} = E\{(X - E\{X\})(Y - E\{Y\})\}. \tag{2.58}$$

Similarly, the cross-covariance describes the covariance of two RVs and the autocovariance is the covariance between a RV and itself. The autocovariance characterises any statistical relationship between different samples of a RV, and can be used as a form of prior information in the estimation of a RV. The application of covariances to inverse problems in image processing is discussed in Chapters 5, 6 and 8.

2.6 Handling real world data

In order to process real world data on a digital computer, several operations must first be performed on the data. This section discusses these processes including sampling, quantisation, and the associated practical problems. The characteristics of real world signals, such as them having finite energy, are also presented. In processing real astronomical data, there are other factors that have to be taken into account. These are discussed in Chapters 4 and 8.

2.6.1 Sampling and aliasing

The sampling process is inherent to digital image processing where a real world, analogue signal is converted into a discrete sequence of samples. The process of area sampling is discussed here where the signal is convolved with the point spread function of the detector cell before sampling by the Dirac comb defined in Section 2.2. The sampled image $f_s(x, y)$ of the continuous image $f(x, y)$ is defined as:

$$\begin{aligned} f_s(x,y) &= f(x,y) \odot h^c(x,y) \times \text{comb}(x,y) \\ &= \left(f(x,y) \odot \text{rect}(x,y) \right) \sum_{m=-\infty}^{\infty} \sum_{n=-\infty}^{\infty} \delta(x - mX_0, y - nY_0). \end{aligned} \tag{2.59}$$

$h^c(x,y)$ describes the point spread function of a detector cell with uniform sensitivity. In the Fourier domain, the comb function is written as [4]:

$$\mathfrak{F}\{\text{comb}(x,y)\} = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \sum_{m=-\infty}^{\infty} \sum_{n=\infty}^{\infty} \delta(x - mX_0, y - nY_0) \exp\{-i2\pi(ux + vy)\} dx dy$$

$$= \frac{4\pi^2}{X_0 Y_0} \sum_{m=-\infty}^{\infty} \sum_{n=\infty}^{\infty} \delta(u - \frac{2\pi m}{X_0}, v - \frac{2\pi n}{Y_0}) \qquad (2.60)$$

and the Fourier transform of the signal can be written as $F(u, v)$. Then using the convolution-multiplication property of the Fourier transform, the FT of the sampled signal is:

$$\mathfrak{F}\{f_s(x,y)\} = F(u,v) \times \mathfrak{F}\{h^c(x,y)\} \odot \mathfrak{F}\{\text{comb}(x,y)\}$$

$$= \frac{4\pi^2}{X_0 Y_0} \sum_{m=-\infty}^{\infty} \sum_{n=\infty}^{\infty} F(u - \frac{2\pi m}{X_0}, v - \frac{2\pi n}{Y_0}) \ \text{sinc}(u - \frac{2\pi m}{X_0}, v - \frac{2\pi n}{Y_0})$$

$$= F_s(u, v). \qquad (2.61)$$

Since the sampling process does not record every point of the original signal, there is information loss in the process. However, provided that the samples are taken sufficiently close, it is possible to recover the signal exactly. This is called the Nyquist criterion and states that the sampling rate must be at least twice that of the highest frequency component of the signal [9]. Mathematically it is expressed as [85]:

$$X_0 = \frac{1}{2f_H}$$

$$= \frac{D}{2\lambda f_l} \qquad (2.62)$$

where f_H is the highest frequency component of the data and X_0 is the sampling rate. The second equation relates to a practical imaging situation, where D is the diameter of the aperture, f_l is the focal length of the optics, and λ is the wavelength of the optical wavefront.

Fig. 2.9 illustrates the phenomenon of aliasing. (This is the same as Fig. 1.4 and is reproduced here for completeness.) Fig. 2.9(a) is the spectrum of a 1D bandlimited signal, and Fig. 2.9(b) shows the spectrum of the Nyquist sampled signal. In order to retrieve the signal spectrum, a window can be used to lowpass filter the desired portion of the replicated spectrum. However, if the signal is sampled beneath the Nyquist rate, some of the high frequency components are shifted and added to the existing low frequency components. As a result, the separation of these frequency components is difficult.

(a)   (b)   (c)
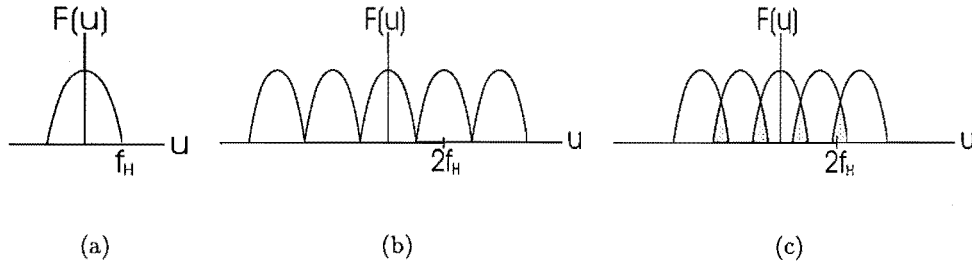
Figure 2.9: *Illustrating the problem of aliasing. (a) Spectrum of a one-dimensional band-limited continuous signal. (b) Spectrum of the critically sampled signal. (c) Spectrum of an undersampled signal. The grey areas indicate the higher frequency components being shifted to overlap with the lower frequency components.*

Another example illustrating the effect of aliasing is to consider two sinewaves at 1 rad/s and 11 rad/s. To meet the Nyquist criterion, they must be sampled at at least 2 rad/s and 22 rad/s respectively. Fig. 2.10 shows the effect of sampling both signals at 10 rad/s: the sample points are denoted by the symbol 'O' and the samples for both signals coincide. Generally, the low frequency signal is reconstructed from these samples, and the information regarding the high frequency signal is lost.
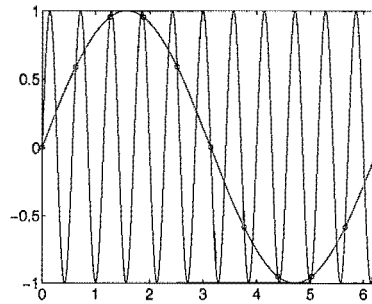
Figure 2.10: *The aliasing phenomenon. Two sinewaves of 1 rad/s and 11 rad/s produce identical sample values when sampled at 10 rad/s.*

2.6.2   Quantisation

After the signal has been sampled, the sampled values are usually stored by a digital instrument. Since the instrument can only take on a discrete number of values, the sampled
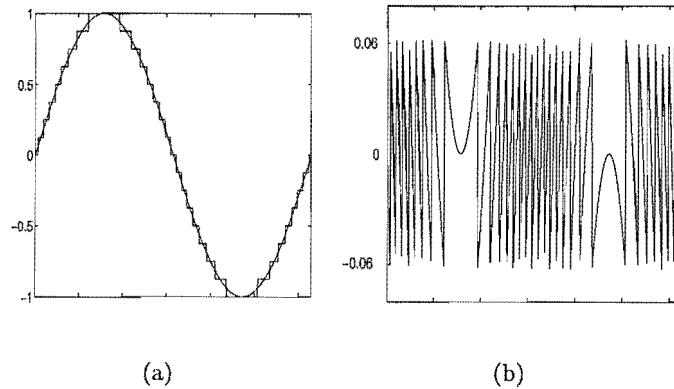
Figure 2.11: *(a) A sine wave and its quantised version. (b) Associated quantisation error.*

values must be quantised. Quantisation is therefore defined as the conversion of an analogue number to a digital number where the possible values are predetermined by the configuration of the instrument. The device which performs quantisation is called an analogue to digital converter (ADC).

The process of quantisation differs from that of sampling in the sense that a range of possible amplitudes are converted to a single digital number within the computer. It is a many-to-one mapping where information is lost. However, if the signal is sampled by a digital instrument, then the two processes are analogous.

In a quantisation process, the amplitude of the analogue signal is compared to a set of decision levels. Each sample is then assigned a number which is closest to the analogue value. The process of quantisation is characterised by two variables: the step size Δ, and the number of steps. In a gray-scale image, it is common to use 8-bits (with $2^8 = 256$ levels) to represent the image.

The major problem associated with uniform quantisation is the quantisation noise, or quantisation error. Fig. 2.11(a) shows a sine wave and its quantised version using 17 levels. The difference between the two signals is shown in Fig. 2.11(b). This error is approximately uniformly distributed and has a variance of $\Delta^2/12$ [78].

2.6.3   Noise on the detector

The corruption of the measurements by noise is inevitable. There are two major sources of noise in a photoelectronic system. Firstly, there is random thermal noise as a result of the resistive circuit elements. This noise process is randomly distributed about zero, and can be described by a Gaussian PDF (see Section 2.5.2). It is therefore often referred to as white Gaussian noise. Other terminologies that are used interchangeably include read noise and thermal noise.

The second source of noise arises from the statisical nature of light and is particularly relevant at low light levels. Due to the small number of photons, this noise process is usually modelled by a Poisson PDF as defined in Section 2.5.2. As the light level increases, the Poisson noise process can be approximated by the Gaussian noise process.

2.6.4   Finite energy, bandlimited images and compactness

Real world signals have certain characteristics that are important in image processing. Since the signals can only have finite amplitude, this implies that they can only have finite energy also. The energy of a signal $f(x, y)$ is:

$$E_f = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} |f(x, y)|^2 dx dy. \tag{2.63}$$

Parseval's theorem states that the energy of a signal in the image domain is the same as that in the Fourier domain [160]:

$$\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} |f(x, y)|^2 dx dy = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} |F(u, v)|^2 du dv \tag{2.64}$$

where $|F(u, v)|^2$ is also known as the power spectrum. This property is crucial in the design of image processing algorithms, because it ensures that an error which has a small energy in the image domain also has a small energy in the frequency domain. A common measure of the quality of an image is the signal-to-noise ratio (SNR), defined as the ratio between the energy of the signal $f(x, y)$ to that of the noise $n(x, y)$:

$$SNR = 10 \log_{10} \left\{ \frac{E_f}{E_n} \right\} \quad dB. \tag{2.65}$$

An image with a finite frequency content is called a bandlimited image. In reality, all images are bandlimited because they do not have infinite resolution. A signal which is finite
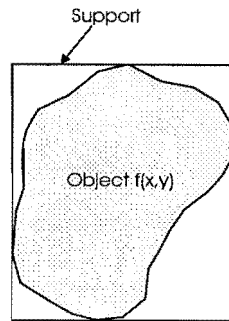
Figure 2.12: *The support of an object in the image domain.*

in extent in the image domain is said to be compact. The extent of such an object can be defined by a support, which is the smallest rectangle that encloses the entire object. An example is shown in Fig. 2.12. The support of an object has been used as an *a priori* constraint in many image restoration algorithms [90,102,105,123,168,183]; these are discussed further in Chapter 6.

2.6.5 Positivity

The intensity or light emitted from an object must be positive; however when these values are recorded by an instrument, the read noise on the detector can result in a negative value being recorded. Nevertheless, the prior knowledge that intensity values must be positive can be used when reconstructing the object from the measurements. Enforcing positivity in a reconstruction increases the computational complexity, but reduces the possible number of solutions to the problem, and is a common constraint used in image restoration algorithms. This constraint is used in a number of algorithms investigated in this thesis.

Chapter 3

Inverse problems: an overview

The ultimate goals of modelling a physical system are to both understand its workings and predict its behaviour. These systems range from an electrical circuit, atmospheric turbulence, the human mind, the stock market, to the universe itself. In order to achieve this goal, we must create a model of the system and in this process, an approximation is inevitable since we can never model a real system perfectly. The quality and the accuracy of a model to the real world system are difficult to assess. In general, if a model can predict what happens in the real world, then it is a good model. Usually, a model is used until it breaks down, or fails to predict the behaviour of the real system; then a better model must be devised. For example, as long ago as 340 B.C., the Greek philosopher Aristotle proposed that the Earth was the centre of the universe. Ptolemy developed this model further, which was reasonably accurate for describing the movements of the planets. Although he did realise that there was a flaw in the model, it remained the best model available until 1514 A.D. when Copernicus postulated the idea of the Earth and the planets orbiting the Sun. Many years later, with much controversy, this model was finally accepted, as it provided a more consistent description of the real system compared with the astronomical measurements taken.

In this chapter, we formalise the process of modelling real world systems through the investigation of inverse problems. Inverse problems are defined in Section 3.1 and the feasibility of solving such problems are examined in Section 3.2. An overview of the techniques for solving inverse problems is given in Section 3.3. Methods devised for dealing with real world

systems, where the ideal properties attributed to theoretical systems no longer hold, are discussed in Section 3.4.

3.1   What exactly are inverse problems?

Any real system has an infinite number of parameters, but the amount of measurable data is finite. As a result, only finite models can be used. Consider the situation depicted in Fig. 3.1, where the relationship between a real world system, a model of the system, and the data (or measurements) is illustrated. From our available data, we would like to infer a description of the real system via a solution of an inverse problem (dashed line). Unfortunately, this is not possible, because the real world system is imperfectly known and can never be defined by a finite number of measurements. As a result, we devise a finite model of the real system, and from the measurements, we infer the model parameters. This is the inverse problem that we solve. On the other hand, the forward problem is the prediction of the data values given the model parameters, and a good understanding of the forward problem facilitates solving the inverse problem.

Since a real world system has an infinite amount of information, and that only a finite amount of it is measurable, there is always a loss of information in the process. This leads to the question: is the data available sufficient for determining a solution via the model? And if so, is the solution unique? Consider a simple example: $a + b = 6$. There are two unknowns in the system, and only one data point which is their sum. This is analogous to the problem of aliasing, where two different frequency components of a data set sum to form the same frequency measurement. It is obvious that there are an infinite number of frequency pairs that are consistent with the data.

In order to find a solution to this problem, extra information is required. For example, the probability distribution of the values of a and b might be known. In this case, we can find the most likely values of a and b consistent with the observation. Sometimes, the physical limitation of a system can be used to restrict the solution set, an example of which in astronomical imaging is the aperture of the telescope. Since the physical dimensions of the aperture cut out all the light outside it, it is reasonable to assume that the image formed at the telescope is from the light captured within the aperture only. It should be emphasised that, although in many situations, there exists an infinite number of solutions
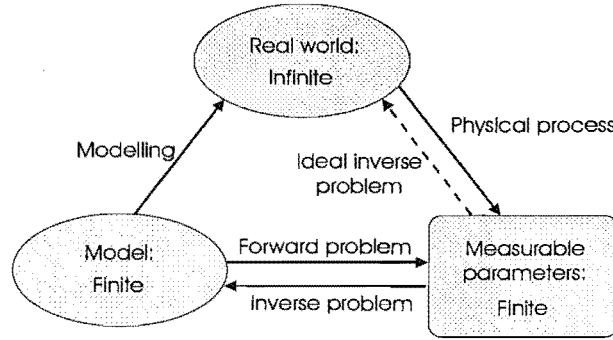
Figure 3.1: *A depiction of the relationship between the real world and the model we use to approximate the real world.*

to an inverse problem, often these solutions are very similar; and for practical purposes, they are essentially the same. It is when such a *representative* solution for the problem is possible, that inverse problems can be viewed as solvable.

Once it is established that there is a representative solution to an inverse problem, the next issue is the quality of the solution. Consider a linear problem characterised by

$$\mathbf{d} = \mathbf{Hf},\tag{3.1}$$

where $\mathbf{d}$ is the measurement, $\mathbf{f}$ is the unknown to solve for, and $\mathbf{H}$ is the matrix describing the transformation. Firstly, the solution to this problem should be stable. This means that small changes in the data values $\mathbf{d}$ should not produce large changes in the solution for $\mathbf{f}$. This is important because in a real system, there are always random errors on the data due to the presence of noise, and it is only possible to find a representative solution if this noise is not amplified excessively when estimating $\mathbf{f}$. Secondly, the solution should be meaningful. For example, if the solution to a problem does not depend on the data at all, then even though the uniqueness and stability requirements are met, it gives no information about the system at all. (In fact, the system is so stable that no matter what happens to the data, the solution remains the same!) In other words, the solution must still be influenced by the data. A balance must be achieved between these often conflicting requirements. A solution which fulfils these criteria is called a well-posed solution, and the corresponding problem is called a well-posed problem.

These ideas can be expressed mathematically as follows [68, 69]. Consider the same system as described by Eq. (3.1), where $\mathbf{f}$ and $\mathbf{d}$ belong to respective Hilbert spaces $\mathbf{H}_1$ and $\mathbf{H}_2$,

and that **H** is a bounded linear operator. This problem is defined to be well-posed if:

(1) for every element **d** $\in$ **H₂** there exists a solution in the space **H₁**;

(2) the solution is unique;

(3) the problem is stable on all spaces (**H₁, H₂**), i.e. the solution depends continuously on the data; and that small perturbations in the data do not cause large changes in the solution.

Otherwise the problem is ill-posed.

It is a well-posed problem that we would like to solve, but many systems are ill-posed. In order to deal with these problems, they must be regularised first. Katsaggelos defined regularisation as 'the analysis of an ill-posed problem via an analysis of an associated well-posed problem, whose solution yields meaningful answers and approximations to the given ill-posed problem' [95]. Essentially, the application of regularisation should either reduce the number of possible solutions, and/or stabilise the solution. The restriction of an object to have only positive intensity values is one example of reducing the solution set, while assuming that the object is smooth is one way of stabilising the solution. Ultimately, the objective of using regularisation is to combine data information and prior information to achieve a smooth solution which remains faithful to the data. It should be clarified that the term smooth does not imply flat and broad; instead it implies that the noise is not amplified. This is especially true for astronomical imaging where the object may be point sources.

3.2 Are inverse problems possible?

Consider an image of the binary star Castor captured by a charge-coupled device (CCD) camera as shown in Fig. 3.2(a). This is one of a series of images captured. The object has been degraded by the effect of the atmosphere and the optics, and insufficient sampling was employed in capturing the data. As a consequence it is impossible to distinguish the binary star in the data. By applying inverse problem techniques, a clear binary object (Fig. 3.2(b)) is recovered.

At first glance, it seems impossible to achieve the reconstruction. In order to prove the feasibility of such an inverse problem, let us first examine a noiseless case. Let $f(x, y)$ be the object of interest of size $M \times M$, $h_k(x, y)$ be the k^{th} realisation of the point spread

(a)   (b)
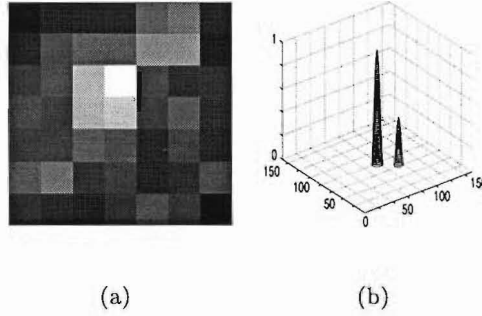
Figure 3.2: *(a) A sample frame of the binary star Castor. The binary star cannot be distinguished. (b) The application of inverse problem techniques has reconstructed a binary star.*

function, each of size $N \times N$, and $d_k(x, y)$ be the k^{th} convolution:

$$d_k(x, y) = f(x, y) \odot h_k(x, y). \tag{3.2}$$

Each of the K convolutions is therefore of size $(M + N - 1)^2$. In the inverse problem of blind deconvolution, where f and h are unknown, there are a total of $M^2 + KN^2$ unknowns. On the other hand, the number of known variables is

$$
\begin{aligned}
K(M + N - 1)^2 &\approx K(M + N)^2 \\
&= KM^2 + KN^2 + 2KMN \\
&> M^2 + KN^2 \quad \forall\ K.
\end{aligned}
\tag{3.3}
$$

From this simple analysis, it appears that this inverse problem is possible, because there are more data points than unknowns.

Since it is easier to determine the number of unknowns compared to the number of measurements, instead of classifying a problem to be well-posed or ill-posed, often linear problems are defined as either overdetermined or underdetermined. An overdetermined system is one which has more measurements than unknowns, such as the example above; while the converse is true for an underdetermined system. An exactly determined system is when the number of unknowns equals the number of measurements; this does not usually happen in real life but is a useful means for analysing associated problems.

Returning to the simple analysis above: although the number of unknowns is less than the number of measurements, this analysis is deceptive because the noise is not taken into

account. When noise is present, the measurements with values below the noise floor are not representative, and the number of measurements is reduced as a result. Moreover, it is possible that not all measurements are independent. In this case, a meaningful solution is only possible if some measures are taken to ensure that the problem is not ill-posed.

In view of real world systems, where the presence of noise is inevitable, the feasibility of an inverse problem must be re-evaluated. In this case even in an overdetermined system, regularisation may need to be included in order to obtain a good solution. On the other hand, it is also possible for an underdetermined system to have a solution. This can be achieved by employing prior information in the reconstruction, such as the aliasing example mentioned earlier on in this chapter, or by parameterising a solution. If the object we wish to reconstruct has N unknowns, but is essentially smooth and can be represented by fewer parameters $p \ll N$, then a solution can be obtained. An example is the parameterisation of the denominator in the DWFS framework to obtain an object intensity distribution, which is discussed in Chapter 9.

Sometimes, only a partial model can be reconstructed. For example, the reconstruction of an image is of lower-resolution than desired due to the lack of data. This is another example of a solution obtained from an underdetermined system, and illustrates the concept of a representative solution.

A more rigorous approach to determining the feasibility of the inverse problem discussed above is to use the concept of zero sheets [107]. It has been shown that for a convolution with a dimension greater than unity, blind deconvolution on a noiseless frame of data is possible in the absence of any prior information. This method can be summarised as follows. For a two-dimensional object $f(x, y)$, there exist a corresponding Fourier transform $F(u, v)$ and a Z-transform $\mathcal{F}(\zeta, \gamma)$ [127]. Assuming that all the quantities in all these spaces have compact support, then there exists a zero map of $\mathcal{F}(\zeta, \gamma)$ which completely characterises $f(x, y)$ to within a complex multiplicative constant. In other words, for any K, $\mathcal{F}(\zeta, \gamma)$ is zero on a surface in $2K$-dimensional Z space. This surface is called the zero sheet and is of dimension $2K - 2$. The recovery of the object $f(x, y)$ given its zero sheet is achieved by firstly evaluating the Fourier spectrum $F(u, v)$ point-by-point, and then taking an IFT. An alternative evaluation can be obtained via the method of Curtis et al. [28]. Although this means of analysing the inverse problems produces a number of theoretical results, the presence of noise causes difficulties for such an analytic approach. In practice, a statistical

approach is better suited to real world data.

3.3 Classes of solution to inverse problems

Most of the existing methods for solving inverse problems can be classified into one of the three categories of algorithms: maximum likelihood (ML), maximum *a posteriori* (MAP), and maximum entropy (ME) methods. In this section, these three categories are defined, and the relationship between them is explained.

3.3.1 Maximum likelihood methods

If we know that our individual errors and fluctuations
follow the magic bell-shaped curve exactly,
then the resulting estimates are known to have
almost all the nice properties that people have been able to think of.
- J.W. Tukey, 1965.

As early as 1816, the method of maximum likelihood was already proposed, but was not known by this name [70]. The first item of literature that recorded the maximum likelihood method was the paper by Fisher in 1912 [37]. His postulation is as follows. Let f be a function involving arbitrary elements $\vartheta_1, \vartheta_2, \ldots, \vartheta_r$, and that x is the abscissa. If f is a function of unit area, then $p = f \delta x$ is the probability of an observation falling within the range δx. Since δx is independent of the underlying function we are trying to estimate, the probability of any particular set of ϑ's is proportional to P, where

$$\log P = \sum_1^n \log f, \tag{3.4}$$

with n being the number of observations. This means that the most probable set of values for the ϑ's will make P a maximum. Essentially, the maximum likelihood approach is used when we wish to infer the underlying function when only a set of data is given, and that no prior information about the function is available.

From a more pragmatic point of view, the maximum likelihood method is essentially a

'best-fit' method. Consider a system described by:

$$\mathbf{d} = \mathbf{Hf} + \mathbf{n} \tag{3.5}$$

where $\mathbf{f}$ is the object of interest, $\mathbf{H}$ is the point spread function matrix, $\mathbf{n}$ is the noise on the system, and $\mathbf{d}$ is the noisy measurement. Given $\mathbf{d}$ and $\mathbf{H}$, we wish to find an estimate of $\mathbf{f}$. If we assume no information about the object, the best we can do is to minimise some metric of the noise on the data, for example the square of the L_2 norm:

$$\min \|\mathbf{d} - \mathbf{H}\hat{\mathbf{f}}\|^2 \tag{3.6}$$

where the circumflex denotes the estimate of the quantity. This makes sense because it utilises the information available and assumes as little as possible about the unknowns. It can also be viewed as a procedure of ensuring that the object estimate is as consistent as possible with the observed data.

If we assume the statistics of the noise, we can formalise this reasoning to try and estimate the most likely noise distribution. In other words, we wish to maximise the probability $P(\mathbf{d}|\hat{\mathbf{f}})$. Let us assume the covariance of the noise to be $\mathbf{K}_{nn}$, then the expression to be maximised is:

$$E = \exp\{-(\mathbf{d} - \mathbf{H}\hat{\mathbf{f}})^T \mathbf{K}_{nn}^{-1}(\mathbf{d} - \mathbf{H}\hat{\mathbf{f}})\}. \tag{3.7}$$

To simplify this expression, the logarithm can be taken without affecting the location of the maximum of the metric. Since maximising Eq. (3.7) is equivalent to minimising the negative of it, the expression to be minimised is:

$$\begin{aligned} E &= (\mathbf{d} - \mathbf{H}\hat{\mathbf{f}})^T \mathbf{K}_{nn}^{-1}(\mathbf{d} - \mathbf{H}\hat{\mathbf{f}}) \\ &= \mathbf{d}^T \mathbf{K}_{nn}^{-1}\mathbf{d} - \mathbf{d}^T \mathbf{K}_{nn}^{-1}\mathbf{H}\hat{\mathbf{f}} - \hat{\mathbf{f}}^T \mathbf{H}^T \mathbf{K}_{nn}^{-1}\mathbf{d} + \hat{\mathbf{f}}^T \mathbf{H}^T \mathbf{K}_{nn}^{-1}\mathbf{H}\hat{\mathbf{f}}. \end{aligned} \tag{3.8}$$

In order to find the minimum, we can differentiate this expression with respect to $\hat{\mathbf{f}}$ and equate it to zero:

$$\begin{aligned} \frac{\partial E}{\partial \hat{\mathbf{f}}} &= -\mathbf{d}^T \mathbf{K}_{nn}^{-1}\mathbf{H} - \mathbf{H}^T \mathbf{K}_{nn}^{-1}\mathbf{d} + 2\mathbf{H}^T \mathbf{K}_{nn}^{-1}\mathbf{H}\hat{\mathbf{f}} \\ &= 0. \end{aligned} \tag{3.9}$$

This can be rearranged to give an estimate of $\mathbf{f}$:

$$\hat{\mathbf{f}} = (\mathbf{H}^T \mathbf{K}_{nn}^{-1}\mathbf{H})^{-1}\mathbf{H}^T \mathbf{K}_{nn}^{-1}\mathbf{d}. \tag{3.10}$$

This is called the maximum likelihood estimate, the minimum-variance unbiased estimate, or the Gauss-Markov estimate [115]. It is desirable to be able to obtain a solution in closed form such as Eq. (3.10), however in practical situations, an iterative solution is a more realistic approach. The iterative form of Eq. (3.10) is [4]:

$$\hat{\mathbf{f}} = E\{\mathbf{f}\} + \mathbf{H}^T \mathbf{K}_{nn}^{-1} (\mathbf{d} - \mathbf{H}\hat{\mathbf{f}}) \qquad (3.11)$$

where $E\{\cdot\}$ denotes the expectation operator. The second term of Eq. (3.11) is in fact a matched filtering operation. This becomes apparent if we let $\mathbf{K}_{nn} = \mathbf{I}$; then multiplying $(\mathbf{d} - \mathbf{H}\hat{\mathbf{f}})$ by $\mathbf{H}^T$ is equivalent to correlating the former with $\mathbf{H}$.

It should be pointed out that in the limiting case where the statistics of the noise is unknown, the covariance matrix $\mathbf{K}_{nn}$ is usually assumed to be an identity matrix, in which case the maximum likelihood solution simplifies to a least-squares or pseudoinverse solution. This means that the least-squares solution can be viewed as a special case in the maximum likelihood framework.

The pseudoinverse solution also led to an important result derived by Hunt (1972) [84]. He showed that if $\mathbf{H}$ is a circulant matrix, then it can be decomposed by the DFT, thus conceptually linking the algebraic and Fourier domain solutions to the restoration problem [71].

The above expressions are valid for a system described by Eq. (3.5) with Gaussian statistics. If the system is dominated by Poisson statistics, then the probability $P(\mathbf{d}|\hat{\mathbf{f}})$ is given by [78]:

$$P(\mathbf{d}|\hat{\mathbf{f}}) = \frac{\mathbf{d}^k}{k!} \exp\{-\mathbf{d}\}, \quad k = 0, 1, \ldots \qquad (3.12)$$

and the metric to be minimised is:

$$E = \sum \mathbf{H}\hat{\mathbf{f}} - \mathbf{d} \ln[\mathbf{H}\hat{\mathbf{f}}]. \qquad (3.13)$$

For the derivation of this equation, the reader is referred to the book by Snyder and Miller [158]. It should be pointed out that due to the nonlinearity involved in the minimisation, a closed form solution is not available, and an iterative procedure is necessary for obtaining a solution.

The ML estimator is also associated with the expectation-maximisation (EM) method, which was first proposed by Dempster et al [31]. The basic idea is to infer a 'complete

data set' from an 'incomplete data set' via a two-step procedure: the expectation step, and the maximisation step. Despite the mathematically involved derivations, the EM algorithm is essentially a procedure for obtaining the maximum likelihood formulation for Poisson statistics. A popular application of the EM method is positron emission tomography [96, 156].

3.3.2   Maximum *a posteriori* methods

In the maximum likelihood estimate, no information was available regarding the object to be estimated. However, if such information is available, it is intuitively reasonable to include it in the solution. This is the basic idea of a maximum *a posteriori* (MAP) solution.

An alternative approach is to consider an exactly determined maximum likelihood solution for a real system. It is obvious that there is noise on the system, however the maximum likelihood solution implies that the mostly likely situation is when there is no noise (provided that the noise is zero-mean). From a practical viewpoint, this paradox means that the noise on the system is manifested as part of the object reconstruction which is undesirable. Therefore, if prior knowledge about the object is available, it can be used to differentiate between an object reconstruction that is consistent with the prior information, and one that is corrupted by the noise.

The following political example illustrates the application of the MAP method. Consider a hypothetical situation, where New Zealand was about to have a referendum on the issue of lowering the legal drinking age from 20 to 18. Suppose that at this date, a general election had just taken place, and that the Labour party had won, with the proportion of voters voting for them at 52%, and that the National party had taken the remaining 48% of the votes. (This situation disregards all the other parties for simplicity only.) With the numerous opinion polls conducted at the time, it can be taken that 85% of Labour supporters and 55% of National supporters intended to vote 'Yes', with the remainder of the population voting 'No'. Suppose that we met a person who intended to vote 'Yes'. Then given all the information above, what is the probability of the person being a Labour

supporter? This problem can be solved by Bayes' theorem, where:

$$
\begin{aligned}
P(\text{Labour}|\text{Yes}) &= \frac{P(\text{Yes}|\text{Labour})P(\text{Labour})}{P(\text{Yes}|\text{Labour})P(\text{Labour}) + P(\text{Yes}|\text{National})P(\text{National})} \\
&= \frac{0.85 \times 0.52}{0.85 \times 0.52 + 0.55 \times 0.48} \\
&= 0.63.
\end{aligned}
\tag{3.14}
$$

This is a simple demonstration of a Bayesian estimator, and illustrates how a measurement can be used indirectly provided it is related to the desired quantity. The key issue is that we must know the prior probabilities, which, although are simple to obtain for political leanings, are very difficult in the case of objects.

In a more general inverse problem, consider the situation described by Eq. (3.5). The MAP estimate of $\mathbf{f}$, $\hat{\mathbf{f}}$, is one which maximises the *a posteriori* probability $P(\hat{\mathbf{f}}|\mathbf{d})$:

$$
\max_{\hat{\mathbf{f}}} P(\hat{\mathbf{f}}|\mathbf{d}) = \max_{\hat{\mathbf{f}}} \frac{P(\mathbf{d}|\hat{\mathbf{f}})P(\hat{\mathbf{f}})}{P(\mathbf{d})},
\tag{3.15}
$$

where $P(\mathbf{d}|\hat{\mathbf{f}})$ is the conditional *a priori* probability, and $P(\hat{\mathbf{f}})$ and $P(\mathbf{d})$ are the prior probabilities. To simplify this expression, the logarithm can be taken without affecting the position of the maximum. This results in:

$$
\max_{\hat{\mathbf{f}}} \log[P(\hat{\mathbf{f}}|\mathbf{d})] = \max_{\hat{\mathbf{f}}} \log[P(\mathbf{d}|\hat{\mathbf{f}})] + \log[P(\hat{\mathbf{f}})].
\tag{3.16}
$$

Note that the term $P(\mathbf{d})$ has been discarded, as it is independent of $\hat{\mathbf{f}}$.

The first term in Eq. (3.16) is in fact the ML estimator which describes the consistency between the observed data and the object. In other words, the ML estimator is when the prior probability distribution is the same as the posterior probability distribution, where:

$$
P(\hat{\mathbf{f}}|\mathbf{d}) = P(\mathbf{d}|\hat{\mathbf{f}}).
\tag{3.17}
$$

The second term in the equation expresses the prior probability, or simply the prior, of the object. This can be a frequency distribution of $\mathbf{f}$, which is usually unavailable; alternatively, it can be any reasonable assumption about the object. Obviously, the choice of prior depends on the problem. Examples of prior information used for astronomical imaging are a major topic of discussion throughout this thesis.

It should be emphasised that, in obtaining a solution via the MAP framework, a balance must be achieved between the two terms in Eq. (3.16). If the weighting on the first term is

too high, then the object estimate will fit to the noise, and a noisy reconstruction results. On the other hand, if the weighting of the second term is too high, the solution will be too smooth, obscuring any features in the reconstruction. Fig. 3.3 illustrates this concept. A binary object is shown in Fig. 3.3(a), and the image of the binary star distorted by atmospheric turbulence and corrupted by noise is shown in Fig. 3.3(b). A Wiener filter is used to reconstruct the object and the prior information is the signal-to-noise ratio. When too much weight is put on the prior, the reconstruction is too smooth (Fig. 3.3(c)); on the other hand, insufficient prior weighting results in a noisy reconstruction (Fig. 3.3(d)). The best reconstruction with the optimal weighting is shown in Fig. 3.3(e).

Although there are algorithms for finding the optimal weighting between the consistency measure and the prior term, more often than not, the method of trial and error is employed. Nevertheless, it should be pointed out that there is a range of values for the weighting that produce a similar reconstruction, and provided the result is not unduly sensitive to the assumptions made, the result is justified.

3.3.3   Maximum entropy methods

Both the maximum likelihood and the maximum *a posteriori* methods are linear reconstruction schemes. A different approach to the problem of object reconstruction is to apply non-linear methods, one of which is the maximum entropy method.

The concept of applying maximum entropy (ME) to image restoration was introduced in the 1970s. Although all the variations of the methods involve estimating unmeasured Fourier coefficients, there exist various views of the workings of the ME method. Ables [1] described ME as a maximally non-committal technique with regard to the unmeasured data, while Gull and Daniell [66, 67] defined ME to give the 'most probable' solution consistent with the data. Jaynes [89] explained that ME is useful for describing a class of problems where orthodox statistics fail to account for the prior information available, while Andrews and Hunt [4] had taken a combinatorial approach to the problem. Meanwhile, Högbom [79] and Subrahmanya [164, 165] treated ME as a penalty term in the restoration framework. As there are many different views regarding the maximum entropy approach, in this thesis, the two most widely accepted views are defined and discussed. The first approach is that of Jaynes, and Andrews and Hunt, where the problem formulation involves the maximisation

(a)   (b)

(c)   (d)   (e)
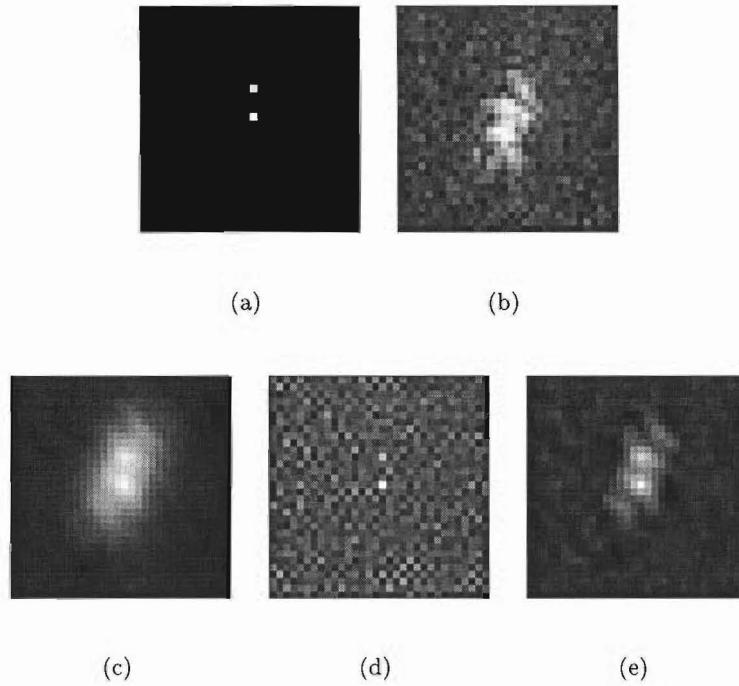
Figure 3.3: *A demonstration of the importance of weighting the prior term correctly. (a) A binary object. (b) Image of the binary star distorted by atmospheric turbulence and corrupted by noise. (c) Reconstruction with too high a weighting on the prior. (d) Reconstruction with too low a weighting on the prior. (e) Reconstruction with the optimal weighting on the prior.*

of:

$$-\mathbf{f}^T \ln(\mathbf{f}). \tag{3.18}$$

Secondly, there is a group of 'maximum entropy priors' which is used as a penalty term (or prior probability) in a reconstruction. This definition encompasses the function described in Eq. (3.18), as well as functions such as $\mathbf{f}^{1/2}$ and $\ln(\mathbf{f})$ [129]. For clarification, the first definition is referred to as the ME method, while the second definition is called the ME priors.

ME method

The maximum entropy method can be interpreted as the method which gives a solution with the most disorder. In other words, it favours the solution which is most likely to have produced the observed measurements by assuming as little as possible. It is also the

solution which allows for the incorporation of prior information that does not conform to the traditional prior described by frequency distributions [89], which is by far the most commonly encountered in practice. However, only testable information can be incorporated mathematically into the problem formulation. This means that there is a procedure which can determine unambiguously whether the information agrees with the prior. For example, the statement 'the mean value of the parameter ϑ in a set of measurements is 2.56' is testable, while the statement 'the crowd was indifferent when the value of ϑ is 675' is not easily testable, even though it is information about the parameter we are trying to estimate.

The formal definition of the maximum entropy method can be derived from a combinatorial approach [4]. Consider the situation where an object is captured by a CCD camera; this is depicted mathematically by Eq. (3.5). It is assumed that there are M^2 elements in the object distribution, and that a total number of P units of energy are allocated among the M^2 elements of the object distribution:

$$\sum_{j=1}^{M^2} f_j = P. \tag{3.19}$$

As for the noise, let the values of n_j be a random variable about a positive mean, $n_j = N_j - B$, with B being a positive number, and N_j being another random variable. It is also assumed that the number of noise particles is finite, $\sum_{j=1}^{M^2} N_j = N_0$. Furthermore, it is assumed that the object and noise particles are distributed so that they are identical but distinguishable, and that any particle is equally likely to occupy any of the M^2 elements. The noise and the object are statistically independent. This situation is one of complete ignorance [4, 89].

In order to find the object and noise distributions which lead to the observed image values **d**, we can employ the concept that since the object and noise particles are equally likely to occur in any of the M^2 elements, the object and noise distribution is the one that can be formed from f_j and n_j in the maximum number of ways, provided that the image formation constraint stated in Eq. (3.5) is met.

It can be shown from a combinatorial consideration that the maximum number of ways in which an object can be formed from a fixed number P of distinguishable particles is:

$$W_f = \frac{P!}{\prod_{j=1}^{M^2} f_j!}, \tag{3.20}$$

and similarly, the number of ways the noise can be formed is:

$$W_n = \frac{N_0!}{\prod_{j=1}^{M^2} n_j!}.$$ (3.21)

From this, it can be deduced that the maximum number of ways the object and noise particles can be formed is:

$$\max_{\hat{\mathbf{f}}, \hat{\mathbf{n}}} W_f W_n$$ (3.22)

subject to the image formation process. As before, we can take the logarithm of Eq. (3.22) to give:

$$\ln W_f W_n = \ln P! + \ln N_0! - \sum_{j=1}^{M^2} \ln f_j! - \sum_{j=1}^{M^2} \ln n_j!.$$ (3.23)

The first two terms are constants and can be ignored in the maximisation. Applying Stirling's approximation to the logarithm of a factorial, where:

$$\ln k! \simeq k \ln k,$$ (3.24)

the optimisation problem can be formulated as follows.

$$\max_{\hat{\mathbf{f}}, \hat{\mathbf{n}}} \quad - \sum_{j=1}^{M^2} \hat{f}_j \ln \hat{f}_j - \sum_{j=1}^{M^2} \hat{n}_j \ln \hat{n}_j$$

$$\text{subject to} \quad d = \hat{f} \odot h + \hat{n}$$

$$\text{and} \quad \sum_{j=1}^{M^2} \hat{f}_j = P.$$ (3.25)

This can be solved either via a one-step process or recursively [4]. Although the above analysis is intuitively satisfying, it provides few clues as to how the ME method actually performs when applied to image processing problems. The following section gives a more practical approach to the problem.

ME priors

The ME method has been analysed formally by many authors. However, since it is a non-linear technique derived from many different disciplines of science, many prevailing views exist on this method. In order to examine the characteristics and quantify the properties of the ME priors, Nityananda and Narayan [129] took a pragmatic approach to the problem. They noted that there are many entropy priors that can be viewed under the same framework, and examined the performance of the priors in terms of resolution and noise sensitivity.

To understand the relationship between these entropies, or entropy priors, let us first define the quantity $f(x, y)$ such that its known Fourier coefficients are $\rho_{M,N}$:

$$\int \int f(x, y) \exp\{-i2\pi(Mx + Ny)\}dxdy = \rho_{M,N} \in K, \qquad (3.26)$$

where K is the set of known values. Then the aim of an ME restoration is to maximise a quantity which is the integral of a function $b(f(x, y))$:

$$E(f) = \int \int b(f(x, y))dxdy. \qquad (3.27)$$

The quality of the reconstruction using an ME prior is determined by two parameters: R_{ME}, which measures the compression of the baseline relative to the peaks (when compared to a bandlimited function), and n, which indicates how far this flattening effect extends upwards from the baseline. R_{ME} is defined in terms of the curvature of the function $f(x, y)$:

$$R_{ME} = \frac{b''(f_{min})}{b''(f_{max})}, \qquad (3.28)$$

while the value of n is in the order of 10^1. It should be noted that positivity is enforced if $n \geq 0$. A small n signifies a 'soft' prior, and produces a solution with a very flat baseline but relatively wide peaks, an example of which is the commonly used form of $-f \ln(f)$. As n increases, the prior gets 'harder'. The peaks in the reconstruction become sharper and the resolution increases, however, the baseline starts to have ripples in it. The entropy $\ln f$ is 'harder' than $-f \ln(f)$, while $-\exp(-f \ln(R_{ME}))$ is the hardest of all. Table 3.1 lists a family of entropy priors from the softest to the hardest prior. The baseline flattening and peak sharpening properties of the priors go through a gradation according to the order in Table 3.1.

It should be noted that although the two parameters R_{ME} and n describe the behaviour of the family of ME priors, they are only model parameters and are not guaranteed to predict the behaviour of the algorithm accurately at all times. In this case, numerical simulations are needed to analyse the performance of the algorithm.

3.3.4   Relationship between the three methods

There exists a very close relationship between the maximum likelihood estimator (MLE), the maximum *a posteriori* (MAP) estimator, and the maximum entropy (ME) method. The relationship is illustrated by the Venn diagram in Fig. 3.4.

ME priors
$\mathbf{f}^{3/2}$
$-\mathbf{f}\ln\mathbf{f}$
$\mathbf{f}^{1/2}$
$\ln\mathbf{f}$
$-1/\mathbf{f}^2$
$-\exp\{-\mathbf{f}\ln R_{ME}\}$

Table 3.1: *Different forms of the ME priors. The prior in the top row is the 'softest' prior which produces a solution with a flat baseline and wide peaks. The prior in the last row is the 'hardest' prior, meaning that it allows better resolution but introduces ripples in the baseline.*
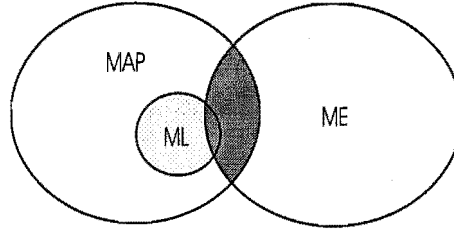
Figure 3.4: *Venn diagram illustrating the relationship between the maximum a posteriori, the maximum likelihood, and the maximum entropy methods.*

Since the MLE estimates the most likely function in the absence of prior information, it is a subset of the MAP estimator. In image processing, the usage of the MLE is most prevalent in situations described by Poisson statistics [96,156]; whereas the MAP estimator tends to be used when the scenario is governed by Gaussian statistics. There is also a 'grey-area' between the ML and MAP formulations. For instance, Schulz proposed a penalised maximum likelihood framework for astronomical blind deconvolution [154], which is essentially a MAP estimator.

The MAP estimator and the ME method are related under certain circumstances [171]. Consider the ME restoration defined by Andrews and Hunt as in Eq. (3.25). The solution to this problem is given by:

$$\hat{\mathbf{f}} = \exp\{-1 - 2\lambda\mathbf{H}^T(\mathbf{d} - \mathbf{H}\hat{\mathbf{f}})\} \tag{3.29}$$

where λ is the Lagrange multiplier. It should be pointed out that this solution also corre-

sponds to a constrained least squares solution.

To arrive at a MAP solution in the form of Eq. (3.29), the probability distributions $P(\mathbf{n})$ and $P(\mathbf{f})$ are required. Assuming that the noise is a zero-mean Gaussian distribution,

$$P(\mathbf{n}) = c_1 \exp\{-\frac{1}{2}(\mathbf{d} - \mathbf{Hf})^T \mathbf{K}_{nn}^{-1}(\mathbf{d} - \mathbf{Hf})\} \tag{3.30}$$

where c_1 is a constant and $\mathbf{K}_{nn}$ is the noise covariance matrix. In the case where the noise is stationary and uncorrelated, this matrix is an identity matrix $\sigma^2 \mathbf{I}$, where σ^2 is the variance of the noise process. With a judicious choice of $P(\mathbf{f})$, where

$$P(\mathbf{f}) = c_2 \exp\{-\beta \mathbf{f}^T \ln \mathbf{f}\}, \tag{3.31}$$

c_2 being another constant, it leads to the solution

$$\hat{\mathbf{f}} = \exp\{-1 + \frac{1}{\beta\sigma^2}\mathbf{H}^T(\mathbf{d} - \mathbf{H}\hat{\mathbf{f}})\} \tag{3.32}$$

which is equivalent to Eq. (3.29) except for the multiplicative constants. Therefore, although the methods of ME and MAP have been derived from different fundamental assumptions, they are actually very similar when certain probility distributions are considered.

Since the MLE is a subset of the MAP estimator, and the latter is equivalent to the ME method for certain situations, it is fair to say that the MLE is also related to the ME method. However, this is the extent of the equivalence. It should be pointed out that in Frieden's paper, where he stated that 'maximum likelihood implies maximum entropy' [47], used a different definition of maximum likelihood to that in the other literature, where maximum likelihood is used in conjunction with the likelihood function derived from a parametric description of an underlying probability function [176].

It is intuitive that certain classes of solutions would perform better than others under different conditions. An interesting example is the state of complete ignorance. The question is: what prior distribution describes the state of complete ignorance? To this question, Laplace replied, 'when the probability of a simple event is unknown, we may suppose all values between 0 and 1 as equally likely'. This point of view was shared by Bayes [89]. However, this does not describe many practical situations. Suppose a student finds a jar of crystals in a laboratory, and wishes to know whether they are soluble in water. If the student is able to dissolve one crystal, then it can be inferred that all the crystals are soluble. This means that the prior distribution should have greater weight at the end points $\vartheta = 0, 1$, which can be realised by using a maximum entropy framework.

3.4 Regularisation

At the beginning of the 20^{th} century, pure mathematicians such as Hadamard [68, 69] stated that physical situations only lead to well-posed problems, and that ill-posed problems are mathematical anomalies. It was not until the mid-60s that 'well-posed extensions', or regularisation methods were introduced to ill-posed problems. For example, in 1963, Tikhonov [169] proposed that the solution should have 'regular' properties, while Franklin [41] suggested the use of *a priori* statistics in the model in 1970.

According to Tikhonov [170], the regularisation method consists of finding regularisation operators that operate on the data, and determining the regularisation parameter from information pertaining to the problem. The regularisation operator must depend continuously on the data, and that as it approaches zero (or as the noise approaches zero), the regularised solution must approach the true solution. Many concepts have been proposed for regularising the ill-posed solution. These include the restriction of the data; the use of regularisation operators, and well-posed stochastic extensions of ill-posed problems.

3.4.1 Condition number

To determine whether a system requires regularisation (or how much regularisation is needed), one can examine the condition number of a matrix. The condition number of the matrix **H**, cond(**H**), is defined as [95]:

$$\text{cond}(\mathbf{H}) = \frac{\mu_1}{\mu_r} \tag{3.33}$$

where $\mu_1, \dots, \mu_n$ are the singular values of **H**, r is the rank of **H**, and that $\mu_1 \geq \mu_2 \geq \dots \geq \mu_r > \mu_{r+1} = \dots = \mu_n = 0$. If the value of cond(**H**) is small, then a small change in d does not produce a large change in **f**, i.e. the system is stable. As a result, only a small amount of regularisation is needed. However, if the condition number is large, then a small perturbation in the data produces a large perturbation in the solution. Such a matrix is called an ill-conditioned matrix, and the associated system is an ill-posed system. In this case, it is crucial that regularisation is included in the solution.

3.4.2   Forms of regularisation

There are many well-established regularisation methods, and a review of these methods is given below. Consider an imaging system

$$d_k(x,y) = f(x,y) \odot h_k(x,y) + n_k(x,y) \tag{3.34}$$

where $f(x,y)$ is the unknown, and k denotes the k^{th} realisation of the respective quantities. The unregularised solution to the problem is obtained by solving for the maximum likelihood solution with the error metric:

$$E(\hat{f}(x,y)) = \sum_k \sum_{x,y} (\hat{f}(x,y) \odot h_k(x,y) - d_k(x,y))^2. \tag{3.35}$$

We call the term to be minimised $E(\hat{f}(x,y))$ the error function, or the objective function. It should be noted that although the ML solution associated with Gaussian statistics is used here, the discussion on regularisation can be generalised to different noise statistics. When regularisation is included in the solution, the error function is modified to:

$$E(\hat{f}(x,y)) = \sum_k \sum_{x,y} (\hat{f}(x,y) \odot h_k(x,y) - d_k(x,y))^2 + \lambda \times \text{regularisation term} \tag{3.36}$$

where λ is the Lagrange multiplier and determines the amount of regularisation in the error function. This type of regularisation, where an extra term is added to the metric, is called explicit regularisation. Implicit regularisation is also possible, but here the regularisation is not obvious from the error metric. One example is the termination of iterations, where an iterative algorithm is terminated before complete convergence to prevent noise amplification [104]. However, the point at which the algorithm should be terminated is not well-defined, and user input may be required. Moreover, it is difficult to analyse its performance analytically. As a result, it is preferable to use explicit regularisation in an algorithm. Several types of explicit regularisation are discussed below.

Energy and smoothness constraints

Both of these constraints modify the error function to the form:

$$E(\hat{f}(x,y)) = \sum_k \sum_{x,y} (\hat{f}(x,y) \odot h_k(x,y) - d_k(x,y))^2 + \lambda \sum_{x,y} (\hat{f}(x,y) \odot c(x,y))^2 \tag{3.37}$$

where $c(x,y)$ is the regularisation kernel. With the energy constraint, also known as the Tikhonov-Miller regularisation, the kernel is simply the identity matrix. As for the smooth-

ness constraint (or the Laplacian), the kernel is essentially a high-pass filter:

$$c(x,y) = \begin{pmatrix} 0 & -0.25 & 0 \\ -0.25 & 1 & -0.25 \\ 0 & -0.25 & 0 \end{pmatrix}.$$

This kernel ensures continuity in the reconstruction and is more suited to truncated images compared to the energy constraint.

It is important to achieve a balance between the consistency error and the regularisation term (first and second terms in Eq. (3.37) respectively), as discussed in Section 3.3.2.

Positivity and support

In incoherent imaging (defined in Chapter 4), it is the intensity values that are observed. Since the intensity values must be positive, it is reasonable to enforce the reconstructed values to be positive. When the object is of finite extent, it is also possible to define a support that it must be enclosed in. There are two ways of implementing these two constraints. Firstly, they can be treated as a regularisation term [105]:

$$E = \lambda \sum_{\Omega_f} |f(x,y)|^2 \tag{3.38}$$

where Ω_f denotes the points where $f(x,y)$ violates the positivity or support constraints, i.e. the points where f is negative or outside the support. This is essentially a penalty method, which is also called a method of 'loose' constraints, and is a linear method. An alternative approach is to use a projection based method [183], where the prior information is expressed through constraint sets. Basically, this method involves setting the negative pixels or pixels outside the support to zero after each iteration and is a non-linear method. Positivity can also be enforced by parameterisation where $f(x,y) = \psi^2(x,y)$ [157, 168]. This however distorts the error surface and a unique solution cannot be guaranteed. This problem is overcome by the quadratic programming method [94] which preserves the unimodality of a parabolic error surface. This is the method adopted for enforcing positivity in this thesis.

Total variation

In images which contain both smooth areas and sharp edges, it is desirable to use a prior that can preserve both characteristics. In this case, the total variation prior can be used [21,185]. This involves the minimisation of:

$$E = \lambda \sum_{x,y} |\nabla f(x,y)| \tag{3.39}$$

where ∇ is the gradient operator defined as:

$$\nabla f(x,y) = \vec{x}\frac{\partial f(x,y)}{\partial x} + \vec{y}\frac{\partial f(x,y)}{\partial y}$$ (3.40)

and $\vec{\cdot}$ is a direction vector.

Entropy constraint

As discussed in Section 3.3.3, a commonly used form of the maximum entropy prior is $-\mathbf{f}\ln(\mathbf{f})$.

Statistical priors

If the statistics of the unknown object are available, it can be incorporated into the problem formulation via a covariance matrix. An example is the estimation of the atmospheric point spread function, which is characterised by Kolmogorov statistics [147]. If the unknown that we are solving for is $\mathbf{f}$, and its associated covariance matrix is $\mathbf{K}_{ff}$, then the regularisation term is:

$$\log[P(\hat{\mathbf{f}})] = \lambda\frac{1}{2}\hat{\mathbf{f}}^T\mathbf{K}_{ff}^{-1}\hat{\mathbf{f}}.$$ (3.41)

The application of Kolmogorov statistics to astronomical blind deconvolution has been investigated, the results of which are included in Chapter 6.

Multiple frames

The use of multiple frames adopts a different approach to that specified by Eq. (3.36). Instead of adding an extra term to the objective function, using multiple frames provides more data information, and reduces the number of unknowns relative to the known quantities. In the case of single-frame blind deconvolution, the symmetry of Eq. (3.34) means that the estimates of $\hat{f}(x,y)$ and $\hat{h}(x,y)$ can be reversed. There is also a possibilty of obtaining the trivial solution where the object is estimated to be a delta function, and the PSF the data (or vice versa). When multiple frames are used, the problem of the reversal of the estimates is eliminated, and the possibility of the trivial solution is reduced.

Others

This section has given an overview of the regularisation operators that can be used. Of course, the operators are not limited to the list mentioned here. In fact, in a reconstruction, the operators used need not conform to any standard structure; as long as the assumption is reasonable, it is a valid form of operator.

3.5 Summary

This chapter examined the topic of inverse problems. An inverse problem is the deduction of the underlying real function from measured data through a model of the physical process which produced the measurements. Inverse problems occur frequently in reality, and many solutions have been proposed to solve these problems, a review of which has been given.

In solving an inverse problem, only a finite amount of information can be measured to essentially describe a system which has an infinite number of variables. The presence of noise is also inherent to any measurement. As a result, many inverse problems are ill-conditioned, and regularisation must be employed to achieve a meaningful solution to the problem. Common forms of regularisation have been presented in this chapter.

Examples of inverse problems include interpolation and almost all problems in astronomical imaging. Some of these problems are addressed in the subsequent chapters in this thesis.

Chapter 4

Imaging through turbulence

When an object is imaged, its wavefront is degraded by many elements before it is finally recorded. These degradations include the lenses in the optical system, the CCD array, the imaging medium and the noise. All these non-ideal components affect the wavefront differently. A good understanding of these processes facilitates solving the associated inverse problems.

One of the fundamental theories in imaging is that of light propagation, which is divided into geometrical optics and Fourier optics. Geometrical optics employs simple geometrical relationships to predict the propagation of light. Fourier optics provides a more accurate model of light propagation which also accounts for the diffraction effects introduced by an imaging system. Although diffraction effects are not included in the theory of geometrical optics, corrections can be made to the approximation after the process of light propagation. These topics are discussed in Section 4.1.

The phenomenon of diffraction poses as an upper limit to an imaging system. However, diffraction-limited imaging is only realisable if appropriate sampling is employed in the focal plane. The relationship between diffraction, sampling, and the resolution of an optical system is examined in Section 4.2. Although the resolution of the captured images is limited by diffraction, a class of post-processing methods known as superresolution algorithms can be applied to the images to exceed this limit, a brief overview of which is also given in this section.

In a real system, the upper limit of diffraction-limited imaging is often not reached due to the aberrations on the incoming wavefront. Further distortions are introduced on the image as a result. The quantification of the performance of an aberrated imaging system, as well as other parameters which characterise such systems are discussed in Section 4.3.

When capturing astronomical images from ground, it is inevitable that the wavefront propagates through the Earth's turbulent atmosphere. The characterisation of such a process is essential for solving the inverse problem of reconstructing the star's intensity distribution from the captured images. The process of imaging through atmospheric turbulence is presented in Section 4.4.

It is important to be able to perform simulations to verify the theory and the performance of the algorithms. However the ultimate objective is to apply the algorithms on real world data. Section 4.5 deals with the practical issues associated with handling real astronomical data. In Section 4.6, a review of the existing algorithms on solving astronomical inverse problems is given.

4.1 Optics theory

Optics deals with the study of light and is of utmost importance in the understanding of wavefront propagation in a system. Optics theory is encountered in many systems, such as telescopes and microscopes, and there exist numerous branches of optics: geometrical optics, Fourier optics and quantum optics [32]. In this section, the theory of geometrical and Fourier optics which is relevant to this thesis is discussed.

It is however necessary to firstly introduce the concept of light rays and a wavefront. Fig. 4.1 shows a point source from which light rays diverge. A wavefront is the surface perpendicular to the direction of light propagation, and is spherical in this illustration.
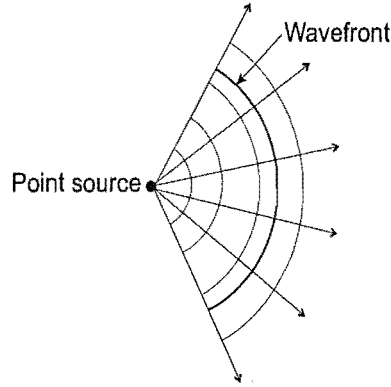
Figure 4.1: *Illustration of an object: a point from which rays of light diverge.*

4.1.1 Geometrical optics

Fundamental laws of geometrical optics

Geometrical optics is the study of the propagation of light using simple geometrical relationships [32]. There are four basic laws of geometrical optics. The first law is the law of rectilinear propagation, meaning that in an isotropic, homogeneous medium, light travels in straight lines. When light encounters a surface between two media, reflection, or refraction, or both phenomena occur. These phenomena are the second and the third laws respectively. Reflection is illustrated in Fig. 4.2, where a light ray is reflected off a mirror surface. The size of the angle of arrival θ_1 is the same as that of the angle of departure θ_2. An example of refraction is shown in Fig. 4.3 where the light ray travels from medium 1 to medium 2 with respective refractive indices of n_1 and n_2. The refractive index describes the speed of light in a particular medium relative to its speed in vacuum c which is equal to $3 \times 10^8 \mathrm{ms}^{-1}$. The angles θ_1 and θ_2 are related to the refractive indices by Snell's Law:

$$\frac{n_1}{n_2} = \frac{\sin \theta_2}{\sin \theta_1}. \tag{4.1}$$

For instance, atmospheric turbulence can be viewed as a collection of lenslets with different refractive indices. The fourth law of geometric optics is the law of reversibility. This law states that if the direction of a light ray is reversed, it traces exactly the same path backwards.
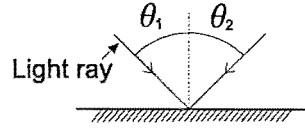
Figure 4.2: *An illustration of reflection.*
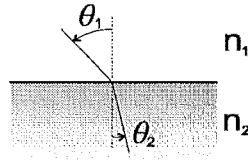
Figure 4.3: *An illustration of refraction.*

Lenses and image formation

The most important components of optical imaging and optical data processing are lenses. Lenses come in many shapes, the most common of which is the converging, or double-convex lens. The discussion that follows is presented from the point of view of using such a lens, but this does not restrict the application of the theory to other lenses.

As shown in Fig. 4.1, the point source is a point from which rays of light diverge. To form an image of such an object, a lens must be used so that the spherical wavefront of the object converges on a point. This is illustrated in Fig. 4.4 [91]. Two approximations are employed here: the thin lens approximation, which states that there is negligible translation of the light ray within the lens in the direction other than that of light propagation, and the paraxial approximation, which states that only portions of the wavefront that lie near the lens axis are considered. These approximations are valid for a variety of circumstances. For instance, a telescope can be modelled as a thin lens. The properties of a lens can be combined into a single parameter called the focal length f_l where:

$$\frac{1}{f_l} = (n-1)\left(\frac{1}{R_1} - \frac{1}{R_2}\right) \tag{4.2}$$

and R_1 and R_2 are the curvature of the lens surface on the left-hand and right-hand sides respectively. n is the refractive index of the lens. The distance d_i behind the lens at which
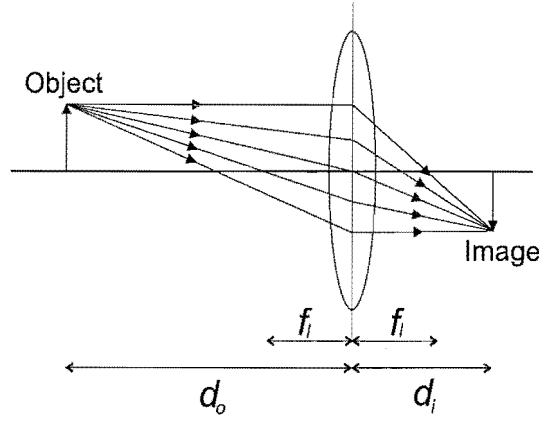
Figure 4.4: *Image formation using a thin lens.*

the image forms is defined by the thin lens equation as

$$\frac{1}{d_o} + \frac{1}{d_i} - \frac{1}{f_l} = 0. \tag{4.3}$$

In the case where d_o is infinite, for example, when imaging a star, $d_i = f_l$. A useful property of a converging lens is that there exists a Fourier transform relationship between the wavefront at the lens and the image at the focal plane [61]. This is useful in analysing image formation in an optical system.

In astronomical imaging, the major distortion on a wavefront from a star is the phase distortion introduced by the Earth's atmosphere. One way of correcting for this distortion is to use an adaptive optics system. If a single turbulent layer is situated at distance d_0 above ground as in Fig. 4.4, then a deformable mirror (DM) can be placed at a distance d_i behind the lens, called the conjugate plane, to cancel the effect of the phase distortion. The corrected wavefront is then reflected off the mirror to be captured by an imaging device. This concept is not only limited to astronomical imaging, but can also be used for applications such as retinal imaging [112, 122].

General imaging system

A general imaging system is shown in Fig. 4.5. The point at which the wavefront of the object enters the system of imaging elements is called the entrance pupil, while the point at which the wavefront leaves the imaging elements is called the exit pupil. Obviously,
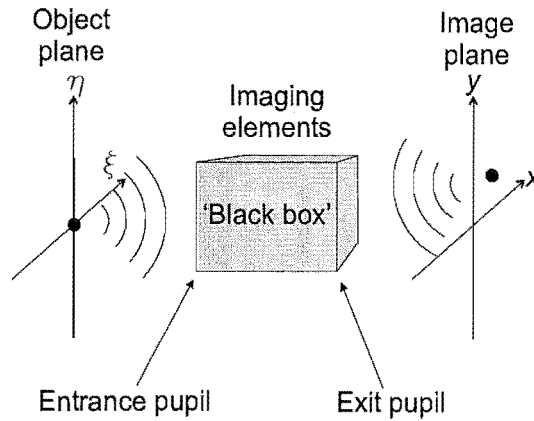
Figure 4.5: *A general imaging system.*

the resolution of the image is dependent on the size of the smaller of the two pupils (or apertures).

The system of imaging elements depicted as a black box in Fig. 4.5 may be only a lens, or it may comprise several components. In astronomical imaging, a typical configuration of the imaging elements includes an aperture with a central obscuration, a high-speed shutter and converging lenses. If both speckle images and wavefront sensing images are captured, a beam splitter and more lenses are also used.

4.1.2 Fourier optics

Diffraction has been defined by Sommerfeld [159] as 'any deviation of light rays from recti-linear paths which cannot be interpreted as reflection or refraction'. Indeed, there are many situations where the theory of geometrical optics is insufficient in describing the properties of light, and it is due to diffraction effects which are inherent to all optical systems. In this case, Fourier optics must be used to describe essentially the wave nature of light. Nevertheless, it should be pointed out that because the theory of geometrical optics is a sufficiently close approximation to Fourier optics that the diffraction effects can be taken into account after image formation via a convolution operation with a point spread function determined by diffraction from the aperture.

Diffraction

Any monochromatic optical wave $u(P, t)$ at position P and time t can be expressed as

$$u(P, t) = \Re\{U(P)\exp\{-i2\pi\nu t\}\} \tag{4.4}$$

where ν is the optical frequency and $\Re\{\cdot\}$ denotes the real part. $u(P, t)$ must satisfy the scalar wave equation

$$\nabla^2 u - \frac{1}{c^2}\frac{\partial^2 u}{\partial t^2} = 0 \tag{4.5}$$

where c is the speed of light in a vacuum, and ∇^2 is the Laplacian operator

$$\nabla^2 = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} + \frac{\partial^2}{\partial z^2}. \tag{4.6}$$

If the time dependence is removed, and $u(P)$ is expressed as a complex function $U(P) = |U(P)|\exp\{-i\vartheta\}$, then $U(P)$ must also obey the Helmholtz equation [61] $(\nabla^2 + k^2)U = 0$ where k is the wave number defined as

$$k = \frac{2\pi}{\lambda} \tag{4.7}$$

and λ is the wavelength.

The diffraction problem is illustrated in Fig. 4.6. Consider a finite aperture in an infinite opaque screen, which is taken to be planar and adopts a Cartesian coordinate system (x_1, y_1). The observation plane is a distance z away from and parallel to the aperture plane, with the coordinate system (x_0, y_0). Using the Helmholtz equation and applying Green's theorem with a judicious choice of Green's function [163] results in the Rayleigh-Sommerfeld formula which describes the wavefront at the observation plane [61]:

$$U(x_0, y_0) = \frac{1}{i\lambda}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} U(x_1, y_1)\frac{\exp\{ikr_{01}\}}{r_{01}}\cos(\theta)dx_1 dy_1. \tag{4.8}$$

θ is the angle formed by the direction of r_{01} with the z-axis. Eq. (4.8) can also be written as

$$U(x_0, y_0) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} h(x_0, y_0; x_1, y_1)U(x_1, y_1)dx_1 dy_1 \tag{4.9}$$

where

$$h(x_0, y_0; x_1, y_1) = \frac{1}{i\lambda}\frac{\exp\{ikr_{01}\}}{r_{01}}\cos(\theta). \tag{4.10}$$

It should be noted that Eq. (4.9) is the Huygens-Fresnel principle. Physically, it means that the optical field at any point in the observation plane can be expressed as the superposition of many spherical wavefronts originating from the aperture plane.
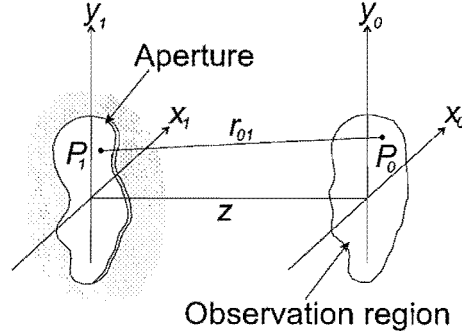
Figure 4.6: *Diffraction geometry.*

Fresnel diffraction

The Huygens-Fresnel principle can be simplified under certain circumstances for easier mathematical manipulation. These approximations, which are usually made when dealing with propagating wavefronts, are called the Fresnel and Fraunhofer approximations.

It is firstly assumed that the observation distance z between the aperture and observation plane is much greater than the maximum linear dimension of the aperture. Moreover, only a finite region about the z-axis in the observation plane is of interest. These assumptions mean that the angle θ is negligibly small such that $\cos(\theta) \cong 1$. The actual distance r_{01} between the points can also be approximated by z. However the r_{01} term in the exponential cannot be approximated by z. This is because the small difference between r_{01} and z can introduce large errors when multiplied by k which is large, and can generate phase errors that are much greater than 2π radians.

However, the term r_{01} in the exponential can be simplified by the Fresnel approximations. The geometric distance between the points (x_0, y_0) and (x_1, y_1) is

$$
\begin{aligned}
r_{01} &= \sqrt{z^2 + (x_0 - x_1)^2 + (y_0 - y_1)^2} \\
&= z\sqrt{1 + \left(\frac{x_0 - x_1}{z}\right)^2 + \left(\frac{y_0 - y_1}{z}\right)^2}.
\end{aligned}
\tag{4.11}
$$

Using the binomial expansion where

$$
\sqrt{1 + b} = 1 + \frac{1}{2}b - \frac{1}{8}b^2 + \dots \qquad |b| < 1
\tag{4.12}
$$

and approximating Eq. (4.11) with the first two terms results in

$$r_{01} \cong z\left[1 + \frac{1}{2}\left(\frac{x_0 - x_1}{z}\right)^2 + \frac{1}{2}\left(\frac{y_0 - y_1}{z}\right)^2\right]. \tag{4.13}$$

This assumption allows h to be written as

$$h(x_0, y_0; x_1, y_1) = \frac{\exp\{ikz\}}{i\lambda z}\exp\left\{\frac{ik}{2z}[(x_0 - x_1)^2 + (y_0 - y_1)^2]\right\}. \tag{4.14}$$

When the distance z is sufficiently large for the assumption to be accurate, the observer is said to be in the region of Fresnel diffraction, where

$$z^3 \gg \frac{\pi}{4\lambda}[(x_0 - x_1)^2 + (y_0 - y_1)^2]^2_{\max}. \tag{4.15}$$

Essentially, the Fresnel approximation has replaced the spherical wavefronts with quadratic wavefronts. This is intuitively reasonable because the region around the turning point of a parabola closely approximates a sphere. Obviously, the validity of the approximation depends on the relative sizes of the apertures, the observation region, and the distance z.

It is interesting to note that the superposition integral can be written in two alternative forms. Firstly, $U(x_0, y_0)$ can be viewed as a convolution between $U(x_1, y_1)$ and the weighting h, giving the expression:

$$\begin{aligned} U(x_0, y_0) &= \frac{\exp\{ikz\}}{i\lambda z} \\ &\times \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} U(x_1, y_1)\exp\left\{i\frac{k}{2z}[(x_0 - x_1)^2 + (y_0 - y_1)^2]\right\}dx_1dy_1. \end{aligned} \tag{4.16}$$

Alternatively, the quadratic terms in the exponent can be expanded to give:

$$\begin{aligned} U(x_0, y_0) &= \frac{\exp\{ikz\}}{i\lambda z}\exp\left\{i\frac{k}{2z}(x_0^2 + y_0^2)\right\} \\ &\times \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} U(x_1, y_1)\exp\left\{i\frac{k}{2z}(x_1^2 + y_1^2)\right\}\exp\left\{-i\frac{2\pi}{\lambda z}(x_0x_1 + y_0y_1)\right\}dx_1dy_1. \end{aligned} \tag{4.17}$$

This means that the function $U(x_0, y_0)$ can be found from a Fourier transform of

$$U(x_1, y_1)\exp\left\{i\frac{k}{2z}(x_1^2 + y_1^2)\right\}. \tag{4.18}$$

Fraunhofer diffraction

The Fraunhofer, or far field approximation is made when the spherical waves have travelled so far that they have become almost planar and can be approximated by plane waves. The observer is said to be in the Fraunhofer region if the condition

$$z \gg \frac{k(x_1^2 + y_1^2)_{max}}{2} \tag{4.19}$$

is met. Effectively, the quadratic phase factor in the integral in Eq. (4.17) is approximately unity over the entire aperture, and the field at $U(x_0, y_0)$ is a Fourier transform of the aperture distribution itself:

$$U(x_0, y_0) = \frac{\exp\{ikz\}}{i\lambda z} \exp\left\{i\frac{k}{2z}(x_0^2 + y_0^2)\right\}$$
$$\times \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} U(x_1, y_1) \exp\left\{-i\frac{2\pi}{\lambda z}(x_0 x_1 + y_0 y_1)\right\} dx_1 dy_1. \tag{4.20}$$

With a change of variables of:

$$\begin{aligned}
x &= x_0 \\
y &= y_0 \\
u &= \frac{x_1}{\lambda z} \\
v &= \frac{y_1}{\lambda z},
\end{aligned} \tag{4.21}$$

Eq. (4.20) can be written as

$$U(x, y) = c_1 \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} U(\lambda z u, \lambda z v) \exp\{-i2\pi(xu + vy)\} du dv, \tag{4.22}$$

with c_1 being the exponential scalars. This clearly illustrates the Fourier transform relationship between the wavefronts at the aperture and image planes.

When imaging stellar objects, the Fraunhofer approximation can be employed, as the wavefronts from the stars have travelled lightyears to reach the surface of the Earth and can be considered planar.

Angular spectrum

A slightly different approach to diffraction theory is to Fourier analyse a complex field and examine what is called the angular spectrum. Through this analysis, it is found that the

field amplitude at any point is a sum of contributing plane waves, including the phase shifts introduced by propagation.

Consider a complex wavefront in a Cartesian coordinate system (x, y, z). The wavefront is at the plane $z = 0$ and is denoted by $U(x, y, 0)$. Its two-dimensional Fourier transform is given by:

$$A_0(f_X, f_Y) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} U(x, y, 0) \exp\{-i2\pi(f_X x + f_Y y)\} dx dy. \qquad (4.23)$$

For a unit-amplitude plane wave propagating with direction cosines (α, β, γ), its wavefront is denoted by:

$$B(x, y, z) = \exp\left\{ i\frac{2\pi}{\lambda}(\alpha x + \beta y + \gamma z) \right\} \qquad (4.24)$$

where $\gamma = \sqrt{1 - \alpha^2 - \beta^2}$. Therefore at $z = 0$, the complex exponential in Eq. (4.23) can be viewed as a plane wave propagating with direction cosines:

$$\begin{aligned}
\alpha &= \lambda f_X \\
\beta &= \lambda f_Y \\
\gamma &= \sqrt{1 - (\lambda f_X)^2 - (\lambda f_Y)^2}.
\end{aligned} \qquad (4.25)$$

Because the complex amplitude of the plane wave component is simply $A_0(f_X, f_Y) df_X df_Y$ evaluated at $(f_X = \alpha/\lambda, f_Y = \beta/\lambda)$, the quantity:

$$A_0\left(\frac{\alpha}{\lambda}, \frac{\beta}{\lambda}\right) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} U(x, y, 0) \exp\left\{ -i2\pi\left(\frac{\alpha}{\lambda} x + \frac{\beta}{\lambda} y\right) \right\} dx dy \qquad (4.26)$$

is called the angular spectrum of the wavefront $U(x, y, 0)$.

If the wavefront is propagated for a distance z, the angular spectrum is modified to:

$$A_z\left(\frac{\alpha}{\lambda}, \frac{\beta}{\lambda}, z\right) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} U(x, y, z) \exp\left\{ -i2\pi\left(\frac{\alpha}{\lambda} x + \frac{\beta}{\lambda} y\right) \right\} dx dy. \qquad (4.27)$$

The relationship between the angular spectra before and after propagation is [61]:

$$A_z\left(\frac{\alpha}{\lambda}, \frac{\beta}{\lambda}, z\right) = A_0\left(\frac{\alpha}{\lambda}, \frac{\beta}{\lambda}\right) \exp\left\{ i\frac{2\pi}{\lambda}\sqrt{1 - \alpha^2 - \beta^2} z \right\}. \qquad (4.28)$$

For the wavefront propagation considered in this thesis, $\alpha^2 + \beta^2 < 1$, meaning that the propagation over a distance z has simply introduced a relative phase change in different components of the angular spectrum.

If the aperture has a transmittance function $t(x, y)$ on the $z = 0$ plane, then the resultant angular spectrum is the convolution of the incident angular spectrum with the Fourier transform of the transmittance function $T(\alpha/\lambda, \beta/\lambda)$. This is denoted by:

$$A_t\left(\frac{\alpha}{\lambda}, \frac{\beta}{\lambda}\right) = A_i\left(\frac{\alpha}{\lambda}, \frac{\beta}{\lambda}\right) \odot T\left(\frac{\alpha}{\lambda}, \frac{\beta}{\lambda}\right). \tag{4.29}$$

From a practical viewpoint, the Fourier transform of the aperture function $T(\alpha/\lambda, \beta/\lambda)$ is analogous to a point spread function, which has the effect of broadening the angular spectrum of the incident wavefront. The smaller the aperture is, the more spreading it introduces to the angular spectrum. This is a direct manifestation of the scaling property of the Fourier transform as discussed in Chapter 2.

One application of the angular spectrum is the propagation of the wavefront from the back of the lenslet array in a Shack-Hartmann wavefront sensor to the detector plane. This is used in the forward modelling of the wavefront sensor for a myopic deconvolution scheme [146].

Diffraction at an aperture

Simple geometrical optics predicts that an image of a point object is also a point. When diffraction effects are taken into account, it is found that the image is spread, resulting in a Fraunhofer diffraction-limited pattern. The pattern depends on the shape of the aperture, and here the two most common apertures: the square and the circular apertures, are investigated.

To facilitate the discussion, let us first introduce the generalised pupil function $A(u, v)$ defined mathematically as:

$$A(u, v) = \begin{cases} P(u, v) \exp\{i\phi(u, v)\} & \text{inside aperture} \\ 0 & \text{outside aperture} \end{cases},$$

where $P(u, v)$ is the pupil function:

$$P(u, v) = \begin{cases} 1 & \text{inside aperture} \\ 0 & \text{outside aperture} \end{cases}.$$

The pupil function describes a diffraction-limited pupil for imaging an object, while the generalised pupil function accounts for any wavefront aberrations in the exit pupil. The

pupil phase $\phi(u, v)$ is a function of the overall wavefront aberration Φ defined in Section 4.3.3. For notational convenience, the coordinates $(\lambda z u, \lambda z v)$ are shortened to (u, v) when there is no confusion.

Recall the Fourier transform relationship between the aperture and the observation plane in Eq. (4.22). The simplest form of $U(u, v)$ is the pupil function $P(u, v)$. Applying the Fourier transform relationship to this results in the system image response $\tilde{h}(x, y)$:

$$\tilde{h}(x, y) = \mathfrak{F}\{P(u, v)\}. \tag{4.30}$$

In an aberrated system, the pupil function is replaced by the generalised pupil function. It should be clarified that $\tilde{h}(x, y)$ is the wavefront on the image plane, and $h(x, y) = |\tilde{h}(x, y)|^2$ is an intensity distribution which defines the point spread function (PSF). The Fourier transform of the PSF is called the optical transfer function (OTF).

We are now in a position to examine the diffraction patterns at the aperture. Consider first a rectangular aperture described by:

$$P(u, v) = \text{rect}(u)\text{rect}(v). \tag{4.31}$$

Recalling that the aperture plane and the image plane have a direct Fourier transform relationship, the pattern observed at the focal plane is:

$$\tilde{h}(x, y) = \text{sinc}(x)\text{sinc}(y). \tag{4.32}$$

The intensity pattern observed is

$$|\tilde{h}(x, y)|^2 = \text{sinc}^2(x)\text{sinc}^2(y) \tag{4.33}$$

and is shown in Fig. 4.7(a). This pattern decays at a rate proportional to the reciprocal of the square of the distance from the origin. A situation where this pattern can be observed is when a Shack-Hartmann (SH) sensor is used to image a point source at or above the Nyquist rate. Each lenslet in the SH sensor is square due to the need to configure them into an array. In order to Nyquist sample this intensity pattern, two samples are required across the main lobe of the spot [17].

Now consider the diffracting aperture to be circular with radius a. Then the pupil function in the absence of aberrations is:

$$P(u, v) = \begin{cases} 1 & u^2 + v^2 \leq a^2 \\ 0 & \text{otherwise.} \end{cases}$$
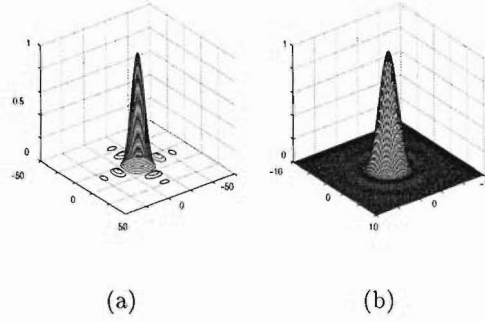
Figure 4.7: (a) Fraunhofer diffraction pattern of a square aperture. (b) Fraunhofer diffraction pattern of a circular aperture.

Due to its circular symmetry, it can be expressed as:

$$P(\rho) = \begin{cases} 1 & \rho \le a \\ 0 & \text{otherwise} \end{cases}$$

which is also equal to the circle function $\text{circ}(\rho/a)$. Instead of taking a Fourier transform, a Hankel transform can be used and this results in the diffraction pattern described by:

$$\tilde{h}(r) = \frac{J_1(2a\pi r)}{ar} \tag{4.34}$$

where J_1 is the first order Bessel function of the first kind. The intensity pattern observed at the focal plane is the square of the above expression, denoted by:

$$|\tilde{h}(r)|^2 = \left(\frac{J_1(2a\pi r)}{ar} \right)^2. \tag{4.35}$$

This pattern is called the Airy disk and is shown in Fig. 4.7(b). The main difference between the Airy disk and the diffraction pattern of a square aperture is that the zero-crossings of the former are not equally spaced: the further away from the origin, the closer the zero-crossings are (Table 4.1). The Airy disk is useful for characterising the resolution limit of a system and is discussed in the next section.

4.2   Resolution, sampling and superresolution

The resolution of an image captured at the image plane is limited by the diffraction effects introduced at the aperture. However, this upper limit is only reached if sufficient sampling

r	$\left(\frac{J_1(2a\pi r)}{ar}\right)^2$	max or min
0	1	max
3.833	0	min
5.136	0.0175	max
7.016	0	min
8.417	0.0042	max
10.174	0	min

Table 4.1: *The first few maxima and minima of the Airy disk.*

is employed. The two topics of resolution and sampling are therefore closely linked, but are seldom addressed together. This section examines the relationship between them.

Although the resolution of the captured image is limited by diffraction effects, post-processing methods called superresolution methods can be applied to exceed this limit in the reconstruction. The feasibility of superresolution is examined, and an overview of the existing methods is presented.

4.2.1 Resolution and sampling

A common criterion for resolution comparison is the ability to resolve two closely spaced point sources. One of the most established criteria is Rayleigh's resolution criterion, where two point sources of equal strength are 'barely resolved' if the centre of the Airy disk of one point source falls on the first zero of the Airy disk of the other point source [61]. Although this criterion is somewhat arbitrary, for historical purposes, it is still used as a benchmark for resolution comparison.

Fig 4.8 illustrates Rayleigh's resolution criterion for a circular aperture, where the peak-to-saddle ratio is 0.735 [13]. It should be noted that the peak-to-saddle ratio changes for different apertures. The angular resolution of an image according to Rayleigh's criterion is [61]:

$$\alpha_{min} = 1.22\frac{\lambda}{D} \tag{4.36}$$

where D is the diameter of the aperture and λ is the wavelength. This corresponds to a
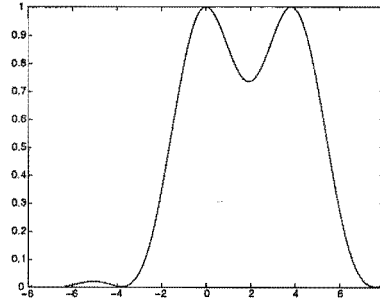
Figure 4.8: *Intensity pattern of two point sources imaged through a circular aperture which just meet Rayleigh's resolution criterion.*

(circular) aperture with no central obscuration. For an aperture with a central obscuration (which is the case for most telescopes), also known as an annular aperture, the angular resolution is in fact improved [13,161]. This is reasonable considering that the obscuration essentially acts as a highpass filter. For instance, for an obscuration with a diameter equal to half of the telescope diameter, the resolving power is increased to λ/D.

In astronomical imaging, angular resolution is defined as simply λ/D. However, whether this arose as a convenient approximation to $1.22\lambda/D$, or to account for the effect of the central obscuration is unclear. Assuming the definition of λ/D, for a 1-metre telescope imaging at 500nm, an angular resolution of

$$
\begin{aligned}
\alpha_{\min} &= \frac{\lambda}{D} \\
&= 5 \times 10^{-7} \text{ rad} \\
&= 2.86 \times 10^{-5 \circ} \\
&= 0.1 \text{ arcsec}
\end{aligned}
\tag{4.37}
$$

is achievable in the absence of atmospheric turbulence. Nevertheless, this resolution is only easily attainable if the Nyquist criterion is also satisfied (Subsection 2.6.1). The Nyquist sampling rate is essentially the arc length (since the angle is very small) which subtends the angle $\lambda/2D$, and should match the pixel size of the camera. A realistic pixel size is approximately $10\mu\text{m} \times 10\mu\text{m}$. Finally the focal length of the lens can be determined, which is simply the ratio of the pixel size to the angular resolution $\lambda/2D$ in radians for a Nyquist sampled system.

4.2.2 Superresolution

Superresolution is commonly referred to as the phenomenon of extending the object spectrum beyond the optical cutoff frequency. Whether superresolution is achievable was a topic of discussion for a long time, and its feasibility was only accepted recently. This can be proved by analytic continuation [61], the essence of which is as follows. Firstly, the two-dimensional Fourier transform of a spatially-bounded object is an analytic function in the u-v plane. Secondly, if any analytic function in the u-v plane is known exactly in an arbitrarily small but finite region of that plane, then the entire function can be found uniquely by means of analytic continuation.

There exists much literature on superresolution [44,54,85,86,114,118,119,139,155–157,179]. One of the earliest superresolution algorithms was proposed by Gerchberg [54]. Basically, the object is projected between the spatial and the frequency domains, and constraints derived from prior knowledge are imposed in each domain. Although this seems *ad hoc*, the algorithm is shown to be convergent. Walsh and Delaney [179] modified this algorithm and proposed a non-iterative solution to superresolution by solving a set of linear equations.

A superresolution algorithm for Poisson statistics was proposed independently by Richardson [139], Lucy [114], and Shepp and Vardi [156]. Richardson and Lucy took a Bayesian approach to the problem, while Shepp and Vardi derived the algorithm via an expectation-maximisation (or maximum likelihood) framework.

Hunt and Sementilli [86] extended these algorithms to the Poisson MAP (PMAP) method which used the idea of supersampling, allowing the object and the PSF to be on a finer grid than the data. This explicit implementation of superresolution ensures that there is 'room' for the higher frequency content. Gerwe *et al.* [55] also employed supersampling to process adaptive optics compensated imagery.

Hunt [85] provided a different perspective on this topic, stating that superresolution and interpolation are linked in a simple manner: if the frequency content of the data increases, the sampling grid must be more finely spaced. As a result, algorithms which use interlacing or image registration techniques can also be viewed in the framework of superresolution [2,22,57,73–75,109].

Fried [44] also viewed superresolution differently from most of the existing literature. He defined it as the recovery of spatial frequency components where the OTF is zero, regardless of the components being above or beneath the optical cutoff of the instrument. Although for astronomical imaging, the OTF is usually non-zero everywhere, this is not so for non-regular baseline interferometry. In Fried's paper, this phenomenon was demonstrated on the CLEAN algorithm [79], with the key to the success of superresolution being the small extent of the object. It is interesting to point out that the CLEAN algorithm is very similar to a modified Richardson-Lucy algorithm [106], which as discussed above has the ability to superresolve.

Although it is now generally accepted that superresolution is achievable, the standard for comparing different algorithms is not unified. For example, the definition of resolution is not standardised. As discussed in the previous subsection, it is unclear whether λ/D is used to account for the effect of the central obscuration, or is just a convenient approximation to the exact value. This introduces a possibility of misinterpreting the frequency content of the reconstruction and the extent of superresolution.

In image processing, the resolution of the image is limited by the optics and the CCD sampling. While it is reasonable to define superresolution as overcoming the limits imposed by the optics when the optics is the main limiting factor, it is necessary to reconsider this definition when the CCD sampling poses as the major limitation. Therefore, in this thesis, the term 'superresolution' is used to describe the situation when the frequency limit imposed by the major limiting factor is overcome. Practical algorithms where superresolution is achieved are discussed in Chapters 5 and 7.

4.3   Characterising an optical system

4.3.1   Coherent and incoherent imaging

In this thesis, all the images dealt with are formed assuming incoherent imaging. However, it should be pointed out that this is not the case for all imaging systems, and here a distinction is made between coherent and incoherent imaging.

Coherent illumination implies that the optical field is perfectly correlated or deterministic.

This means that all the object points have a fixed phase relationship with respect to each other. In coherent illumination, the response of each object point can be added on a complex amplitude basis. It is therefore said to be linear in complex amplitude. An example of a coherent light source is a laser.

The transfer function of a coherent system is the Fourier transform of the space-invariant impulse response:

$$\tilde{H}(u,v) = \mathfrak{F}\{\tilde{h}(x,y)\}. \tag{4.38}$$

The impulse response is also a function of the pupil function (Eq. (4.30)). This leads to:

$$\begin{aligned} \tilde{H}(u,v) &= \mathfrak{F}\{\mathfrak{F}\{P(u,v)\}\} \\ &= P(-u,-v). \end{aligned} \tag{4.39}$$

This means that the coherent transfer function is the rotated pupil function.

In the case of incoherent imaging, each object point is assumed to be statistically independent of each other. As a result, the summation of the responses from the object points is performed on an intensity basis. Incoherent imaging is therefore said to be linear in intensity. Self-luminous objects such as stars, and objects illuminated by thermal light are examples of incoherent light sources.

The incoherent transfer function is once again the Fourier transform of the corresponding impulse response. However, due to the property of linearity in intensity, the transfer function is:

$$\begin{aligned} H(u,v) &= \mathfrak{F}\{|\tilde{h}(x,y)|^2\} \\ &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} P(u',v')P(u'+u,v'+v)du'\,dv' \\ &= P(u,v) \otimes P(u,v) \end{aligned} \tag{4.40}$$

where $\otimes$ denotes a correlation. This means that the incoherent transfer function is the autocorrelation of the (generalised) pupil function.

Resolution of coherent and incoherent imaging

It is of interest to determine whether coherent or incoherent imaging provides a better resolution. Of course, for astronomical imaging all light waves concerned are incoherent.

However, in areas such as microscopy, a choice can be made as to what light source is used.

For a coherent imaging system, the cutoff frequency is related to Eq. (4.39). Since the pupil function P takes on the value of either one or zero, the same is true for the coherent transfer function. In the absence of aberrations, there exists a finite passband where all frequency components are transmitted unaltered. This frequency limit is denoted by

$$f_0 = \frac{D}{2\lambda d_i} \qquad (4.41)$$

where D is the diameter of a circular aperture or the width of a square aperture, and d_i is the distance between the lens and the image plane as defined in Section 4.1.1.

Since the transfer function of an incoherent imaging system is the correlation of the pupil function as in Eq. (4.40), the cutoff frequency occurs at $2f_0$. However, this does not imply that an incoherent imaging system has twice the resolving power of a coherent system [61]. In fact, the two cutoff frequencies cannot be compared directly because the coherent system frequency cutoff is defined for image amplitude, while that for the incoherent system is used for image intensity. The resolution of the two imaging systems therefore cannot be compared directly; the performance of the systems depends strongly on the structure of the function to be imaged, particularly its phase distribution.

Nevertheless, it should be pointed out that the effects of coherent and incoherent imaging do have similarities if the impulse responses $\tilde{h}(x, y)$ and $|\tilde{h}(x, y)|^2$ are examined. Consider the simple case of $\tilde{h}(x, y)$ being a sinc function; then the mainlobes in the two cases are not very different. The degree of blurring introduced by each function to an object is therefore similar.

4.3.2   Parameters for measuring image quality

There are many parameters that can be measured to determine the quality of an optical system. A popular parameter used in astronomical imaging is the full width at half maximum (FWHM) which measures the width of an intensity peak at half its maximum intensity [143]. However, a parameter that is more commonly used is the Strehl ratio, which compares the ratio of the peak PSF values for a system, with and without aberrations. It
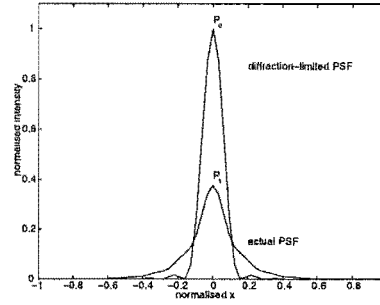
Figure 4.9: *The Strehl ratio is the ratio of the peak of the PSF in the presence of aberrations, to that in a diffraction-limited system.*

is illustrated in Fig. 4.9 and is mathematically defined as:

$$\text{Strehl} = \frac{\max(h(x,y)_{\text{actual}})}{\max(h(x,y)_{\text{diffraction-limited}})}. \tag{4.42}$$

The Strehl ratio is always between 0 and 1. In a real optical system, it is impossible to obtain a Strehl ratio of 1. A Strehl ratio of 0.8 or greater is effectively diffraction-limited, and is known as Maréchal's criterion.

One drawback of the Strehl ratio is its inability to differentiate between a broad and a narrow PSF as shown in Fig. 4.10; the dashed line is a diffraction-limited PSF, while the solid line is a Gaussian, both of which have the same Strehl ratio. Nevertheless, the Strehl ratio is a commonly used performance measure for different optical systems.

A phenomenon closely linked to the Strehl ratio is the core and halo of a point source. A point source reconstructed from Shack-Hartmann wavefront sensing data comprises two terms: a core which is essentially the main lobe of the diffraction-limited pattern, and a halo whose width is generally of the order of the seeing disk diameter. A more detailed discussion is included in Section 4.4.2.

4.3.3 Wavefront aberrations

The diffraction limit poses as an upper bound to the quality of an image. However, in many cases, this limit is not reached because of the presence of aberrations. Aberrations can arise due to defects in the optical system such as the limited resolution of a deformable mirror in an adaptive optics system or the misalignment of the components; or from a spatially and
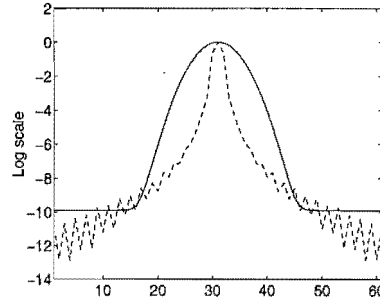
Figure 4.10: *The Strehl ratio cannot differentiate between a broad PSF (solid line) and a narrow PSF (dashed line).*
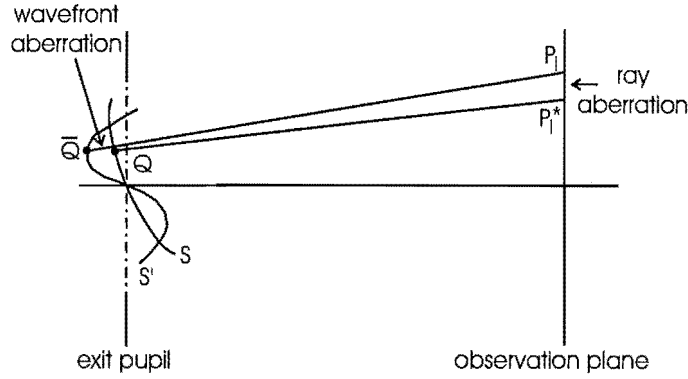
Figure 4.11: *The geometry of wavefront and ray aberrations.*

temporally varying imaging medium such as the turbulent atmosphere of the Earth when imaging a star. Aberrations are classified into two broad categories: chromatic aberrations and monochromatic aberrations. Whilst the former are affected by the wavelength of the light, the latter are not. In this thesis, only monochromatic aberrations are of interest.

Consider the optical system described by Fig. 4.11. The ideal spherical wavefront is denoted by S. In the absence of aberrations, a point P_0 which is also denoted by Q is mapped onto the point P_i^* on the observation plane. When aberrations are present, the wavefront is no longer spherical. The distorted wavefront is denoted by S' and the corresponding point P_0, now denoted by $\bar{Q}$, is mapped onto P_i on the observation plane. The distance between P_i and P_i^* is called the ray aberration. The path length between Q and $\bar{Q}$, points on the ideal and aberrated wavefronts respectively, is called the wavefront aberration $\Phi = \Phi(x_1, y_1; x_0, y_0)$ where (x_1, y_1) corresponds to the object plane, and (x_0, y_0) refers to the image plane.

By converting Φ into polar coordinates, and fixing one of the variables, the wavefront aberration can be expressed in terms of $\rho = \sqrt{x_0^2 + y_0^2}$, $\theta = \tan^{-1}(y_0/x_0)$ and y where y is the position of the object in the field of view [13]. This gives:

$$
\begin{aligned}
\Phi(\rho, \theta, y) &= a_1\rho^2 + a_2\rho y \cos(\theta) \\
&+ \; b_1\rho^4 && \text{spherical aberration} \\
&+ \; b_2\rho^3 y \cos(\theta) && \text{coma} \\
&+ \; b_3\rho^2 y^2 \cos^2(\theta) && \text{astigmatism} \\
&+ \; b_4\rho^2 y^2 && \text{field curvature} \\
&+ \; b_5\rho y^3 \cos(\theta) && \text{distortion} \\
&+ \; \text{other higher order terms} && (4.43)
\end{aligned}
$$

where a_1, a_2 and $b_1, \ldots, b_5$ are constants. These five terms are known as the primary or Seidel aberrations.

Spherical aberration, coma and astigmatism prevent the formation of point images. Spherical aberration causes the rays near the optical axis to be focused in a different location from those passing near the edge of the lens, meaning that there is no real focus in the system, while coma produces a comet-like asymmetric blur. On the other hand, field curvature and distortion allow point images to be formed, but the relative positions of the points are altered.

4.3.4   Zernike polynomials

Zernike polynomials were initially proposed for describing the aberrations of an optical telescope [13]; their application to atmospheric turbulence was later proposed by Fried [42] and extended by Noll [131].

Zernike polynomials are a set of basis functions defined on the unit circle. As a result, they are usually defined in polar coordinates as the product of radial and angular terms. Using the ordering scheme of Noll the Zernike polynomials are defined as:

$$
\begin{aligned}
Z_{\text{even } j}(\rho, \theta) &= \sqrt{n+1}\, R_n^m(\rho)\sqrt{2}\cos(m\theta), \; m \neq 0 \\
Z_{\text{odd } j}(\rho, \theta) &= \sqrt{n+1}\, R_n^m(\rho)\sqrt{2}\sin(m\theta), \; m \neq 0 \\
Z_j(\rho) &= R_n^0(\rho), \; m = 0, && (4.44)
\end{aligned}
$$

where

$$R_n^m(\rho) = \sum_{s=0}^{(n-m)/2} \frac{(-1)^s(n-s)!}{s![(n+m)/2-s]![(n-m)/2-s]!}\rho^{n-2s}. \tag{4.45}$$

Here n is the radial order and m is the azimuthal order, both of which are non-negative integers where $m \leq n$ and $n - |m| =$ even. The index j is a mode ordering number and is a function of n and m.

Zernike polynomials are useful for expanding a wavefront over a circular aperture. The expansion of such a wavefront $\phi(R\rho, \theta)$, where R is the radius of the aperture, is given by

$$\phi(R\rho, \theta) = \sum_j a_j Z_j(\rho, \theta), \tag{4.46}$$

where the coefficients a_j are

$$a_j = \int d^2\rho \mathcal{W}(\rho)\phi(R\rho, \theta)Z_j(\rho, \theta) \tag{4.47}$$

and $\mathcal{W}(\rho)$ is a unit volume weighting function

$$\mathcal{W}(\rho) = \left\{ \begin{array}{ll} \frac{1}{\pi} & \rho \leq 1 \\ 0 & \rho > 1 \end{array} \right. .$$

The first mode of the Zernike polynomials is piston. Since this does not affect the formation of and cannot be measured from images formed at the telescope, it is not included in the analysis of atmospheric turbulence. The dominant modes are the second and the third modes as they cause approximately 90% of the degradation to the wavefront [42]. These correspond to tip and tilt, and produce random motion of the images in the image plane, but do not affect the spatial resolution of the image. One application that utilises this property is the Shack-Hartmann (SH) wavefront sensor. The displacement of the image captured by a lenslet in a SH sensor is directly proportional to the mean slope (or the tip and tilt components) over that lenslet or subaperture. When this concept is extended to the entire aperture across all the lenslets, an estimate of the wavefront can be calculated. A more detailed discussion on wavefront reconstruction from wavefront sensing data is presented in Chapter 8.

4.4   Imaging through atmospheric turbulence

Light propagates from the stellar object of interest to ground through the turbulent atmosphere, and the image formed is captured by a ground-based telescope. In order to study the inverse problem of recovering the object distribution from the collected image data, it is essential to understand firstly the characteristics of the distortions introduced by the atmosphere on the propagating wavefront.

The distortion of the atmosphere is not deterministic in nature. It can be thought of as a collection of lenslets with different and time-varying refractive indices. Mathematically, it can be sufficiently modelled using Kolmogorov statistics. Because the structure of the atmosphere changes in the order of milliseconds, in order to freeze its effects, telescopic images must be captured within this time frame. Such images are called short-exposure images. On the other hand, images captured over an extended period of time, where the effects of the atmophere are averaged, are called long-exposure images. Both types of images are used in object reconstruction algorithms. Short-exposure images are used in blind deconvolution and deconvolution from wavefront sensing (DWFS), while long-exposure images are used for real-time compensation of adaptive optics imagery as well as calibrating an optical system.

4.4.1   Short and long exposure images

Short exposure images are captured within milliseconds and the atmospheric turbulence is essentially 'frozen'. On the other hand, a long exposure image is an average of a large number of independent realisations of the atmosphere. Fig. 4.12 compares a short exposure image, a long exposure image, and a diffraction-limited image for a circular aperture with $D/r_0 = 10$. The long exposure image is the sum of 1000 short exposure images. The short exposure image has a speckly appearance (hence the term speckle image) and contains some high resolution information, while the long exposure image has very little high frequency information.

Fig. 4.13 (reproduced with the permission of the author [93]) examines the relationship between the ensemble OTFs and PSFs for a circular aperture for short and long exposure images under different imaging conditions ($D/r_0 = 0.1, 1, 2$ and 10). These are compared

(a)   (b)   (c)
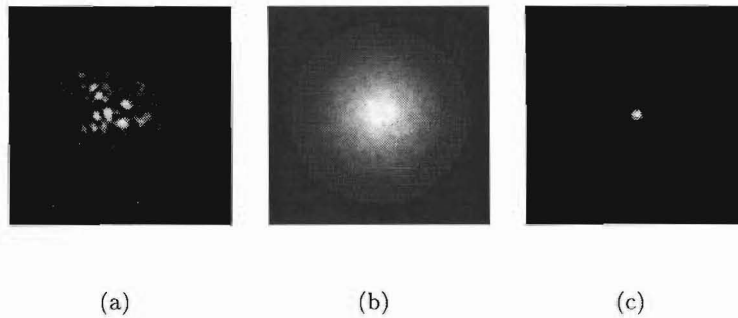
Figure 4.12: *Simulated images of a single star through a circular aperture with $D/r_0 = 10$. (a) Short exposure image. (b) Long exposure image. (c) Diffraction-limited image.*

with the corresponding diffraction-limited quantities. Short exposure images retain some high frequency information, therefore the ensemble OTFs and PSFs do not deviate much from the diffraction-limited quantities when seeing is good. The degradation is only clear when $D/r_0 = 10$. On the other hand, the long exposure OTFs do not have much high frequency information, and as a result, the corresponding PSFs are smoother and broader.

4.4.2   Core and halo structure

As introduced in Section 4.3.2, a point source reconstructed with the use of Shack-Hartmann wavefront sensing data results in a structure comprising two components, the core and the halo. This is a direct result of the inability of the SH-WFS to reconstruct certain high frequency components, causing aliasing in the reconstruction. In this subsection, the core and halo structure of the reconstruction of a point source is examined more closely, and several observations are made.

Fig. 4.14(a) shows the centre portion of a diffraction-limited image of a point source under noiseless conditions. The image is shown on a log scale to highlight the core structure which is clearly visible. Most of the energy is concentrated in the core. The maximum value of the core is -3.3008 on the log scale. Since the core is a narrow compact structure in the image plane, it contains much high frequency information in the Fourier domain.

Figs. 4.14(b), 4.14(c) and 4.14(d) show the images of a point source in the presence of
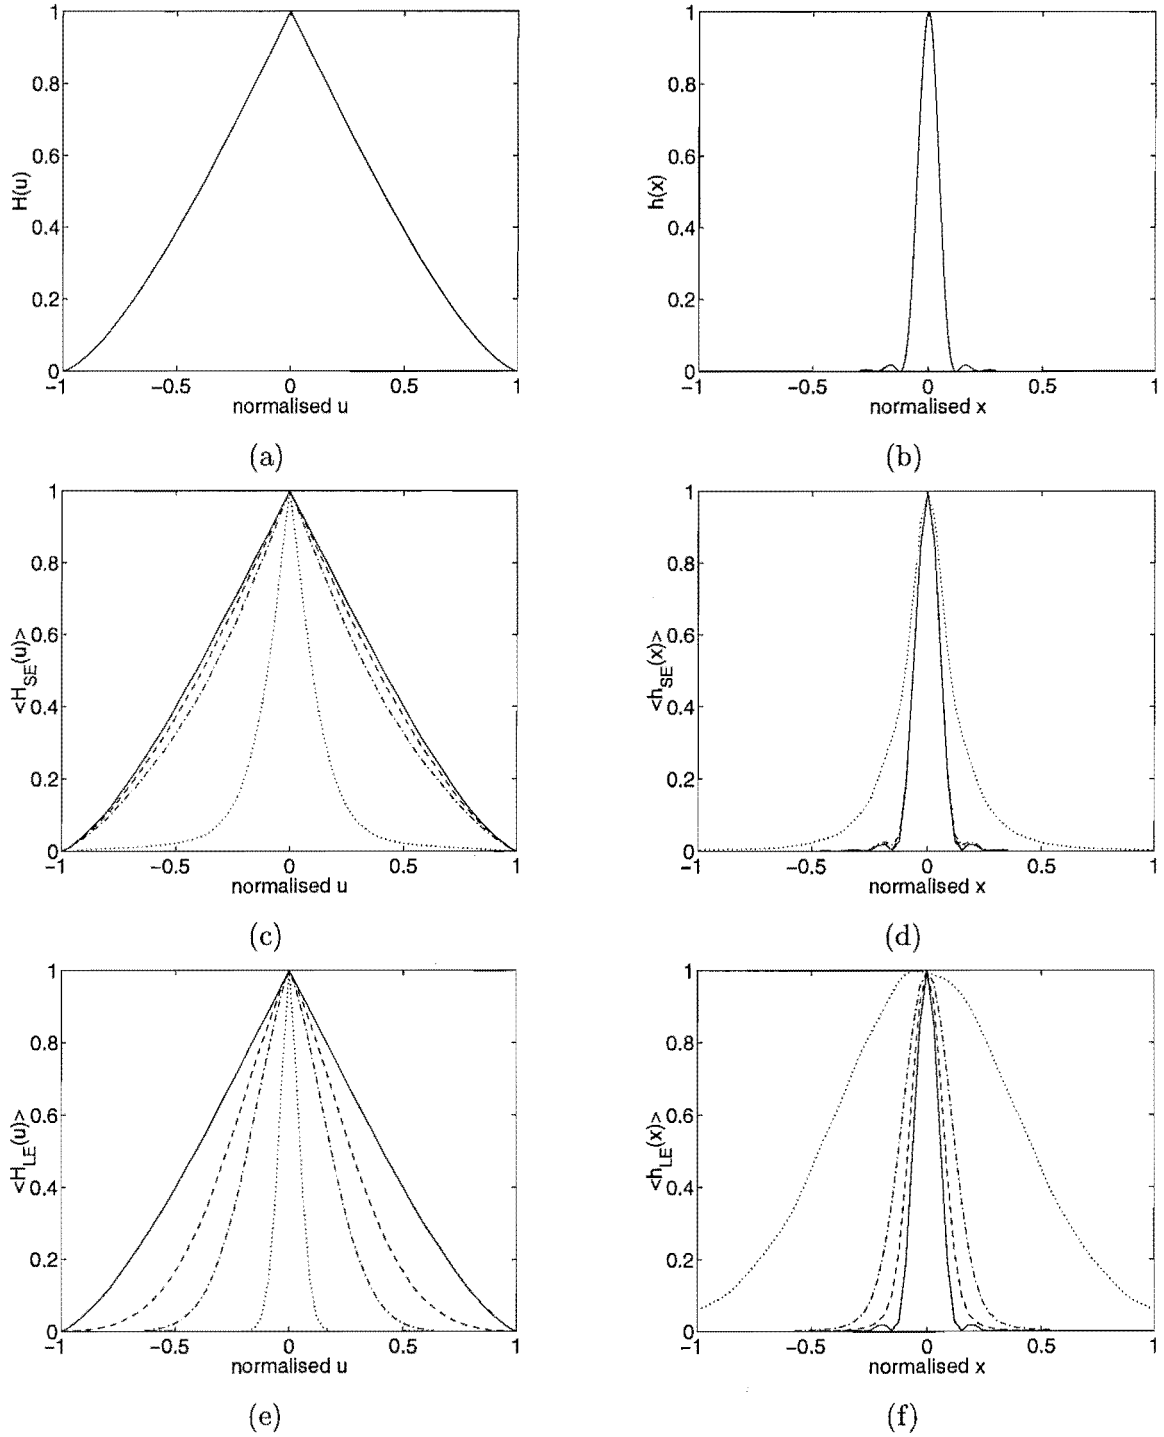
Figure 4.13: *Typical OTFs and PSFs for a circular aperture. (a) Diffraction-limited OTF, (b) diffraction-limited PSF, (c) family of short exposure OTFs, (d) family of short exposure PSFs, (e) family of long exposure OTFs and (f) family of long exposure PSFs. The family of curves correspond to $D/r_0 = 0.1$ (solid), 1 (dashed), 2 (dashed-dotted) and 10 (dotted).*

(a)

(b)   (c)

(d)   (e)
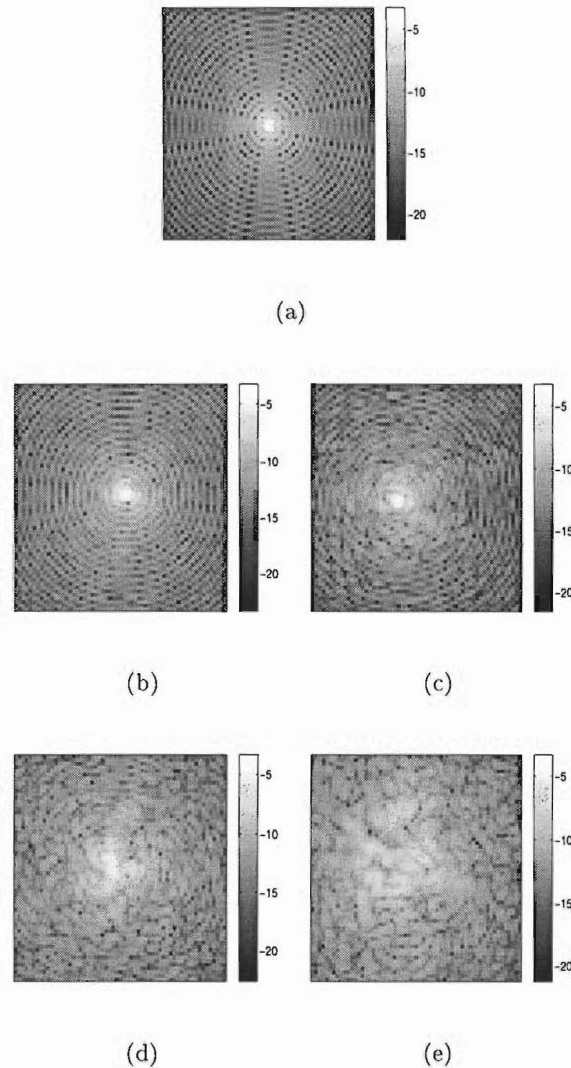
Figure 4.14: *(a) Centre portion of a diffraction-limited image of a point source. Note the prominent core structure. (b,c,d,e) Centre portions of the image of a point source in the presence of atmospheric turbulence with $D/r_0 = 1, 2, 4$ and 8. All images are displayed on a logarithmic scale.*

atmospheric turbulence with $D/r_0 = 1, 2$ and 4. The core is still visible in each case, although the halo starts to dominate. The respective maxima of the cores are -3.5650, -3.9656 and -4.3229. For $D/r_0 = 8$ (Fig. 4.14(e)), the core structure is no longer visible: the halo is the sole component of the image. The maximum in this image is -5.4646.

The effect of the atmospheric turbulence on the image can be compensated by measuring a

simultaneous frame of Shack-Hartmann wavefront sensing data and using it to estimate the phase distortion (see Chapter 8 for the theory of phase estimation from wavefront sensing data). A lenslet array of size 16×16 is used here. For the case of $D/r_0 = 8$, this corresponds to a d/r_0 of 0.5 where d is the subaperture diameter. The point source image is retrieved using a simple Wiener filter. Figs. 4.15(a) and 4.15(b) show the reconstructed point source images for $D/r_0 = 8$ using 100 and 200 Zernike polynomials respectively in the phase reconstruction. The maxima in each case are -3.4640 and -3.4930.

In both cases, the core structure is reconstructed, with little difference between the maximum values. The retrieval of the core structure means that most of the high frequency information is present in the reconstruction. However, the halo is still visible. This leads to three important results. Firstly, because the SH-WFS is unable to reconstruct some of the high frequency components in the phase, they are aliased and are reflected as the halo structure. The inability of the SH-WFS to reconstruct high frequency components is due to the finite number of subapertures in the WFS, leading to only a finite number of Zernike weights being reconstructed [111]. Secondly, once the core is reconstructed, an increase in the number of Zernike polynomials does not improve the reconstruction. Finally, although the level of the halo is lowered, it is difficult to reduce it significantly. These problems affect the methods of deconvolution from wavefront sensing (DWFS) and adaptive optics (AO). As a result, it is desirable to apply post-processing techniques to improve the resolution and the visual quality of the image. A class of algorithms known as myopic deconvolution has been proposed for post-processing AO imagery [24,27,51,126] and is briefly discussed in Chapter 8. A proposed solution motivated by these problems in DWFS is presented in Chapter 9.

4.4.3   Wavefront statistics and structure functions

When the wavefront of the stellar object propagates to ground, the distortion on the wavefront is predominantly a phase component. However, the intensity of the wavefront is also distorted due to the propagation. This intensity distortion, also called scintillation, can be seen as the twinkling of the star when observed from ground by the naked eye. It should be noted that because the scintillation is small compared to the phase distortion, it is usually assumed constant in many algorithms. In this section, the statistics of the phase distortion and the scintillation are discussed; their application in image reconstruction algorithms is

(a)   (b)
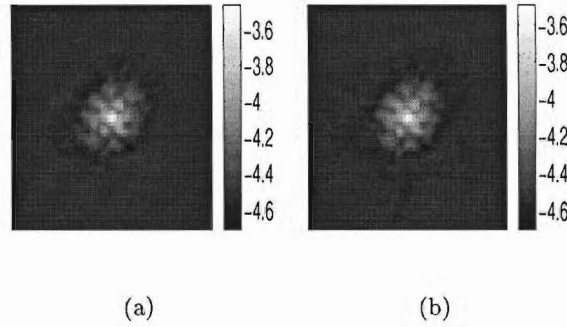
Figure 4.15: *Reconstruction of the point source image for the case of $D/r_0 = 8$ with a 16×16 SH lenslet array using (a) 100 and (b) 200 Zernike polynomials in the phase reconstruction. In each case, the core is retrieved, but the halo is still present.*

presented in Chapter 6.

For simplicity, it is assumed that the wavefront is a monochromatic plane wave at wavelength λ, propagating from a star at the zenith, towards a ground-based observer. Following the thin screen approximation, the layer of turbulence is confined between height h and $h + \Delta h$. The complex wavefront at height $h + \Delta h$ is $A_{h+\Delta h} = 1$. At h, the complex wavefront is:

$$A_h = |A_h(u,v)| \exp\{i\phi_h(u,v)\} \tag{4.48}$$

where the modulus $|A_h(u,v)| = 1$ because it has not been propagated, and $\phi_h(u,v)$ is the phase distortion introduced by the turbulence. According to Roddier [140], the covariance $C(u_1, v_1, u_2, v_2)$ of the complex field is

$$
\begin{aligned}
C(u_1, v_1, u_2, v_2) &= \langle A_h(u_1, v_1) A_h^*(u_1 - u_2, v_1 - v_2) \rangle \\
&= \exp\{-\frac{1}{2} D_\phi(u_1, v_1, u_2, v_2)\}
\end{aligned}
\tag{4.49}
$$

where $\langle \cdot \rangle$ represents an ensemble measurement, and $D_\phi(u_1, u_2, v_1, v_2)$ is the two-dimensional horizontal structure function defined as

$$D_\phi(u_1, u_2, v_1, v_2) = 6.88 \left(\frac{\sqrt{(u_1 - u_2)^2 + (v_1 - v_2)^2}}{r_0} \right)^{5/3}, \tag{4.50}$$

with r_0 being the effective diameter of the telescope [147]. It should be noted that the phase covariance has a Fourier transform relationship with its power spectrum, denoted by:

$$W_\phi(f) = 0.38\lambda^{-2} f^{-11/3} \int_0^\infty dh C_N^2(h)(\cos \pi \lambda h f^2)^2 \tag{4.51}$$

where f is the spatial frequency with dimension radian^{-1} and $C_N^2(h)$ is the refractive index structure constant with a dependence on height.

Even though the covariance in Eq. (4.49) is that of the complex field, it is usually taken to be the phase covariance because the intensity fluctuations are comparatively small. However, this expression cannot be used directly to describe the phase because its DC term is infinite. By confining the calculation over a finite region and removing the mean of the phase, Wallner showed that the covariance of the mean-removed phase C_p is [178]

$$
\begin{aligned}
C_p(u_1, v_1, u_2, v_2) \;=\; & -\frac{1}{2} D_\phi(u_1, u_2, v_1, v_2) \\
& + \frac{1}{2} \int D_\phi(u_1', v_1', u_2, v_2) W_A(u_1', v_1') du_1' dv_1' \\
& + \frac{1}{2} \int D_\phi(u_1, v_1, u_2', v_2') W_A(u_2', v_2') du_2' dv_2' \\
& - \frac{1}{2} \int \int D_\phi(u_1', v_1', u_2', v_2') W_A(u_1', v_1') W_A(u_2', v_2') du_1' dv_1' du_2' dv_2'
\end{aligned}
\tag{4.52}
$$

where W_A is the finite region of the aperture.

The power spectrum of the scintillation is defined by Roddier as [140]

$$
W(f) = 3.9 \times 10^{-2} k^2 f^{-11/3} \int_0^\infty dh C_N^2(h) (\sin \pi \lambda h f^2)^2
\tag{4.53}
$$

and the scintillation covariance is the inverse Fourier transform of Eq. (4.53). One point of interest is that the exponent $-11/3$ in both Eqs. (4.51) and (4.53) is a characteristic of Kolmogorov statistics. The effect of non-Kolmogorov statistics can be investigated by altering this exponent [93].

Since $C(u_1, v_1, u_2, v_2)$ is a random complex field, an alternative to calculating the intensity and phase covariances is to find the covariances of its real and imaginary parts. It is reasonable to assume that both these quantities have the same form due to the random nature of the complex field. In fact, these covariance expressions are defined by Eq. (4.49) [140]. Fig. 4.16 shows the covariances of the scintillation, phase, and real and imaginary parts for a complex field of size 16×16. When using these covariances, the two-dimensional quantities of the complex field are vectorised to size 256×1. It should be emphasised that the phase covariance corresponds to Eq. (4.52) and is calculated analytically; on the other hand, the covariances of the real and imaginary parts are described by Eq. (4.49). These covariance structures have been generated assuming a turbulent layer at 5km above ground, and a telescope aperture with diameter $D = 1$m. The effective diameter of the

telescope r_0, also known as Fried's parameter, is set to 0.25m here. The covariances of the scintillation and the real and imaginary parts are symmetric Toeplitz matrices, while the phase covariance is a symmetric block Toeplitz matrix.

A common parameter for describing the severity of atmospheric turbulence is D/r_0, where the larger the value, the more turbulent the atmosphere. When $D < r_0$, the resolving power is limited by the telescope; whereas when $D > r_0$, it is limited by the atmosphere [140]. Due to this characteristic, the configuration of a Shack-Hartmann wavefront sensor is usually designed to have approximately $D/r_0 = 1$ over each lenslet to optimise between the accuracy of the wavefront slope calculation and the severity of the turbulence that can be measured.

(a)   (b)

(c)   (d)
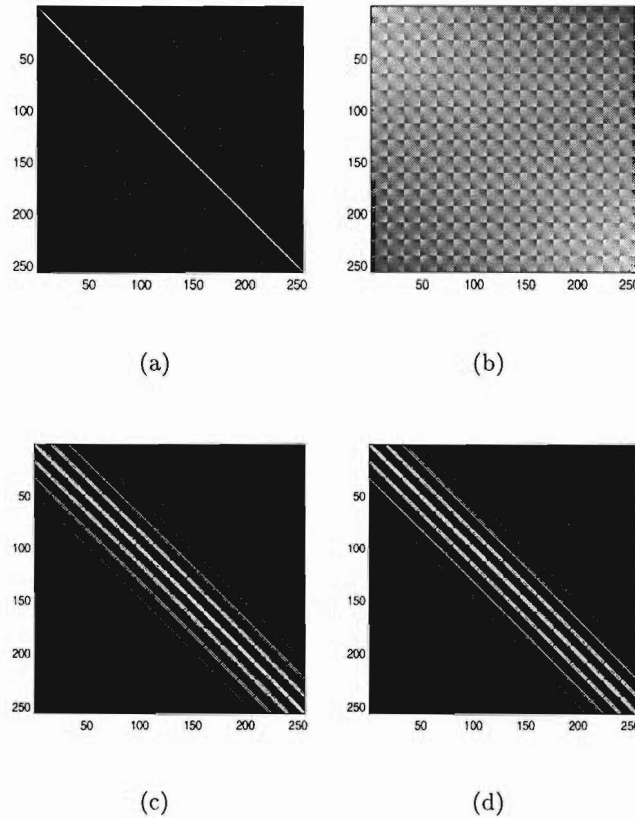
Figure 4.16: *The covariances of the (a) scintillation, (b) phase, and (c) real and (d) imaginary parts for a complex field of size 16×16. When using these covariances, the two-dimensional quantities of the complex field are vectorised to size 256×1. The severity of the atmospheric turbulence simulated for these quantities corresponds to $D/r_0 = 4$. A black pixel has a value of 0, and a white pixel has a value of 1.*

4.5 Dealing with real astronomical data

In general, when one performs simulations to test a theory or idea, many aspects that are present in the real world are ignored temporarily to simplify the situation. This is the correct approach because it isolates the factor that is being investigated, and a direct relationship can be identified between the cause and the effect. The aspects that are usually ignored include noise and other linear and non-linear effects that affect the observed data. However, we do not live in an ideal world. In order to handle real data obtained using real instruments, the effects of noise and those imposed on the data by the instruments must be taken into account.

In astronomical imaging, the wavefront at the aperture of the telescope undergoes several stages of distortion introduced by the data-capturing process. Before the data can be used, the distortions must be corrected. This section discusses the main elements involved in processing the raw data captured by the telescope to a usable form.

4.5.1 Data-capturing process

The real data for this thesis was captured at the Observatoire de Lyon, France during a collaborative research visit. This section therefore discusses the data-capturing process of the SPID (SPeckle Image by Deconvolution) instrument developed at the observatory [3].

The SPID instrument has four modules: interface, wavefront sensor, imaging and calibration (Fig. 4.17). The interface module allows the instrument to be installed to any telescope by collimating the beam from the telescope and making an image of the pupil. The wavefront sensor (WFS) is a Shack-Hartmann sensor, with 5 interchangeable arrays of lenslets (9×9, 14×14, 20×20, 28×28, 40×40). The wavefront sensing data is captured on a CCD camera running at video rate. The imaging module consists of two channels at two spectral bandwidths, producing four realisations of any speckle image captured on the CP40 photon-counting detector. The bandwidths range between 400nm and 720nm, and the two channels correspond to the blue and green sections of the spectrum. This set-up is used to overcome the problem of the photon-counting hole [167]. Both the shutters on the wavefront sensing channel and the data channels can take on six different speeds ranging from 2ms to 20ms. To minimise the human interface during an observation run, the instrument is as automated

as possible, requiring only several input parameters from the user prior to data-capturing. The calibration module consists of a square grid of point sources for calibrating the CP40 detectors, a reference wavefront for calibrating the wavefront sensor, and a binary source at a fixed distance on a rotatable platform to calibrate the scaling between the two CP40 channels and the CCD camera. For a more detailed discussion of the SPID instrument and the associated software, the reader is referred to Ref. [3].
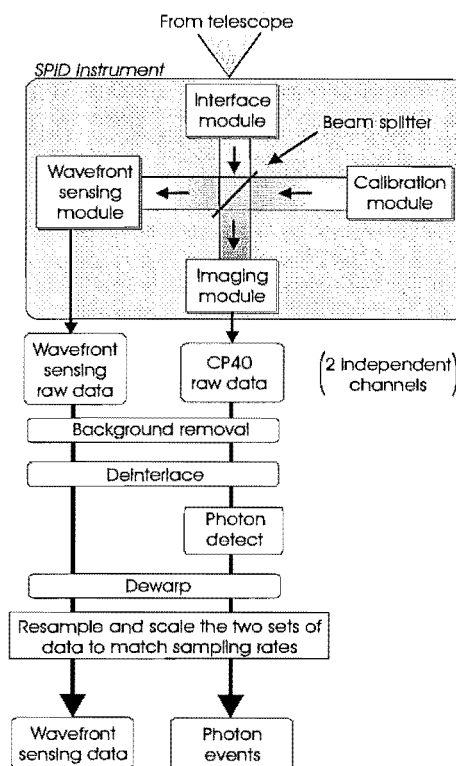
Figure 4.17: *Block diagram of the raw-data collection of the SPID instrument at the Observatoire de Lyon, and the processing steps required to convert this data into a usable form.*

After the raw data has been captured by the CP40 and CCD cameras, the distortions introduced by the non-ideal effects must be compensated for before the data is in a usable form. Although it is a general rule that the data should not be manipulated in order to avoid information loss, in this case it is necessary to perform some pre-processing steps. There are two main reasons for this. Firstly, the raw data is extremely noisy so that it is difficult to extract useful information without dealing with the noise initially. This becomes apparent by examining the sample speckle image and wavefront sensing data frame in Figs.

4.18(a) and 4.18(b). Secondly, there is a strong deterministic distortion in the data channels. Although it is possible to incorporate this into the model for the calculations, it is easier to correct for this distortion before other processing is performed. These pre-processing steps are illustrated in Fig. 4.17 and are discussed in further detail in the following sections.

(a)   (b)
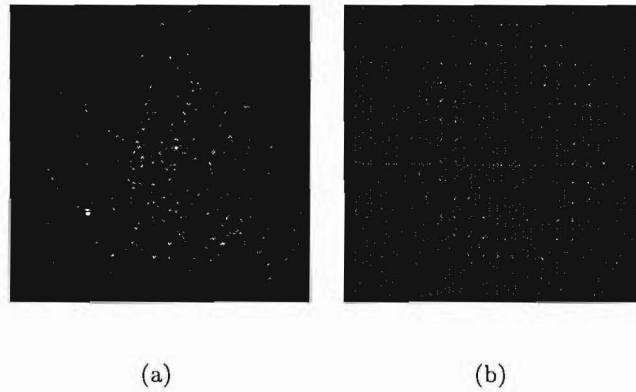
Figure 4.18: *(a) A sample speckle image and (b) a sample wavefront sensing data frame captured at the telescope. These images represent the raw data captured, and as a result no pre-processing has been performed. It can be seen that the data frames are very noisy.*

Background removal

A background is always present on the CP40 and the CCD cameras due to thermal noise, even when the cameras have been cooled with liquid nitrogen. This background measurement must be recorded prior to data-capturing so that the bias it introduces can be subtracted from the astronomical data. A different background measurement is needed before each observing run.

Deinterlacing

Each CP40 channel has 576 × 384 interlaced pixels, each of size 16.8μm (vertical) ×9.2μm (horizontal). The deinterlacing operation involves taking every second pixel in the vertical direction, thus producing two frames corresponding to the odd and even columns of the

captured data, each of size 288×384 pixels. For further details, the reader is referred to Ref. [38].

Photon detection

The CP40 photon-counting detector captures four realisations of the speckle image at two spectral bandwidths. Each of the four images is corrupted by noise and the presence of ions. The noise component is a combination of Poisson and Gaussian processes. The Poisson noise is a result of the random fluctuations of the number of photons on the detector surface, whereas the Gaussian noise arises from the random thermal motion of the electrons in the circuit elements [4, 19]. The ions present in the atmosphere produce very bright spots compared to a photon. Although the arrival of photons is a statistical process, to distinguish the photon events from the ions and the noise components, three deterministic steps must be performed. Firstly, a basic thresholding procedure selects the points that are between the upper and lower thresholds. These thresholds can be determined from the data. Secondly, the potential photon events must be a maximum over a 3×3 patch, and that no negative values must be adjacent to such a maximum because such an occurrence implies a very low probability of a photon event. Finally, the potential photon events are centroided over 3×3 pixels to improve the accuracy of their position to a subpixel level. The accuracy of the centroiding operation is examined in Chapter 7.

One of the intrinsic problems of a photon-counting detector is that it is unable to distinguish between overlapping photon events. This defect is particularly relevant when computing an image correlation for speckle interferometry techniques. This problem is overcome by using more than one realisation of the speckle, hence the four channels used by the camera [167].

Dewarping

Due to the use of a photo-cathode in the CP40 photon-counting detector, an electric field is present, resulting in a warped image being captured. This is illustrated by imaging a square grid of point sources, the result of which is shown in Fig. 4.19. The fact that the image should be a square grid provides a way of dewarping the CP40 channels using a least-squares estimate [115]. If the Cartesian coordinates of the n warped points are

denoted as (x_i, y_i), $i = 1 \ldots n$, and the expected location of the corresponding points are (u_i, v_i), $i = 1 \ldots n$, then a third order relationship between these two sets of points is:

$$u_i = a_0 + a_1 x_i + a_2 y_i + a_3 x_i^2 + a_4 x_i y_i + a_5 y_i^2 + a_6 x_i^3 + a_7 x_i^2 y_i + a_8 x_i y_i^2 + a_9 y_i^3$$
$$v_i = b_0 + b_1 x_i + b_2 y_i + b_3 x_i^2 + b_4 x_i y_i + b_5 y_i^2 + b_6 x_i^3 + b_7 x_i^2 y_i + b_8 x_i y_i^2 + b_9 y_i^3 \quad (4.54)$$

where a_j and b_j for $j = 0 \ldots 9$ are the row and column coefficients respectively. Using these coefficients, the photon-counting images of astronomical subjects can be corrected. Fig. 4.20(a) shows the points corresponding to the peaks of each of the spots in Fig. 4.19. Fig. 4.20(b) shows the grid of regular points for the warped points to be mapped onto, and Fig. 4.20(c) illustrates the dewarped points. The row and column coefficients that produced this transformation are listed in Table 4.2.

Although the warping of the CP40 channels is different for each observing run, the variation is small. This means that even for subsequent observing runs, the first set of coefficients that was computed can be used as a rough correction, while a second application of dewarping can be used to take into account the small variations between observing runs.
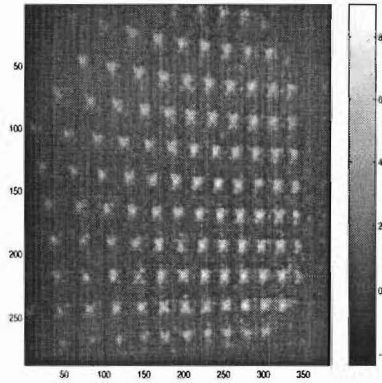
Figure 4.19: *An image recorded by the blue channel (odd) of the CP40 camera when a regular grid of point sources is imaged.*

Scaling and resampling between CP40 and WFS channels

During an observing run, data such as the image of the aperture is captured in the aperture plane by the wavefront sensing CCD, while the photon-counting images are captured on the image plane by the CP40 camera. These two cameras have the same number of pixels,

(a)

(b)   (c)
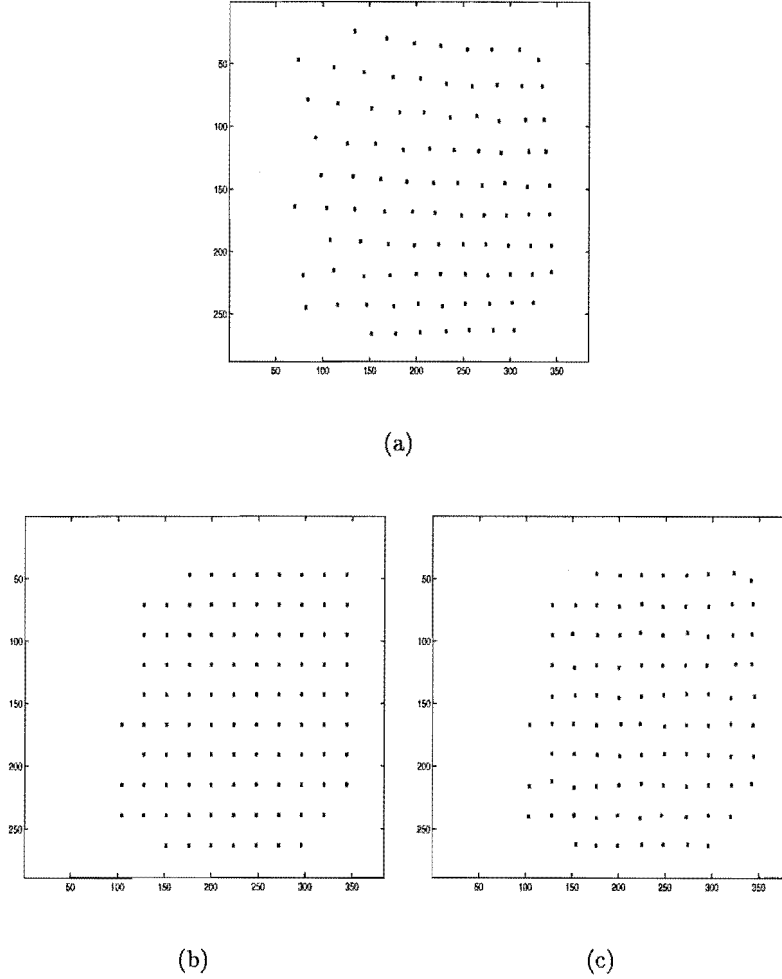
Figure 4.20: *(a) Warped image of a square grid of point sources on a CP40 channel. (b) A regular grid of points to map the warped points onto. (c) Corresponding dewarped image.*

but the sampling in each domain is different. Moreover, the alignment between the two channels is affected by the dewarping operation. As a result, one of the two sets of images must be resampled before they can be used simultaneously. This is an example of image rectification as defined in Chapter 1. This is discussed in further detail in Chapter 8, where resampling is performed before the wavefront sensing data is used as a constraint in a blind deconvolution framework.

Row coefficients $a_0 - a_9$	Column coefficients $b_0 - b_9$
12.3×10^1	8.49×10^1
7.51×10^{-1}	-7.03×10^{-2}
2.43×10^{-2}	9.45×10^{-1}
9.02×10^{-4}	9.52×10^{-5}
1.21×10^{-4}	8.41×10^{-4}
1.55×10^{-4}	7.84×10^{-5}
8.54×10^{-7}	-7.82×10^{-8}
2.40×10^{-8}	-1.49×10^{-6}
-1.40×10^{-6}	-3.99×10^{-8}
3.02×10^{-7}	-9.71×10^{-7}

Table 4.2: *The row and column coefficients for dewarping the CP40 channels.*

4.6 An overview on existing methods

The ability to image a distortion-free astronomical object from ground-based telescopes is an active area of research. The three main approaches to this problem are real time adaptive optics (AO), computer post-processing methods such as blind deconvolution and bispectrum, and hybrid methods such as deconvolution from wavefront sensing (DWFS). This section provides a brief overview of each approach.

4.6.1 Computer post-processing

Blind deconvolution

Blind deconvolution is the process of retrieving both the object of interest and the point spread function simultaneously when only the convolution of these quantities is measured. Mathematically, it is denoted by:

$$d_k(x,y) = f(x,y) \odot h_k(x,y) + n_k(x,y) \qquad (4.55)$$

where $f(x,y)$ is the object, $h_k(x,y)$ is the PSF, $n_k(x,y)$ is the noise, and $d_k(x,y)$ is the measured data. The subscript k denotes the k^{th} realisation of the respective quantities. Due to the ill-posed nature of the problem, the quality of the reconstruction depends on the prior information and constraints available. Fig. 4.18(a) shows a typical frame of data captured at the Observatoire de Lyon. The object of observation is the binary star Castor. Due to
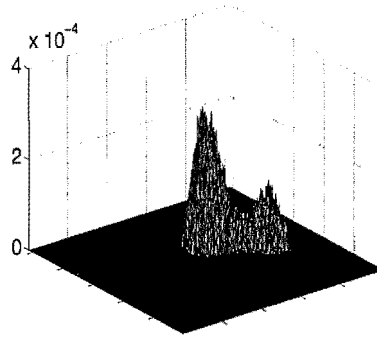
Figure 4.21: *The reconstruction of Castor using multiframe blind deconvolution with prior information.*

the noise and the low light level, it is impossible to distinguish any useful information by inspecting a single frame of image.

By applying a multiframe blind deconvolution algorithm and using an additional constraint (in the form of wavefront sensing data), it is possible to reconstruct the binary object shown in Fig. 4.21. This algorithm is discussed in Chapter 8. There are many other forms of prior information and constraints that can be used in blind deconvolution. A detailed review is given in Chapter 6, where the application of two new forms of prior is also investigated.

Bispectrum

The technique of bispectrum was first proposed by Lohmann *et al.* [113]. It uses the phase closure characteristic and therefore enables the retrieval of both the magnitude and the phase of the object spectrum without the use of a reference star under certain circumstances. Although it is a robust technique, its mathematical complexity means that it is not as popular as the blind deconvolution technique. In Chapter 6, a blind deconvolution algorithm which combines elements of the bispectrum technique is investigated.

4.6.2 Adaptive optics

Adaptive optics (AO) is a means of sensing the aberrations induced by the atmosphere and partially correcting for it in real time. Fig. 4.22 illustrates an AO system. The

incident wavefront is sensed by the wavefront sensor (WFS), and the computer calculates the commands for the actuators to change the deformable mirror (DM) in such a way that the distorting effect of the atmospheric turbulence on the incoming wavefront is minimised.

If an adaptive optics imaging system works perfectly, then diffraction-limited imaging is possible. However, this is impossible to achieve because the measurements do not have infinite precision [147]: the finite signal and measurement noise in the WFS, the finite spatial sampling of the WFS, the finite degrees of freedom in the DM, and the finite temporal response of the AO system all contribute to the fact that only partial correction is possible.
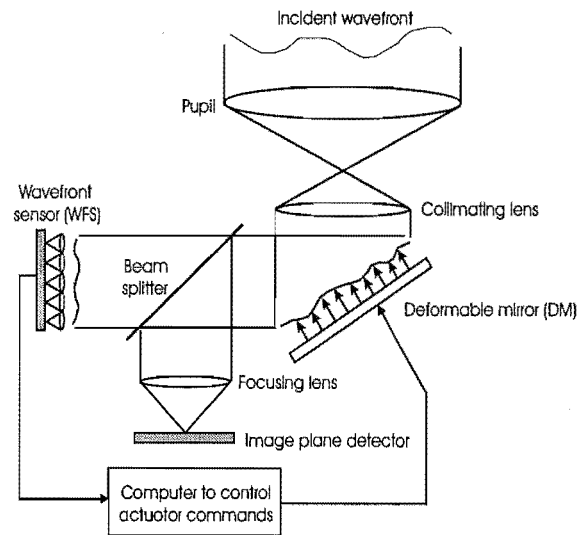
Figure 4.22: A simple illustration of an adaptive optics imaging system.

Another problem suffered by AO systems is low light levels. Using the object of interest to acquire wavefront sensing measurements, known as self-referenced adaptive optics, is sometimes difficult because the object is dim. One solution to this problem is to use a reference star to provide the wavefront sensing measurements. There are two types of reference stars: a natural reference star, and an artificial guide star. The latter is a concept first proposed by Foy and Labeyrie [39] in 1985. Here, a laser beam is focused at the centre of the field of view of the telescope at a finite altitude above ground. When the light is scattered by the atmosphere, some of this light is directed back towards the telescope and is used for wavefront sensing.

However, the use of a natural guide star and an artificial guide star introduces the problems of angular and focus anisoplanatism [45] respectively. Angular anisoplanatism (Fig. 4.23(a)) arises from the wavefront of the natural guide star propagating through slightly different regions of the atmosphere to the telescope than the wavefront of the object. On the other hand, focus anisoplanatism is the result of the artificial guide star being at a finite altitude, compared with the object which is effectively infinitely far from the telescope (Fig. 4.23(b)). Both the problems of anisoplanatism and finite measurements contribute significantly to the performance of the adaptive optics system. It is therefore usual to employ some post-processing techniques on the partially compensated images to improve their quality.

In this thesis, the techniques and performance of adaptive optics systems are not examined; for further details the reader is referred to Ref. [147]. However, because the technique of adaptive optics suffers from similar problems to deconvolution from wavefront sensing (discussed next), some of the techniques proposed in this thesis are also applicable to AO imagery.
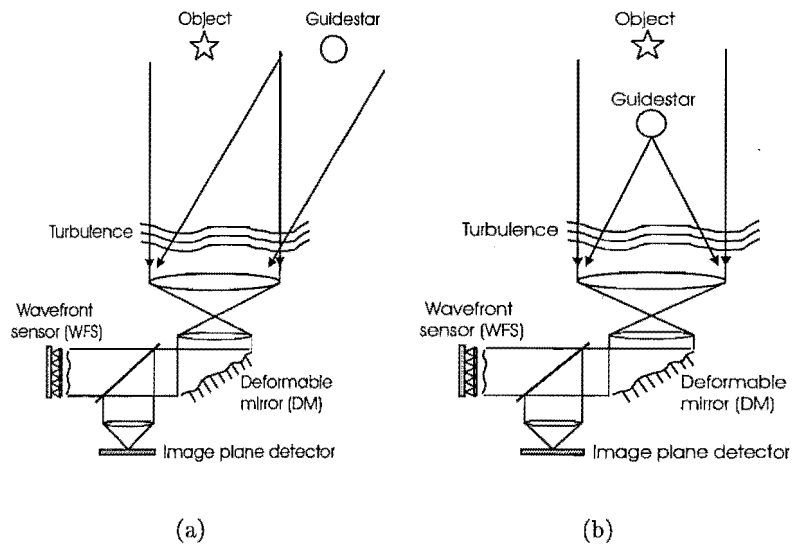
Figure 4.23: (a) The use of a natural guide star suffers from angular anisoplanatism. (b) The use of an artificial guide star suffers from focus anisoplanatism.

4.6.3 Hybrid technique: deconvolution from wavefront sensing

Deconvolution from wavefront sensing (DWFS) is a hybrid imaging technique which combines the elements of adaptive optics and post-detection image processing to obtain high resolution images through atmospheric turbulence. In addition to speckle images, the technique of DWFS uses simultaneous wavefront sensing images to estimate the optical transfer functions in a post-processing step. This technique was first proposed by Fried [46] and later implemented by Primot *et al.* [135]. Fig. 4.24 shows a block diagram of the DWFS method.
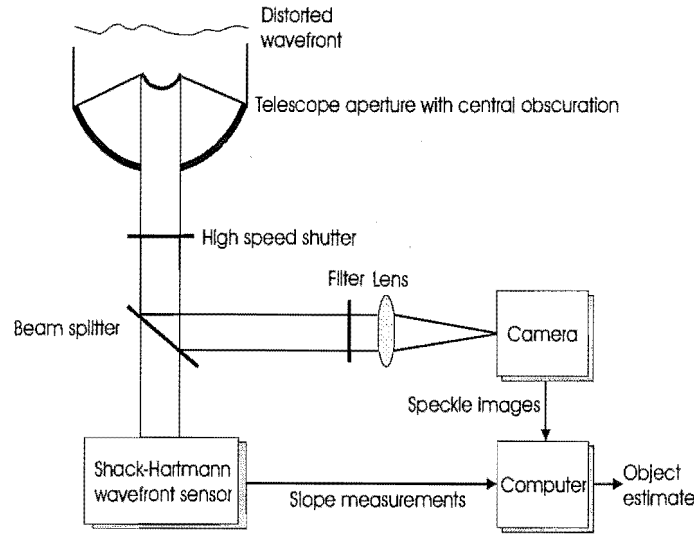
Figure 4.24: *Block diagram of the DWFS method.*

The basic formulation of a DWFS estimator is:

$$\hat{F}(u,v) = \frac{\langle D(u,v)\hat{H}^*(u,v)\rangle}{\langle |\hat{H}(u,v)|^2\rangle + c} \tag{4.56}$$

where the object distribution can be found by taking the inverse Fourier transform of $\hat{F}(u,v)$. This is essentially a vector Wiener filter [133,150]. Fig. 4.25 illustrates an example of applying DWFS. The object of interest is a binary star and it is imaged in the presence of Poisson noise in the image channel and Gaussian noise in the WFS channel. Fig. 4.25(a) shows a sample frame of simulated speckle image, and Fig. 4.25(b) shows one realisation of a wavefront sensing image. The reconstruction of the object using DWFS with 2000 frames is shown in Fig. 4.25(c).

(a)   (b)

(c)
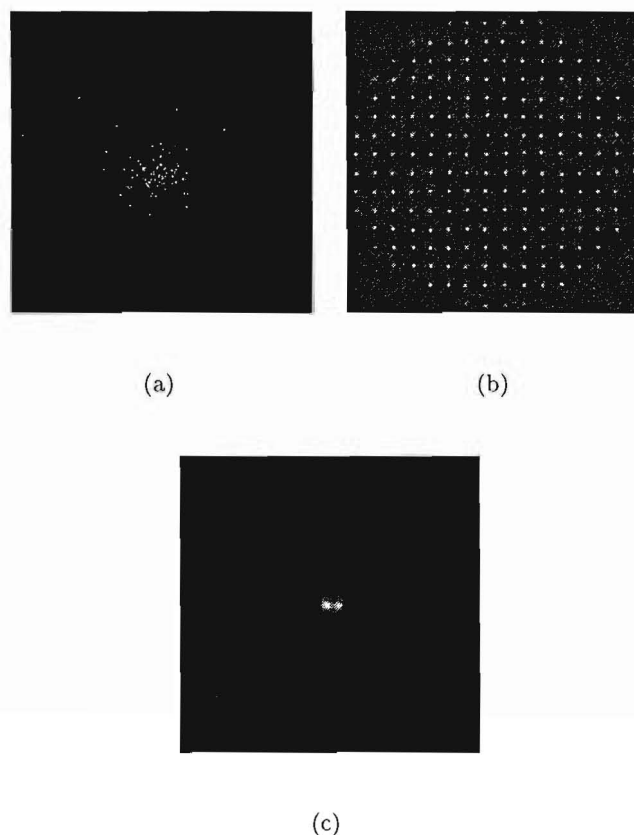
Figure 4.25: *An example illustrating the method of DWFS. (a) Sample frame of a speckle image of a simulated binary star. (b) Sample frame of wavefront sensing data. (c) Object reconstruction using 2000 frames of data.*

The method of DWFS is simple to implement. However, the use of the Shack-Hartmann wavefront sensor means that the magnitude of the object spectrum cannot be reconstructed accurately. A possible solution is to use reference star measurements in addition to images of the object of interest [149]. However, this does not completely remove the magnitude problem and suffers from other difficulties associated with a guide star as discussed in the previous section. In Chapter 9, an alternative solution based on parameterisation is proposed.

Chapter 5

Interpolation and superresolution

Interpolation and superresolution can be viewed as being complementary to each other. Interpolation describes the increase in sampling rate in the spatial domain, while superresolution depicts the extension of the meaningful frequency content of an image. This interpretation is supported by the observation that if the frequency content of the data increases, the sampling grid must be more finely spaced [85]. In this chapter, the problem of interpolation is developed from a statistical point of view. The results are interpreted in terms of the frequency content of the reconstruction.

5.1 Introduction

Many images captured by charge-coupled device (CCD) cameras, scanners, and satellites are of a lower resolution than desired because the pixel size must be large enough to minimise the amount of read noise. The data is often aliased and interpolation is required to recover some high frequency information.

In remote sensing, the images are often captured by different sources and are seldom aligned with each other. In order to compare the images directly, it is convenient to resample the measurements onto the same set of coordinates. Instead of using a geometric transformation, which gives rise to the problem of non-rectangular and skewed pixels, the data can be interpolated onto a finer grid and then downsampled (or decimated) onto the desired set of

rectangular coordinates.

Astronomical imaging also calls for the need for interpolation. When simultaneous speckle images and wavefront sensing data are captured, the sampling rates between the data sets may not match. This is the case with the data captured at the Observatoire de Lyon, which is used in this thesis. It is therefore necessary to perform interpolation on the wavefront sensing data in order to match the sampling rate of the speckle images. This is discussed in further detail in Chapter 8.

Since the need for interpolation arises in many occasions, there are many methods devised for it. These include the common linear interpolator, the sinc and its windowed counterpart, polynomial interpolators [98] and splines [81].

Interpolation is an inverse problem as it deduces the underlying high-resolution object from a set of low-resolution measurements by assuming certain characteristics about the image-acquisition instrument. These assumptions form the model of the system, a key component of which is the sampling process. A property of many interpolation schemes is collocation, where the original sample values are preserved. Whilst this would be optimal for a noise-free point-sampled image, it is clearly not optimal for noisy data or CCD images where a cell integrates over an area to produce a sample. In fact, it has been shown that a better interpolation result is produced when the sampling process is modelled correctly [182]. Fig. 5.1 illustrates the processes of area and point sampling. In Fig. 5.1(a), a section of a standard test image in Matlab [117] is shown. Figs. 5.1(b) and 5.1(c) illustrate the images using area and point sampling respectively. The resolution of the sampled images is a quarter of the original image. It can be seen that the area-sampled image has a smoother appearance. Fig. 5.1(d) shows the respective spectra. The solid line is the true spectrum. The spectrum of the area-sampled image, shown as a dash-dotted line, is closer to the true quantity than the spectrum of the point-sampled image (dotted line).

Another issue associated with sampling is the sensitivity and the fill factor of the sensors in an instrument. The fill factor is the proportion of a CCD cell that is sensitive to light. If the cell is of dimensions $X_0 \times Y_0$ (Fig. 5.2), and the sensitive area is of dimensions $a \times b$, then the fill factor is

$$\frac{a}{X_0} \times \frac{b}{Y_0}. \tag{5.1}$$

If $a = X_0$ and $b = Y_0$, then the sensors are butted. In real instruments, the sensors are

not butted and do not have uniform sensitivity. However, the model complexity increases significantly if these characteristics are to be incorporated. As a result, these factors are not used in the model in many interpolation schemes.

When performing interpolation, often some statistical information is known or can be extracted from the measured data. For example, the direct simulation of atmospheric phase distortion is only practical for an array of size 32 × 32, which is insufficient for many processes [76]. By using a weighted least squares approach for interpolation together with randomisation according to the known statistics of the turbulence, a small phase screen is interpolated with the correct statistics. Although the application of statistical information seems reasonable, it is seldom used in many common interpolation schemes.

In view of the many existing interpolation methods which neither account for the nature of the sampling process nor incorporate statistical information in the reconstruction, in this chapter an optimal Bayesian interpolator is sought as a solution to these problems. This interpolator is optimal in the sense that it is linear, unbiased, and has minimum variance [97]. Classical interpolators such as the sinc and bilinear interpolators are shown to be a subset of the family of derived optimal interpolators.

5.2   Instrument sampling process

The data we wish to capture is continuous, but it must be discretised to be processed digitally. As a result, the high resolution object of size $M \times M$ is denoted by $f(x,y)$, while the low resolution image of the same scene is denoted by $d(x,y)$ (size $N \times N$). Here $N = M/R$, where R is the ratio between the sampling rates of $f(x,y)$ and $d(x,y)$, and is an integer greater than 1. The relationship between the high and low resolution data can then be expressed as:

$$\mathbf{d} = \mathbf{Hf} + \mathbf{n} \qquad\qquad (5.2)$$

where $\mathbf{f}$, $\mathbf{d}$ and $\mathbf{n}$ are lexicographically ordered vectors of the unknown high resolution data, the measured low resolution samples, and additive noise values, respectively. All of these vectors are assumed to be stationary random processes, and the noise in this system is assumed to be zero-mean white Gaussian noise. $\mathbf{H}$ is a matrix describing the sampling process, which is the transformation from the high resolution data to the low resolution

(a)   (b)   (c)

(d)
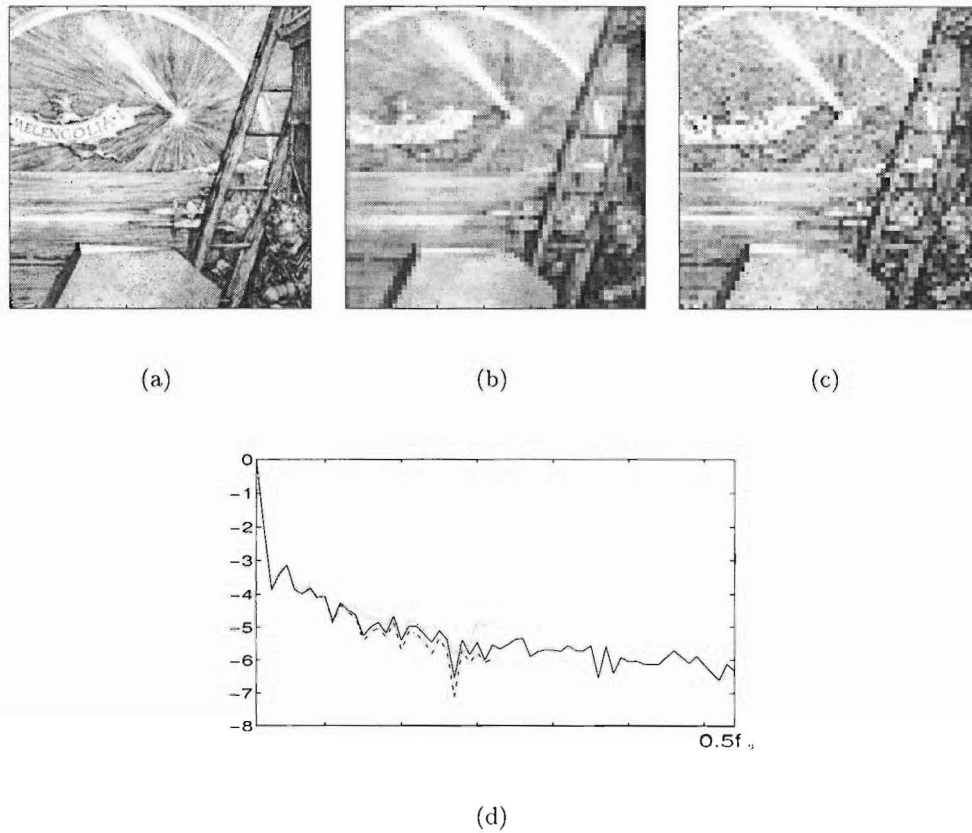
Figure 5.1: *(a) A section of a standard test image in Matlab. It is down-sampled by a factor of four using (b) area sampling and (c) point sampling. (d) The solid line is the spectrum of the true image. The spectra of the area- and point-sampled images are shown as dash-dotted and dotted lines respectively on a logarithmic scale. The area-sampled image is a closer representation of the true image.*

data. The inverse problem here is to find $\mathbf{f}$ when $\mathbf{d}$ is measured. This is an underdetermined problem because there are fewer data points than unknowns. It should be noted that since R can be made arbitrarily high, the reconstructed data can approach continuous data in the absence of noise.

Two possible sampling processes are examined here. Firstly, we consider area-sampling which is the process that models CCD data acquisition. To first order, the CCD array can be modelled as a square grid of light-sensitive elements with no spaces between them, with uniform sensitivity across each cell. This means that a low resolution sample is produced by averaging R^2 samples. On the other hand, in point-sampling, only every R^{th} sample is

Figure 5.2: *A CCD cell with dimensions $X_0 \times Y_0$. The sensitive area is of dimensions $a \times b$. In this case shown here, $a < X_0$ and $b < Y_0$; the sensor cells in the array are therefore not butted.*

taken in each direction. This is the assumption employed in most interpolation schemes. For example, assuming that the high resolution data is of size 8×8 and that the low resolution data is sampled at half the resolution of the high resolution data (i.e., $R = 2$), area-sampling is represented by

$$\mathbf{H}_a = \left(\begin{array}{cc} \left[\begin{array}{cc} \mathbf{B}_a & \mathbf{0} \\ \mathbf{0} & \mathbf{B}_a \end{array} \right] & \mathbf{0} \\ \mathbf{0} & \left[\begin{array}{cc} \mathbf{B}_a & \mathbf{0} \\ \mathbf{0} & \mathbf{B}_a \end{array} \right] \end{array} \right)$$

where each block $\mathbf{B}_a$ is

$$\mathbf{B}_a = \left(\begin{array}{cccccccccccccccc} \frac{1}{4} & \frac{1}{4} & 0 & 0 & 0 & 0 & 0 & 0 & \frac{1}{4} & \frac{1}{4} & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \frac{1}{4} & \frac{1}{4} & 0 & 0 & 0 & 0 & 0 & 0 & \frac{1}{4} & \frac{1}{4} & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \frac{1}{4} & \frac{1}{4} & 0 & 0 & 0 & 0 & 0 & 0 & \frac{1}{4} & \frac{1}{4} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \frac{1}{4} & \frac{1}{4} & 0 & 0 & 0 & 0 & 0 & 0 & \frac{1}{4} & \frac{1}{4} \end{array} \right) .$$

Similarly, for a situation with the same dimensions, point-sampling can be expressed as

$$\mathbf{H}_p = \left(\begin{array}{cccccccc} \mathbf{B}_p & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{B}_p & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{B}_p & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{B}_p & \mathbf{0} \end{array} \right)$$

where each block $\mathbf{B}_p$ is

$$\mathbf{B}_p = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{pmatrix}.$$

Examples of an area-sampled and a point-sampled image are shown in Figs. 5.1(b) and 5.1(c).

When the data-acquisition process of the image-capturing instrument is taken into account, Eq. (5.2) ensures that the high resolution samples would generate low resolution samples that are consistent with the data gathered.

5.3　Optimal interpolation

Following Eq. (5.2), the linear estimate $\hat{\mathbf{f}}$ of $\mathbf{f}$ can be defined as:

$$\hat{\mathbf{f}} = \Theta_{opt}\mathbf{d} \tag{5.3}$$

where Θ_{opt} is the optimal interpolator. The simplest solution is the least squares estimator [115] where

$$\Theta_{opt} = (\mathbf{H}^T\mathbf{H})^{-1}\mathbf{H}^T. \tag{5.4}$$

However, this solution assumes that no knowledge about the underlying high resolution object is available, which is not the case.

In order to utilise all the available information, the optimal interpolator is defined as the best linear unbiased estimator (BLUE) with minimum variance for the high-resolution data. The autocovariance for the high-resolution data is defined as $E\{\mathbf{ff}^T\} = \mathbf{K}_{ff}$, and the auto-covariance of the noise on the measured low-resolution data is $E\{\mathbf{nn}^T\} = \mathbf{K}_{nn}$, where $E\{\cdot\}$ denotes the expectation operator. The data and the noise are assumed to be uncorrelated, i.e., $E\{\mathbf{nf}^T\} = 0$. The optimal interpolator that minimises $E\{\|\hat{\mathbf{f}} - \mathbf{f}\|^2\}$ is given by [97,115]:

$$\Theta_{opt} = \mathbf{K}_{ff}\mathbf{H}^T(\mathbf{H}\mathbf{K}_{ff}\mathbf{H}^T + \mathbf{K}_{nn})^{-1} \tag{5.5}$$

which is essentially a solution to an underdetermined inverse problem, with the prior information expressed in terms of $\mathbf{K}_{ff}$ and $\mathbf{K}_{nn}$.

Equivalently, the criterion of optimality can be viewed as a Bayes estimation problem associated with Eq. (5.2) where we find $\hat{\mathbf{f}}$ such that the *a posteriori* probability is maximised. Applying Bayes rule, we have

$$\max_{\hat{\mathbf{f}}} P(\hat{\mathbf{f}}|\mathbf{d}) = \max_{\hat{\mathbf{f}}} \frac{P(\mathbf{d}|\hat{\mathbf{f}})P(\hat{\mathbf{f}})}{P(\mathbf{d})}. \tag{5.6}$$

Assuming that both $\hat{\mathbf{f}}$ and $\mathbf{n}$ are multivariate Gaussian processes, their respective probabilities are

$$P(\hat{\mathbf{f}}) = \exp\left\{-\frac{1}{2}\hat{\mathbf{f}}^T \mathbf{K}_{ff}^{-1} \hat{\mathbf{f}}\right\}$$
$$P(\mathbf{n}) = \exp\left\{-\frac{1}{2}(\mathbf{d} - \mathbf{H}\hat{\mathbf{f}})^T \mathbf{K}_{nn}^{-1}(\mathbf{d} - \mathbf{H}\hat{\mathbf{f}})\right\}. \tag{5.7}$$

Substituting Eq. (5.7) into Eq. (5.6), taking the log and discarding terms that are not a function of $\hat{\mathbf{f}}$, the expression we wish to maximise is

$$\ln P(\hat{\mathbf{f}}|\mathbf{d}) = -\frac{1}{2}(\mathbf{d} - \mathbf{H}\hat{\mathbf{f}})^T \mathbf{K}_{nn}^{-1}(\mathbf{d} - \mathbf{H}\hat{\mathbf{f}}) - \frac{1}{2}\hat{\mathbf{f}}^T \mathbf{K}_{ff}^{-1} \hat{\mathbf{f}}, \tag{5.8}$$

which again leads to Eq. (5.5) for the optimal interpolator [4].

There are a number of issues associated with the solution of Eq. (5.5). Firstly, for the inverse to exist, $(\mathbf{H}\mathbf{K}_{ff}\mathbf{H}^T + \mathbf{K}_{nn})$ must be of full rank. In practice, the invertibility is guaranteed because $\mathbf{K}_{ff}$ is a positive semidefinite autocovariance matrix and as a consequence $\mathbf{H}\mathbf{K}_{ff}\mathbf{H}^T$ is positive definite. Provided the noise on each CCD element is independent, the noise autocovariance $\mathbf{K}_{nn}$ takes the form of a diagonal matrix. This ensures the bracketed term in equation Eq. (5.5) is positive definite and easily invertible.

Secondly, the uniqueness of the solution can be shown by using Eq. (5.8). If we only consider the first quadratic term, the solution to the problem is not unique because $\mathbf{H}$ is not of full rank, and consequently multiple solutions exist. However, together with the prior term $\hat{\mathbf{f}}^T \mathbf{K}_{ff}^{-1} \hat{\mathbf{f}}$, uniqueness is ensured because the covariance $\mathbf{K}_{ff}$ is of full rank and positive semidefinite.

It should be noted that the low resolution data is not zero-mean (because it is captured on a CCD), while the autocovariance by definition is a zero-mean statistic. For interpolation this is not a problem, since in the initial model assumption, $\hat{\mathbf{f}}$ is assumed to be a stationary random process. This means that the vector $\hat{\mathbf{f}}$ is a set of random fluctuations about a constant mean. As a result, the interpolation is performed as though the data is zero-mean

plus a constant bias, and the interpolated image is thus a zero-mean process with an additive constant [4].

In the presence of noise, an exact fit of the interpolated values to the data is often undesirable. From Eq. (5.5), we see that the noise process is incorporated in the design of the estimator. In the extreme case when the sample is very noisy, the term $\mathbf{K}_{nn}$ dominates the solution. This leads to the seemingly surprising result that the optimal interpolator tends to zero as the noise increases. In practice, it means that the data is so noisy that any conventionally interpolated value is dominated by noise, and it is best not to interpolate at all. This demonstrates the balance between the weighting on the data and the prior information, a property common to many inverse problems. In this case, the noise may be such that the resolution of the existing samples should be reduced by smoothing.

Most image-capturing instruments do not meet the Nyquist sampling criterion. This means that different frequencies of $\mathbf{f}$ can be aliased to the same frequency in $\mathbf{d}$. However, different frequency components in the high resolution data correspond to different values in the autocovariance matrix $\mathbf{K}_{ff}$, and as a result, a meaningful separation of the aliased frequencies can be made by the optimal interpolator in Eq. (5.5).

In summary, the main advantages of the optimal interpolator over the existing techniques are as follows. The sampling process and the statistics of the data and the noise are modelled and incorporated in the interpolator. The explicit inclusion of the noise term in the interpolator means that a balance can be achieved between the fidelity of the interpolated image to the measured data, and the amount of noise present on the data. The optimal interpolator can also partially overcome the problem of aliasing, since although two frequencies may produce identical samples, they are not equally probable.

5.3.1 Application of the optimal interpolator to geostatistics

The optimal interpolator has wide applications, one of which is for interpolating geostatistical data. Geostatistics is an area in which a limitation on the available samples results in a need for interpolation. The existing technique for this is known as kriging [25, 88, 177]. Despite the popularity of kriging, it has two main disadvantages. Firstly, the sampled points are assumed to be noise-free, and collocation is used. As a result, no weighting is assigned

to the quality of the samples. Moreover, the invertibility of the solution is not guaranteed.

The second problem is that a different covariance matrix is calculated for each point of interpolation, and the process of kriging is slow. An improvement is to use a generic covariance model, such as the Markov model, to describe the underlying statistics of the data. In this way, the covariance matrix only needs to be calculated once. Although this may not be as exact as the model used in kriging, the proposed method will be much faster.

Finally, it should be noted that in kriging, the samples are usually irregular. However if regular samples were taken, the optimal interpolator can be applied directly to the data.

5.4   Autocovariances of the high-resolution data and the noise

The statistical interpolator relies on the knowledge of the autocovariance $\mathbf{K}_{ff}$ of the high resolution samples. Obviously, this information may not be available with real world data, but it is clear that certain assumptions are more realistic than others. For comparison, several typical autocovariance models are selected for interpolation. The performance of these models is compared with that of the actual autocovariance extracted from real experimental data in the next section.

Since the power spectrum is the Fourier transform of the autocovariance (assuming zero-mean data), several models of autocovariances can be selected according to the different assumptions in the frequency domain. A white spectrum which assumes that all frequencies are equally likely results in a delta function for the autocovariance:

$$K_{ff}^{\delta}(x,y) = \delta(x,y).$$

This is equivalent to the smoothness constraint in smoothing splines [138,152,173], where Eq. (5.8) is maximised with $\mathbf{K}_{ff}$ as a diagonal matrix. The sinc autocovariance corresponds to band-limited data with a uniform power spectrum and can be expressed as:

$$K_{ff}^{\text{sinc}}(x,y) = \text{sinc}\left(\frac{x}{R}\right) \text{sinc}\left(\frac{y}{R}\right).$$

The Gaussian autocovariance,

$$K_{ff}^{\text{Gaussian}}(x,y) = \exp\left\{ -\frac{(x^2+y^2)}{R^2} \right\},$$

(a)   (b)   (c)
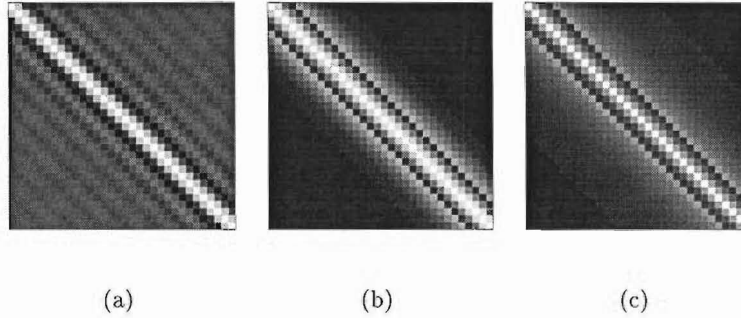
Figure 5.3: *Autocovariance matrices for the (a) sinc, (b) Gaussian, and
(c) Markov covariance models. The correlation c for the Markov model
is 0.5 here. The images are shown on a gray scale with the white pixels
equal to 1 and the black pixels equal to 0.*

models extreme smoothness and continuity, a characteristic of certain images [88]. However,
Goovaerts [63] suggests that the Gaussian model may suffer from numerical instabilities and
artifacts in interpolation, and recommends the use of a Markov model:

$$K_{ff}^{\text{Markov}}(x,y) = c^{\sqrt{(x^2+y^2)}},$$

where c is the correlation between adjacent pixels. The autocovariance matrices $\mathbf{K}_{ff}$ are
constructed from these basic representations. Fig. 5.3 shows the matrices for the sinc,
Gaussian and Markov autocovariances. The value of c is set to 0.5. The delta autocovariance
matrix is not shown because it is simply an identity matrix.

The autocovariance of the noise is dependent on the correlation between the noise samples.
According to Andrews and Hunt [4], there are two noise processes associated with image
detection with a photoelectronic system: random thermal noise sources in the electronic
circuits, and random fluctuations of the number of photons on the detector surface. The
former is adequately described by a zero-mean Gaussian process with a white spectrum.
The noise emissions of the latter are independent of spatial position on the photoactive
surface, meaning that the corresponding noise process is spatially uncorrelated with a white
spectrum. Therefore it is assumed that the overall detection noise is uncorrelated.

The autocovariance of the noise is also a function of the variance of the quantisation process
used in digital processing. In the case of uniform quantisation, it can be assumed that the
CCD sampler has 8 bits, and the images are captured on a gray scale with 256 levels. This
leads to a quantisation step size Δ of 1, with the variance of the noise being $\Delta^2/12$ [78].

Since the noise is uncorrelated, the autocovariance is

$$\mathbf{K}_{nn} = \frac{\Delta^2}{12}\mathbf{I},$$

where $\mathbf{I}$ is the identity matrix. The actual noise variance is likely to be higher from the contribution of sensor noise. On the other hand, a non-uniform quantisation procedure results in image-dependent noise. In this case, the variance is larger at locations with more photons, and smaller at locations with fewer photons. This means that the optimal interpolator can no longer be applied globally to the data; instead, a different interpolator is required for each point of interpolation. This increases the computational complexity substantially. However, non-uniform quantisation is usually only found in CMOS cameras, and the assumption of linear quantisation is valid in most cases.

5.5 Experimental results

This section presents experimental results obtained using two sets of real data: images acquired by a commercial scanner, and images taken using a SPOT remote-sensing orbiting satellite [20]. These data sets are chosen because the images are available at two resolutions, providing a ground truth for assessing the performance of the interpolators. Fig. 5.4(a) and (b) show the scanned images at two different resolutions, while Fig. 5.4(c) and (d) show the SPOT images. The optimal interpolator is directly applicable to the data acquired from the scanner, whereas the SPOT data requires pre-processing to overcome practical problems such as spectral mismatch and the misalignment of the low- and high-resolution data. These issues have been dealt with in the paper by Bretschneider *et al.* [15]; here the investigation is confined to the application of the optimal interpolator to the data.

The results obtained using the proposed interpolator are compared to the performance of the cubic spline interpolator [117, 172] and the bilinear interpolator as these are the most commonly used interpolators. Before the comparison can be made, it is necessary to generate the interpolating kernels under different noise, autocovariance and sampling assumptions. This is addressed in the following subsection.

(a)

(b)

(c)

(d)
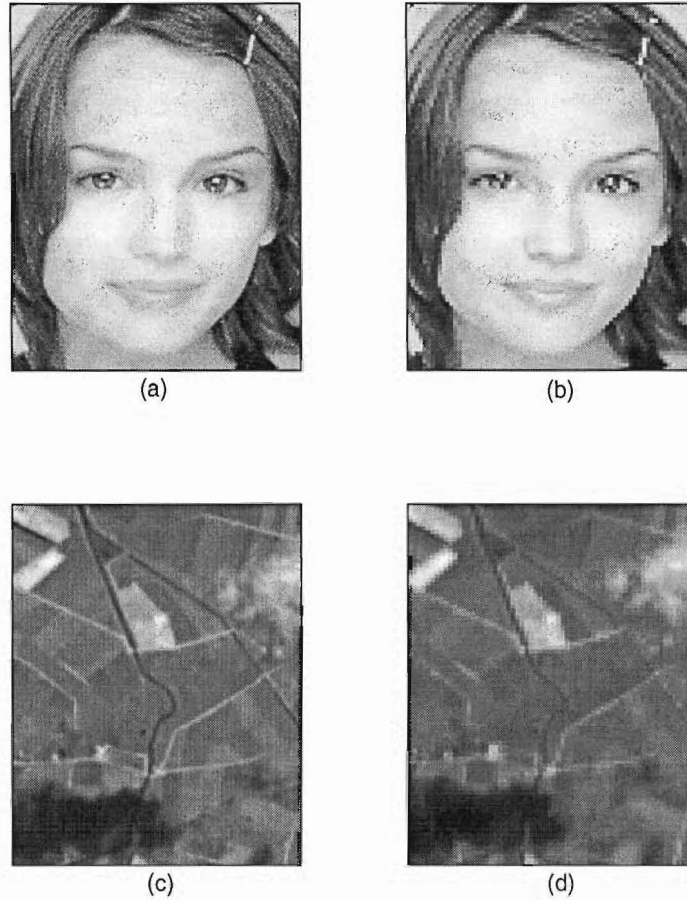
Figure 5.4: *The two sets of data used to evaluate the performance of the optimal interpolator. (a,b) High and low resolution data from the scanner, (c,d) High and low resolution SPOT farm data.*

5.5.1   Interpolating kernels under different assumptions

The characteristics of the interpolating kernels formed from the different sampling, autocovariance and noise assumptions are examined. Fig. 5.5(a) shows the high-resolution portrait (192×128) acquired with a commercial scanner. This is the same as Fig. 5.4(a) and is shown here for completeness. A 32×32 section is extracted (boxed area) and the cross-section of the autocovariance calculated from this section is shown in Fig. 5.5(b). Using the actual autocovariance from Fig. 5.5(b), and modelled autocovariances, the cross-section of the interpolating kernels formed are shown in Fig. 5.6. The left- and right-hand columns contain kernels produced by an area-sampling and point-sampling assumption respectively.
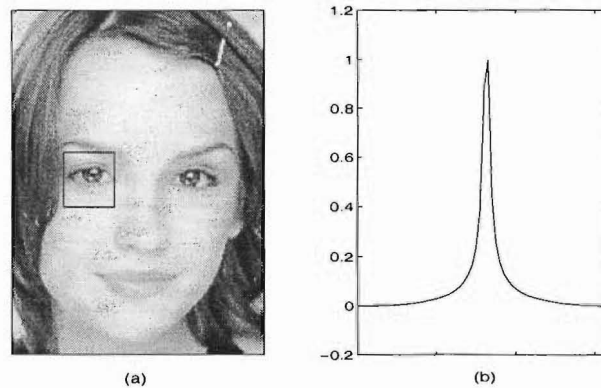
Figure 5.5: *Portrait data from which the autocovariance is extracted for comparison with the modelled autocovariances in Fig. 5.6. (a) Large high-resolution data. The boxed area is the extracted section, size 32x32. (b) Cross-section of the autocovariance calculated from the boxed area, plotted over 64 points.*

The solid lines represent the kernels obtained when no noise is present and the dash-dotted lines are kernels formed in the presence of 8-bit quantisation noise. Fig. 5.6(a) and (b) show the kernels generated using the actual autocovariance of the scanned high resolution data. It can be seen that under a noiseless assumption, the Markov kernels in (c) and (d) are the best model of the true kernels. When noise is present, the Markov kernels tend to spread out more than the true kernels, thus attenuating high frequencies to a larger extent than the true kernels in the reconstruction. Nevertheless, it should be noted that this spreading behaviour does not apply to all kernels; for example, the noisy and noiseless kernels are almost identical for the delta point-sampled model (Fig. 5.6(j)). Therefore, whether the noise has a significant effect on the interpolator depends on the underlying statistics of the image.

It is interesting to note that variations of the statistical interpolator approach classical interpolators. Three of the shown point-sampling interpolators (Fig. 5.6(b), (d) and (j)) are similar to the bilinear interpolator, with the delta-point interpolator being identical to it. On the other hand, the white point-sampling kernel (Fig. 5.6(h)) is a sinc interpolator. The bilinear interpolator is used extensively for its computational convenience, but it can also be justified statistically under a wide range of circumstances. By contrast, the sinc interpolator corresponds to interpolation according to the Nyquist criterion, but in practice does not perform well as an interpolator because it introduces ringing in the reconstructed signal. This is not unexpected statistically because the assumptions that lead to this interpolator
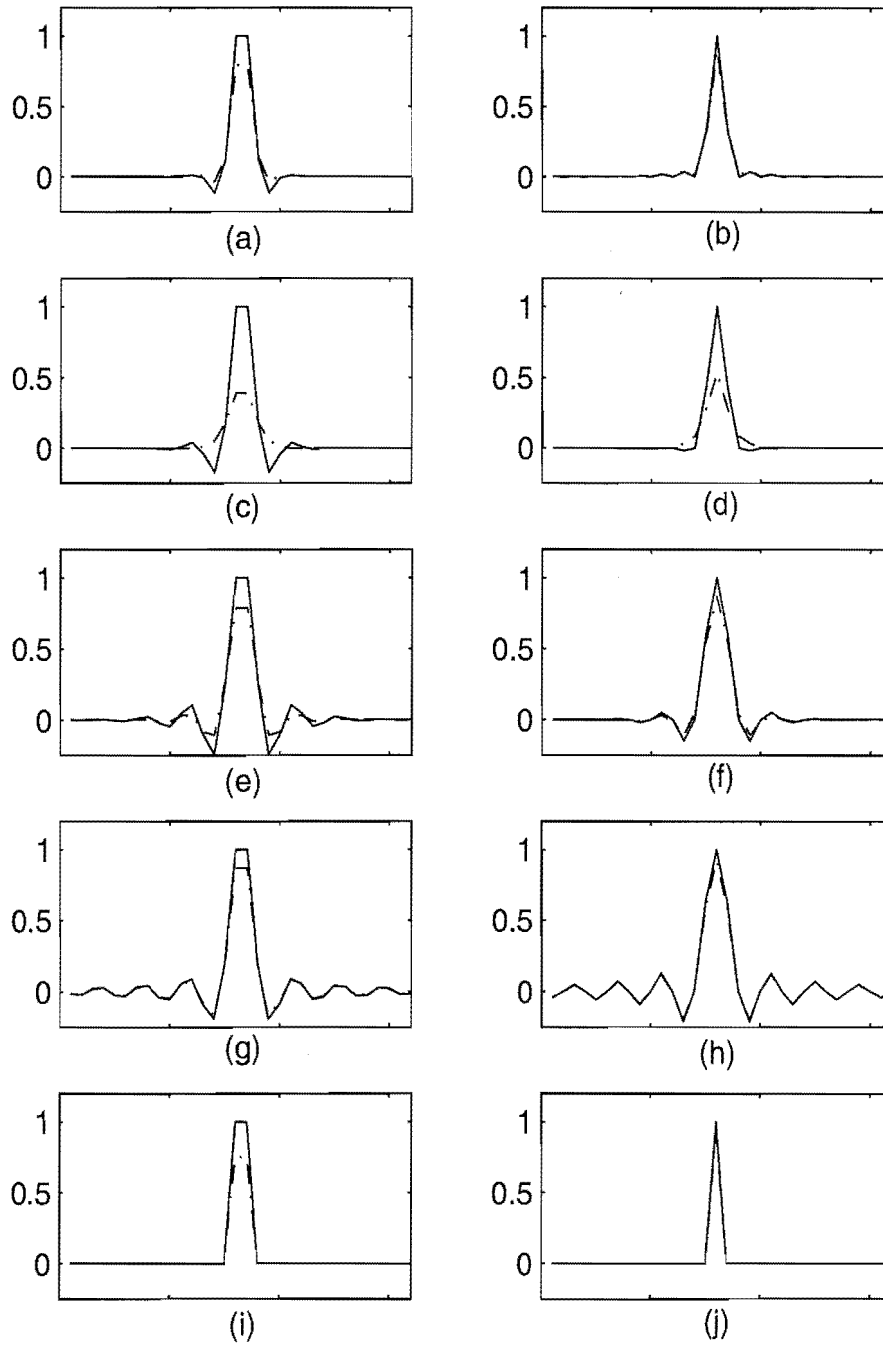
Figure 5.6: *Interpolating kernels under different assumptions. Only the cross-sections are shown here. The solid lines are formed assuming no noise is present, and the dashdot lines are for the assumption of 8-bit quantisation noise. All kernels are plotted over 32 points. (a) Autocovariance from actual high-resolution data, area sampling, (b) Autocovariance from actual high-resolution data, point sampling, (c) Markov area, (d) Markov point, (e) Gaussian area, (f) Gaussian point, (g) White area, (h) White point, (i) Delta area, (j) Delta point.*

are unrealistic and rarely found in images despite its use as a standard interpolator. In summary, classical interpolators can be viewed within the framework of statistical interpolators, and are optimal under particular sampling and statistical assumptions.

It should be noted that in Fig. 5.6, all the interpolators are represented as convolutional kernels. The impulse response of the spline function, the cardinal spline, can also be viewed as a convolution. This means that prior information and the sampling-process model can be incorporated into spline interpolators. In particular, smoothing splines can be incorporated in the framework presented here [172].

5.5.2   Ensemble results

The two data sets are shown in Fig. 5.4. Each of the high-resolution images (Fig. 5.4(a) and (c)) is of size 192×128, and is divided into twenty-four 32×32 images. Correspondingly, the low-resolution images (Fig. 5.4(b) and (d)) are of half the resolution of the high-resolution images (i.e. $R=2$, size 96×64), and are each divided into twenty-four 16×16 images. As a result, for each set of data, there are twenty-four sub-images to be interpolated and compared. This ensures that the statistics between the sub-images of each set of data are very similar.

Each low-resolution sub-image is interpolated by the optimal interpolator under different noise and covariance assumptions, and the normalised MSE is calculated with 8 pixels removed from the edges to eliminate the influence of edge effects. The correlation c between adjacent pixels in the Markov model was estimated from the low-resolution data to be 0.5. The ensemble results for the portrait and SPOT farm images are tabulated in Tables 5.1 and 5.2 respectively. The MSE and the variance for the cubic-spline interpolated result are 2.35×10^{-5} and 1.59×10^{-9} respectively for the portrait. For the SPOT farm image, the corresponding values are 8.41×10^{-6} and 4.04×10^{-11}.

From these results, several conclusions can be drawn. In all cases, the area-based interpolators perform better than their point-based counterparts, confirming that better modelling of the data-acquisition process improves the results. The best results are produced by the extracted-area and Markov-area interpolators, which are essentially comparable in terms of mean-squared error. However, the extracted-point interpolator does not perform as well

Kernel assumption	Noise variance			
	0	0.04	0.0833	0.16
extracted, area	1.94×10^{-5}	1.82×10^{-5}	1.76×10^{-5}	1.74×10^{-5}
extracted, point	8.77×10^{-5}	8.74×10^{-5}	8.74×10^{-5}	8.72×10^{-5}
Markov, area	1.95×10^{-5}	1.81×10^{-5}	1.74×10^{-5}	1.70×10^{-5}
Markov, point	2.76×10^{-5}	2.75×10^{-5}	2.74×10^{-5}	2.74×10^{-5}
sinc, area	2.56×10^{-5}	2.56×10^{-5}	2.55×10^{-5}	2.55×10^{-5}
sinc, point	8.32×10^{-4}	8.32×10^{-4}	8.32×10^{-4}	8.32×10^{-4}
white, area	2.43×10^{-5}	2.39×10^{-5}	2.37×10^{-5}	2.40×10^{-5}
white, point	6.63×10^{-3}	6.19×10^{-3}	6.85×10^{-3}	6.19×10^{-3}
Gaussian, area	2.12×10^{-5}	1.93×10^{-5}	1.85×10^{-5}	1.79×10^{-5}
Gaussian, point	2.42×10^{-5}	2.37×10^{-5}	2.34×10^{-5}	2.30×10^{-5}

Table 5.1: *Mean squared error values for the portrait data set. The variances range from 7.64×10^{-10} to 3.72×10^{-6}. The MSE for the cubic-spline interpolated result is 2.35×10^{-5} with a variance of 1.59×10^{-9}. The noise variance of 0.0833 corresponds to 8-bit quantisation noise.*

Kernel assumption	Noise variance			
	0	0.04	0.0833	0.16
extracted, area	5.64×10^{-6}	6.28×10^{-6}	6.33×10^{-6}	5.42×10^{-6}
extracted, point	1.83×10^{-4}	1.83×10^{-4}	1.83×10^{-4}	1.83×10^{-4}
Markov, area	6.15×10^{-6}	6.18×10^{-6}	6.24×10^{-6}	6.34×10^{-6}
Markov, point	8.88×10^{-6}	8.91×10^{-6}	8.94×10^{-6}	8.99×10^{-6}
sinc, area	6.43×10^{-6}	6.50×10^{-6}	6.44×10^{-6}	6.45×10^{-6}
sinc, point	6.83×10^{-4}	6.83×10^{-4}	6.83×10^{-4}	6.83×10^{-4}
white, area	7.07×10^{-6}	7.08×10^{-6}	7.08×10^{-6}	7.08×10^{-6}
white, point	5.58×10^{-3}	5.58×10^{-3}	5.58×10^{-3}	5.58×10^{-3}
Gaussian, area	6.28×10^{-6}	6.25×10^{-6}	6.25×10^{-6}	6.27×10^{-6}
Gaussian, point	6.28×10^{-6}	6.90×10^{-6}	6.91×10^{-6}	6.93×10^{-6}

Table 5.2: *Mean squared error values for the SPOT farm data set. The variances range from 1.47×10^{-11} to 3.69×10^{-8}. The MSE for the cubic-spline interpolated result is 8.41×10^{-6} with a variance of 4.04×10^{-11}. The noise variance of 0.0833 corresponds to 8-bit quantisation noise.*

as the Markov-point interpolator. This is probably because by using the one-parameter Markov model, we have emphasised the low-frequency components in the reconstruction, thus excessively smoothing the solution. This appears to have compensated for the incorrect data-acquisition process used, and has led to a somewhat better reconstruction than the extracted-point result in terms of MSE. The Gaussian autocovariance model also exhibits this smoothing property, demonstrated by the comparable results obtained using both the area and point interpolators. The white and sinc interpolators performed poorly due to the unrealistic autocovariance assumptions discussed earlier.

Four noise variance assumptions have been investigated; the case where no noise is present while unrealistic is included for completeness. While 8-bit quantisation noise is a realistic minimum, it is appropriate to assume a higher noise variance to allow for the contribution of sensor noise. The best optimal interpolators (extracted-area and Markov-area) are consistently better than the cubic spline interpolator, confirming the benefits of the Bayesian approach.

5.5.3 Superresolution

Since interpolation can be viewed as being complementary to superresolution, the frequency content of an interpolated image should be examined. The high-resolution scanner data is used again, but the ratio R between the resolutions of the high and low resolution data is set to 5. Only one-dimensional interpolation is considered for the convenience of displaying the ensemble average spectrum. The extracted-area, Markov-area and white-point (conventional sinc) interpolators are used to interpolate the data. The results are shown in Fig. 5.7. In Fig. 5.7(a), the ensemble spectra of the data of both resolutions are shown. The sampling frequency of the high resolution data, f_s, is 5 times that of the corresponding low resolution data. The spectral peak at 0.125 f_s in the original spectrum is seen to have aliased into the low resolution spectrum at 0.075 f_s. In Fig. 5.7(b), the extracted-area, Markov-area and white-point reconstructions are shown as dashed, solid and dash-dotted lines respectively. The spectral peak at 0.125 f_s is seen to be partially restored by the extracted-area and the Markov-area interpolators. The meaningful frequency content of the interpolated signal has also been extended, illustrating the phenomenon of superresolution. On the other hand, the reconstruction from the use of the white-point kernel (conventional sinc interpolator) has not restored the high resolution spectrum correctly.
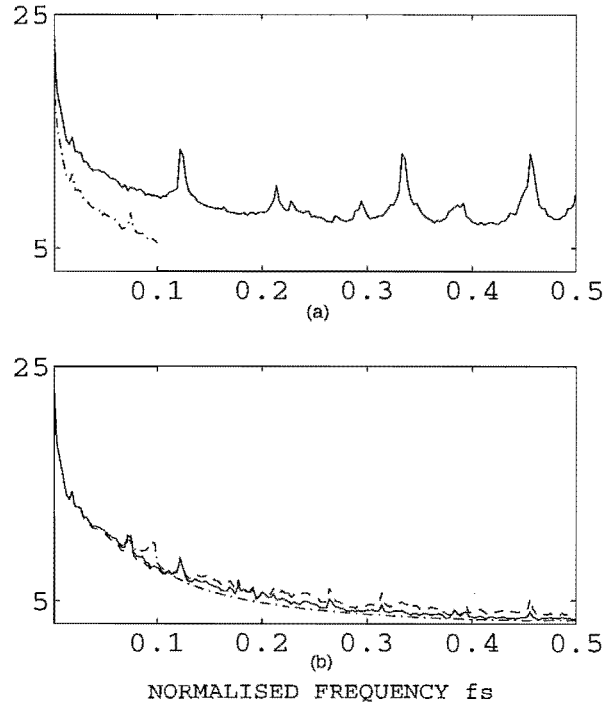
Figure 5.7: *Ensemble spectra of the scanned and interpolated data. All the spectra have a logarithmic vertical scale. (a) Ensemble spectra of the scanned data. The high and low resolution spectra are represented by solid and dash-dotted lines respectively. The spectral peak at 0.125 f_s has aliased into the low resolution spectrum at 0.075 f_s. (b) The solid line is the interpolation using a Markov-area kernel. The peak at 0.125 f_s is partially restored. A similar result is obtained using the extracted-area kernel shown as the dashed line. The dash-dotted line is the reconstruction using a white-point kernel. This reconstruction is less representative of the true high resolution data.*

From this experiment, it is shown that provided the sampling and statistical assumptions are representative of the data, the statistical interpolator is capable of partially restoring aliased sampled signals and achieving superresolution.

5.6   Summary

It has been demonstrated that a framework of optimal interpolators is derivable from a Bayesian statistical basis, and provides a direct solution to the inverse problem of interpolation. In contrast to traditional interpolators, the sampling process is included in the

system model and prior information in the form of autocovariances is included. The solution presented is optimal in terms of linearity, unbiasedness and variance minimisation.

Classical interpolators such as the bilinear and the sinc functions are shown to be derivable in the proposed framework. Thus interpolators usually chosen for computational simplicity such as the bilinear may in fact be optimal under certain assumptions. Finally, the problem of aliasing is dealt with directly using this Bayesian framework. The phenomenon of superresolution, where the meaningful frequency content of the reconstruction is extended beyond the instrument cutoff, is also demonstrated.

The application of the optimal interpolator is not restricted to images; for example, the kriging technique in geostatistics will also benefit from the proposed method.

An extension to the proposed interpolator in terms of resolution improvement is to model the CCD cell sensitivity and to utilise multiple frames of data. One way of implementing this is to use an iterative blind deconvolution framework and is the topic of discussion in Chapter 7.

Chapter 6

Blind deconvolution: review and feasibility study

In astronomical imaging, the ultimate goal is to reconstruct the object distribution. Since the principal degradation introduced by the atmosphere on an object is a phase distortion, this must be estimated and compensated for. Depending on the data captured, the phase or equivalently the point spread function is estimated using different approaches.

When only speckle images are available, iterative blind deconvolution is used to solve the problem. On the other hand, if wavefront sensing data is captured simultaneously, then myopic deconvolution and hybrid methods such as deconvolution from wavefront sensing (DWFS) can be applied. The focus of this chapter is the technique of blind deconvolution; the simultaneous application of wavefront sensing data in object reconstruction is addressed in Chapter 8, and the technique of DWFS is examined in Chapter 9.

6.1 Introduction

In an astronomical blind deconvolution problem, we wish to find both the object and the point spread function (PSF) when only their convolution i.e. the speckle image is measured. This is an ill-posed inverse problem, and as a result, the feasibility and the quality of the

solution depend heavily on the form and amount of prior information available.

There are three areas where prior information can be included: the noise statistics, the object, and the PSF. When the only prior information available is the noise statistics, the blind deconvolution problem is one of maximum likelihood (ML). In this case, the object and the PSF are reconstructed to best fit the data according to the assumed noise model, with the two most common choices being the Poisson and Gaussian models. Unfortunately, knowledge of the noise model alone is insufficient to produce a meaningful solution, because the likelihood is always maximised by a trivial solution where the object is estimated to be a delta function, and the PSF the data (or vice versa). This implies that the situation where no noise is present on the data is the most likely solution, which contradicts the reality that noise is always present in a system. It is therefore essential to include prior information about the object and the PSF to obtain a meaningful reconstruction.

Although it is generally true that in solving such an inverse problem, it is desirable to incorporate as much information into the framework as possible, one important issue that must be addressed is the tradeoff between the potential accuracy of the solution and the ease of which it can be obtained. If a solution is possibly very accurate but the model of the system is very complex, then the gain of the accuracy must be weighed against the computational complexity and more importantly whether a solution can be found. This is the reason behind the popularity of linear models, which may lack accuracy but guarantee a unique solution and facilitate the analysis of the algorithm's behaviour. More complicated models usually have multimodal cost functions, meaning that stagnation of the algorithm in sub-optimal solutions is possible.

In this chapter, iterative blind deconvolution (IBD) of astronomical images is reviewed, and two forms of prior information are investigated. Firstly, since the PSF is a function of the atmospheric turbulence, it is reasonable to constrain the reconstructed PSF to obey the statistics of the atmosphere. Secondly, the Knox-Thompson (KT) phase which is a subset of the bispectrum phase is used to constrain the PSF. This is considered because the bispectrum technique alone is able to reconstruct the object phase by the phase closure property, where the atmospheric phase contribution is eliminated while the phase information of the object is preserved [7]. The combination of the phase closure property with blind deconvolution has the potential to produce reconstructions of good quality.

It should be pointed out that studies performed in this chapter are only used to ascertain the feasibility of these approaches. A more comprehensive investigation is required to completely characterise the performance of these approaches.

6.2 Astronomical setting and mathematical preliminaries

The wavefront of a star propagates undisturbed for lightyears through the vacuum of space, until it reaches the Earth's atmosphere where phase distortions are introduced. When the wavefront subsequently propagates to ground, the principal effect is one of phase distortion, but its intensity is also distorted with a severity that depends on the distance of propagation [140]. The intensity image observed at the image plane of the telescope is modelled by

$$
\begin{aligned}
d_k(x,y) &= f(x,y) \odot h_k(x,y) + n_k(x,y) \\
&= g_k(x,y) + n_k(x,y)
\end{aligned}
\tag{6.1}
$$

where $f(x,y)$ is the stellar object of interest, $h_k(x,y)$ is a space-invariant point spread function (PSF) which accounts for the effects of the atmosphere, $g_k(x,y)$ is the noise-free image, $n_k(x,y)$ is the noise, and $d_k(x,y)$ is a short exposure image. The subscript k denotes the k^{th} realisation of the respective quantities in the K frames captured. The objective of blind deconvolution is to find both $f(x,y)$ and $h_k(x,y)$ when only $d_k(x,y)$ is available. Fig. 6.1 displays the relevant quantities with a simulated binary star and atmospheric turbulence with $D/r_0 = 5$.

Many methods have been proposed for solving the problem of astronomical blind deconvolution. The majority of these methods employ a maximum *a posteriori* (MAP) framework, and differ mainly in the form of prior information used. The MAP framework is presented in the following subsection, and the common forms of prior information are discussed.

The remainder of the chapter is structured as follows. Section 6.3 gives an overview of the existing algorithms. The application of Kolmogorov statistics as prior information is investigated in Section 6.4, while using the bispectrum phase as a constraint is examined in Section 6.5. Conclusions are drawn in Section 6.6.
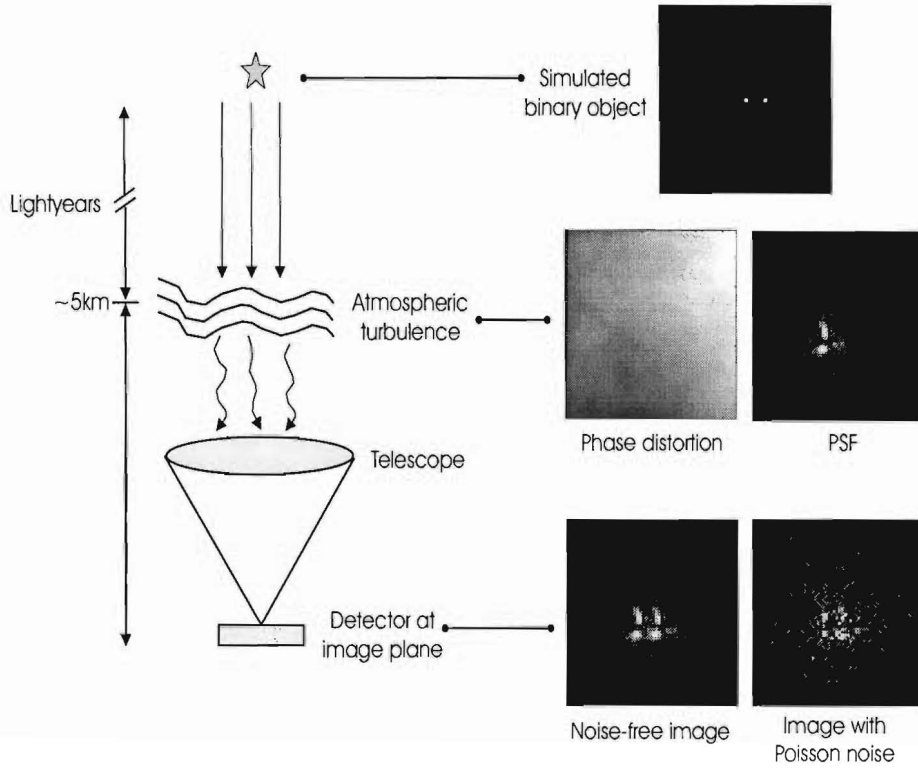
Figure 6.1: *The relevant quantities in the model for blind deconvolution.*

6.2.1 General solution in a MAP framework

The MAP solution to blind deconvolution is denoted by [97]

$$\max_{\hat{f},\hat{h}_k} P(\hat{f}(x,y),\hat{h}_k(x,y)|d_k(x,y)) \tag{6.2}$$

where the *a posteriori* probability is maximised. By applying Bayes' rule, taking the logarithms and discarding terms that are not a function of $\hat{f}(x,y)$ or $\hat{h}_k(x,y)$ result in the log-likelihood equation

$$\begin{aligned}
\max_{\hat{f},\hat{h}_k} \log[P(\hat{f}(x,y),\hat{h}_k(x,y)|d_k(x,y))] &= \max_{\hat{f},\hat{h}_k} \log[P(d_k(x,y)|\hat{f}(x,y),\hat{h}_k(x,y))] \\
&\quad + \log[P(\hat{f}(x,y))] \\
&\quad + \log[P(\hat{h}_k(x,y))].
\end{aligned} \tag{6.3}$$

The blind deconvolution solution is found by maximising this quantity.

The first term of Eq. (6.3) is the consistency measure. It ensures that the estimates are

consistent with the data and directly relates to the noise model assumption. The second and third terms describe the likelihoods of the object and the PSF according to the presumed statistical models or prior information. A balance must be achieved between these three terms for a good solution.

In photoelectronic systems for image detection, the noise depends on two processes: the random fluctuations of the number of photons on the detector surface, and the random thermal sources in the circuits. Depending on the light level, there are two common noise models that encompass these contributing factors.

For high light levels, the Gaussian noise model sufficiently describes the noise process, and the noise is independent of the signal. In this case, the first term of Eq. (6.3) corresponds to

$$E_{\text{con-Gaussian}} = \frac{1}{2} \sum_k \sum_{x,y} \frac{[d_k(x,y) - \hat{g}(x,y)]^2}{\sigma_k^2} \tag{6.4}$$

where

$$\hat{g}_k(x,y) = \hat{f}(x,y) \odot \hat{h}_k(x,y) \tag{6.5}$$

and σ_k^2 is the noise variance which is assumed constant over the image. In order to update $\hat{f}(x,y)$ and $\hat{h}_k(x,y)$, it is necessary to calculate the gradient of $E_{\text{con-Gaussian}}$ with respect to these quantities:

$$\begin{aligned}
\frac{\partial E_{\text{con-Gaussian}}}{\partial \hat{f}(x,y)} &= -\sum_k \frac{d_k(x,y) - \hat{g}_k(x,y)}{\sigma_k^2} \hat{h}_k^*(x,y) \\
\frac{\partial E_{\text{con-Gaussian}}}{\partial \hat{h}_k(x,y)} &= -\frac{d_k(x,y) - \hat{g}_k(x,y)}{\sigma_k^2} \hat{f}^*(x,y).
\end{aligned} \tag{6.6}$$

The object and PSF estimates can be updated using the method of steepest descent or conjugate gradient, the details of which are included in Appendix A.

Under low light conditions, it is more appropriate to use the Poisson noise model where the noise is dependent on the signal. The corresponding error term is [158]:

$$E_{\text{con-Poisson}} = \sum_k \sum_{x,y} d_k(x,y) \log[\hat{g}_k(x,y)] - \hat{g}_k(x,y). \tag{6.7}$$

The update of the quantities $\hat{f}(x,y)$ and $\hat{h}_k(x,y)$ is achieved by the formulation first proposed by Richardson and Lucy [114,139]. The object estimate is updated by:

$$\hat{f}^{q+1}(x,y) = \hat{f}^q(x,y) \left(\sum_k \frac{d_k(x,y)}{\hat{g}_k^q(x,y)} \otimes \hat{h}_k^q(x,y) \right) \tag{6.8}$$

where q denotes the iteration. This is the multiplicative form of the Richardson-Lucy (RL) update. It is important to note that because the new object estimate is a multiplication of the old object estimate with the update (the bracketed term in Eq. (6.8)), the use of a positive starting estimate ensures object positivity for subsequent iterations.

For many optimisation problems, it is convenient to express the update to the object estimate in terms of its effect on the error metric, or a gradient. This allows standard optimisation techniques such as steepest descent and conjugate gradient to be used, as well as providing a means for comparing different algorithms.

The additive form of the RL update is derived as follows:

$$
\begin{aligned}
\hat{f}^{q+1}(x,y) - \hat{f}^q(x,y) &= \hat{f}^q(x,y)\left(\sum_k \frac{d_k(x,y)}{\hat{g}_k^q(x,y)} \otimes \hat{h}_k^q(x,y)\right) - \hat{f}^q(x,y) \\
&= \hat{f}^q(x,y)\left(\sum_k \frac{d_k(x,y)}{\hat{g}_k^q(x,y)} \otimes \hat{h}_k^q(x,y) - 1\right) \\
&= \hat{f}^q(x,y)\left(\sum_k \frac{d_k(x,y) - \hat{g}_k^q(x,y)}{\hat{g}_k^q(x,y)} \otimes \hat{h}_k^q(x,y)\right).
\end{aligned}
\tag{6.9}
$$

In other words, the object estimate can be updated by:

$$
\hat{f}^{q+1}(x,y) = \hat{f}^q(x,y) + \hat{f}^q(x,y)\left(\sum_k \frac{d_k(x,y) - \hat{g}_k^q(x,y)}{\hat{g}_k^q(x,y)} \otimes \hat{h}_k^q(x,y)\right)
\tag{6.10}
$$

where the bracketed term is $\partial E_{\text{con-Poisson}}/\partial \hat{f}^q(x,y)$. The RL update is essentially a steepest descent search with a fixed step size and a search direction modified by $\hat{f}^q(x,y)$. The term $\hat{f}^q(x,y)$ has been interpreted by several authors as a preconditioner [23,26]. However, this is not true because the convergence of the solution is not always improved by its use [93]. Instead, the speed of convergence can be accelerated by increasing the standard RL step size from 1 and using conjugate gradient minimisation [10,11,80,96].

It should be pointed out that the gradient in Eq. (6.10) is very similar to those in Eq. (6.6). The difference lies in the value of the denominator, but in both cases the function of the denominator can be viewed as representing the assumed variance of the noise. In the Gaussian model, this is independent of $\hat{f}(x,y)$ and $\hat{h}(x,y)$, while in the Poisson model it is equal to the convolution of $\hat{f}(x,y)$ and $\hat{h}(x,y)$. This nonlinearity makes the optimisation of $\hat{f}(x,y)$ and $\hat{h}(x,y)$ more difficult for the Poisson model. However, the Poisson model does have an advantage despite the nonlinearity as it produces a photometrically linear restoration over a dynamic range of at least six stellar magnitudes [157].

Nevertheless, the Gaussian noise model is by far the more common of the two models. This is because the Poisson noise model is very sensitive to errors. For example, a small offset in the background level or the presence of any Gaussian noise can only be explained by the Poisson model as the presence of a signal. This can lead to errors in the reconstruction [106]. By contrast, the Gaussian noise model is more robust to the fluctuations in the signal. As a result, the Gaussian noise model is more popular in practice than the Poisson noise model. A combination of the Gaussian and Poisson noise models is also used in some algorithms [52, 106].

6.2.2 Astronomical object prior information

There exist many possible forms of prior information for an astronomical object. However, in order to maximise the scientific value of the data, it is important to minimise the number of assumptions made about the object. Since the light from an astronomical image is incoherent and cannot be negative, it is usual to constrain the object to be positive. It is also reasonable to constrain the solution to conserve light because the number of photons can be measured. If it is assumed that the object is unchanging, the temporal variation of the atmosphere can be used to obtain an ensemble of differently blurred images. This is called the multi-frame constraint. When only one frame is available, i.e. $k = 1$, the symmetry of Eq. (6.1) with respect to the object and the PSF raises a possibility that the reconstructions of $f(x, y)$ and $h_k(x, y)$ are interchanged. A multi-frame constraint eliminates this problem, as well as significantly increasing the amount of information available to estimate a solution. The effect of noise due to low light levels in short exposure images is also reduced by using multiple frames of data [140]. However, the trivial solution is still a possibility for the object.

A support constraint can be applied when the extent of the object is confined to a certain area. Although this provides theoretical justification for the uniqueness of the solution, there are difficulties in applying this constraint to diffuse extended objects because they do not have well-defined edges.

Other possible choices of prior information include the maximum entropy prior [129] (discussed in Chapter 3) and many ad hoc priors, which are primarily used to ensure the uniqueness of the solution. However, the extent of the effect of these priors may not be

obvious. For example, the maximum entropy prior does not only enforce positivity, but also controls the peak widths and the flatness of the baseline. Caution therefore must be exercised when using these priors.

6.2.3   Astronomical PSF prior information

Some of the constraints imposed on the object are also applicable to the PSF. These include the positivity, conservation of light and multi-frame constraints. Similarly, a support constraint is difficult to apply to an astronomical PSF because the extent of a speckle image cannot be determined easily. However, if a support constraint must be used, it should be determined from the SNR of the data, meaning that only data with a level above the noise floor should be inside the support.

The fact that the optical instrument has a finite aperture imposes a bandlimit constraint on the optical transfer function (OTF) [90]. This is a linear constraint and is easily implemented. The process of incoherent imaging which is inherent to astronomical imaging forces the PSF to be an inverse Fourier transform of an autocorrelation [154]. This enforces positivity of the PSF, as well as a bandlimit and conjugate symmetry of the OTF.

The incoherent imaging constraint is made even more powerful by the knowledge that the PSF is produced by an optical instrument aberrated by atmospheric turbulence. This is used in two of the most successful algorithms [90,154]. However, the gain in the power of constraining the solution is accompanied by an increase in nonlinearity and computational complexity.

Other forms of prior information on the PSF include the MEM and anisotropic regularisation (as discussed in Chapter 3). Again, it should be noted that there is a danger in making physically unjustifiable assumptions, as errors in the PSF are reflected in errors in the estimated object.

6.2.4   Prior information for partially corrected systems

The prior information discussed in the previous two subsections is used commonly for iterative blind deconvolution algorithms where only a minimal amount of information is known about the object and the PSF.

In an adaptive optics (AO) system, both the object and the PSF are compensated in real time. However, due to the limitations of the AO equipment such as the degrees of freedom in the deformable mirror and the number of actuators, as well as the time-evolving nature of the atmosphere, the compensation is not exact and a residual error remains. As a result, the object and the PSF are only partially compensated. It is therefore usual to apply post-processing techniques to improve the image quality. It is common to employ blind deconvolution algorithms for this purpose, but since very good starting estimates and prior information are available, this procedure is termed myopic deconvolution [24, 27, 51, 126]. The prior information used includes the power spectral density of the respective quantities, as well as parameterisation of the object. These are discussed in further detail in Section 6.3.2 and Chapter 8.

6.3   Review of existing blind deconvolution algorithms

Since the advocacy of iterative blind deconvolution by Ayers and Dainty in 1988 [6], this problem has been addressed by many authors in the context of both general and astronomical blind deconvolution: Kundur and Hatzinakos [102], You and Kaveh [184], Chan and Wong [21], Yang *et al.* [183], Fish *et al.* [36], McCallum [120], Schulz [154], Conan *et al.* [27], Fusco *et al.* [51], Mugnier *et al.* [126], Christou *et al.* [24], Lane [105], Law and Lane [110], Thiebaut and Conan [168], Jefferies and Christou [90], Sheppard *et al.* [157], Miura *et al.* [123], and many others. These authors all adopt the approach of applying prior information; the difference between the methods is primarily the different forms of prior information used. Table 6.1 summarises the assumptions made by these authors in terms of the noise model, the object and the point spread function. The relative ease of implementation of the commonly used priors and their ability to constrain the solution are illustrated in Fig. 6.2.

Although not all the algorithms are specifically proposed for astronomical applications, they
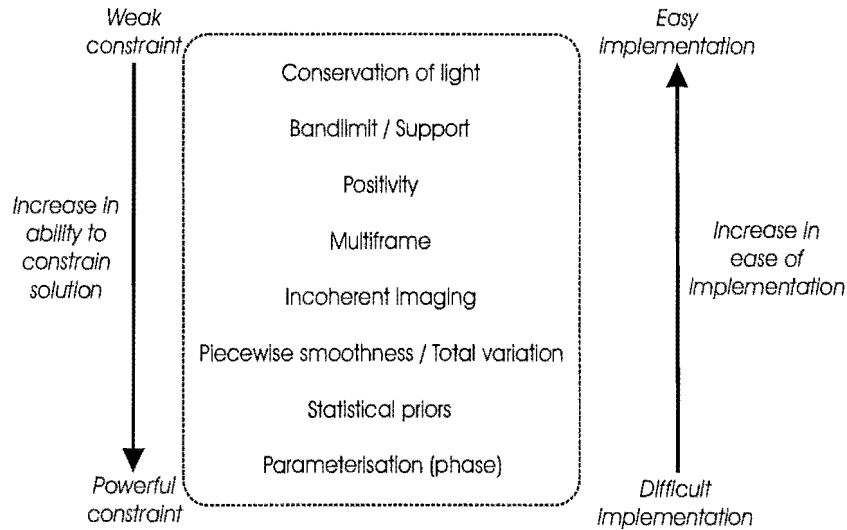
Figure 6.2: *An illustration of the relationship between the ease of imple-*
mentation and the ability to constrain a solution for some commonly
used priors. In general, the more difficult the implementation, the more
powerful the constraint is.

are included here for completeness. Section 6.3.1 discusses the general blind deconvolution
algorithms, while Section 6.3.2 presents the methods specific to astronomical imaging.

6.3.1   General blind deconvolution algorithms

In Kundur and Hatzinakos' NASRIF [102] (nonnegativity and support constraint recursive
inverse filtering) scheme, the inverse of the PSF (instead of the PSF) is estimated and is
expressed as essentially a highpass filter. The filter coefficients are constrained to be greater
than or equal to zero and must sum to unity. Other constraints imposed are the positivity
and support for the object, and that the PSF and its inverse are positive and absolutely
summable. The main advantage of this algorithm is that because explicit knowledge of
the PSF is not required, the object estimate is simply a weighted linear sum of the data,
resulting in easy implementation and analysis. Uniqueness of the solution is ensured by the
convexity of the cost function. However, one problem with this algorithm is that the filter
size is variable, making an optimal size difficult to obtain. This algorithm also suffers from
the lack of regularisation, and visual inspection is sometimes required for termination. This
problem has been addressed by Ng *et al.* [128] who employ eigenvalue truncation and total

variation to alleviate the problem.

You and Kaveh's method [184] employs anisotropic regularisation which assumes piecewise smoothness on both the object and the PSF, and is suitable for reconstructing functions with sharp transitions. However, this algorithm suffers from local minima due to a non-quadratic cost function. The termination of iterations is also sometimes required, which is essentially a form of implicit regularisation. The total variation (TV) method proposed by Chan and Wong [21] is similar to You and Kaveh's method as it also utilises anisotropic regularisation.

Yang et al. [183] proposed a projection-based algorithm, where the prior knowledge of positivity and support of the object and PSF is expressed through convex sets (POCS). In a later paper by Law and Lane [110], it was pointed out that the PSF used in demonstrating the algorithm was a truncated Gaussian and presented many ambiguous solutions to the problem investigated. This is because a Gaussian PSF can be represented as the convolution of two smaller Gaussians which introduces multiple solutions to the problem. Results presented in Ref. [110] also showed no significant advantage of the projection-based method over a statistical approach for a more realistic example. This technique was therefore not considered further.

Fish et al. [36] used a Richardson-Lucy based algorithm and addressed the semiblind deconvolution situation where only the several parameters that characterise the PSF are estimated. As with Yang et al., a Gaussian PSF was used to demonstrate the algorithm which introduced ambiguity in the solution.

McCallum [120] proposed the application of simulated annealing to blind deconvolution. The technique of simulated annealing was originally devised for calculating state functions in statistical mechanics [121]. It has the advantage of being able to overcome the local minima in multimodal cost functions. Positivity and support are also enforced. However, this algorithm is extremely computationally intensive for even a single frame, and is not suitable for use for multiframe blind deconvolution.

6.3.2 Astronomical blind deconvolution algorithms

Schulz [154] proposed two algorithms to bias the solution away from the trivial solution. These have been proposed specifically for multiframe blind deconvolution in the presence of Poisson noise. The first method, the penalised maximum likelihood (ML+) algorithm, uses an *ad hoc* prior on the object:

$$P(\hat{f}(x,y)) = -\beta \sum_{x,y} \log[1 - \hat{f}(x,y)] \tag{6.11}$$

where β is the weighting of the prior, and the object estimate must be normalised to have unit energy. As noted earlier, a trivial solution when using multiple frames for blind deconvolution can only take the form of the object being a delta function, and the PSF the data, but not vice versa. By forcing the object estimate away from a delta function, this prior greatly reduces the possibility of the trivial solution. However, since the object estimate is strongly biased away from a point source, care should be taken to ensure that the underlying object is not actually a point source or close to a point source. Such an occasion can arise when a bright star has a faint companion. This prior also favours objects that are flat and smooth, which must be taken into account when using this prior.

Schulz's second method, the expectation-maximisation Richardson-Lucy (EM-RL) algorithm, parameterises the PSF to arise from the phase errors distributed over an aperture. This powerful prior has been shown to be very successful in constraining the solution. Positivity of the PSF is guaranteed by the phase parameterisation, while positivity of the object is enforced by initialising the RL algorithm with a positive estimate.

Lane [105] proposed an algorithm which incorporates the object and PSF constraints into the error metric, thus reformulating the problem to one of unconstrained minimisation. These constraints include positivity and support. In a later paper by Law and Lane [110], the issues associated with weighting the terms in the error metric, and the derivation of the algorithm from a maximum likelihood framework were addressed.

Jefferies and Christou [90] extended Lane's algorithm by using additional constraints such as the bandlimit constraint on the PSF. Thiébaut and Conan [168] noted that it is necessary to enforce the positivity and support constraints strongly for good results in Lane's method. As a result, they proposed a technique which strictly enforces the positivity and support constraints by reparameterising $f(x,y)$ and $h(x,y)$. The PSF is also constrained to arise

from the optical instrument. A problem with reparameterisation is that it distorts an error surface, resulting in one which may have local minima [94]. A unique solution is therefore not guaranteed. Johnston and Lane [94] proposed a quadratic programming technique to overcome this problem, ensuring positivity as well as preserving the unimodality for a parabolic error surface.

Sheppard et al. [157] extended the work by Thiébaut and Conan [168] and proposed a multiframe blind deconvolution algorithm based on the Bayesian maximum likelihood Richardson-Lucy framework (PMAP). Strict positivity, bandwidth and conservation of light constraints are applied to the PSF through nonlinear reparmeterisation. A strict positivity constraint is also applied on the object via the reparameterisation $\hat{f}(x,y) = \hat{\varphi}^2(x,y)$. This is required because the application of the conjugate-gradient method on data with Poisson noise statistics results in negative values in the estimate if positivity is not enforced explicitly.

Miura et al. [123] proposed an algorithm constraining the Knox-Thompson (KT) spectral phases of the object estimate to fit the measured ones. Other constraints used are positivity, support, multiframe, PSF parameterisation, and conservation of light. This method is shown to be robust against the trivial solution even in low light levels. Unlike most algorithms, the object and the PSFs are estimated alternately with different error metrics: the error metric for the object is based on the difference between the measured and the estimated KT phases, while the PSF error metric is the negative of the log-likelihood based on the Poisson noise model. Since it is unlikely that the local minima in both error metrics coincide, this scheme is more robust compared to single-error metric schemes where stagnation in local minima is a problem.

A class of blind deconvolution algorithms, called myopic deconvolution methods, has recently been proposed as a result of the need to post-process adaptive optics imagery. Good prior knowledge of the PSF and the object is available, which can be used as starting estimates for a blind deconvolution algorithm. Conan et al. [27] were the first to address this problem. They used the mean and the power spectral density (PSD) of the object spectrum, the OTF and the noise spectrum to constrain the solution. The PSD of the object is a function of the photon level, the cutoff frequency and the mean image, while the PSD of the OTF is estimated from either control-loop data (in an adaptive optics system) or from reference star measurements. The images are assumed to have Poisson noise statistics.

Fusco *et al.* [51] extended this approach to the case when the object is known to be made up of point sources. In this case, the object estimation only involves the retrieval of their respective locations and magnitudes which greatly reduces the number of unknowns in the system. Mugnier *et al.* [126] further developed this algorithm to incorporate wavefront sensing data captured simultaneously with the speckle images to constrain the phase and the PSF. This method has produced good results on both simulated and experimental data.

Christou *et al.* [24] also proposed the concept of myopic deconvolution alongside Conan *et al.* The average of the centroided data image is used to initialise the object, while the initial PSF is the average of the centroided images of a bright point source. Other constraints such as positivity, a bandlimit on the PSF, and a speckle holography constraint are used. The success of this algorithm has been demonstrated on simulated and experimental data.

It should be pointed out that the myopic deconvolution schemes deal specifically with astronomical imaging data, and in general the initial estimates are of very good quality. It is of course possible to apply these methods to a blind deconvolution problem, but the performance may not be as good due to the greater number of initial unknowns.

6.3.3 Proposed investigations

In this chapter, two different forms of prior information for astronomical imaging are proposed. Firstly, since the PSF and hence the complex aperture are directly related to the atmosphere, it is reasonable to constrain the latter to obey Kolmogorov statistics. This is compared with Schulz's algorithm where the PSF is constrained to arise from the phase errors distributed over an aperture [154]. Secondly, the property of phase closure which has been used in the well-established technique of bispectrum [113] is used as a constraint on the PSF in a blind deconvolution framework. The respective theories and feasibility investigations are presented in the following sections.

6.4 Using Kolmogorov statistics

As discussed in Chapter 4, the distortions introduced by atmospheric turbulence on a propagating wavefront can be described by Kolmogorov statistics [147]. Using a different notation

Method	Assumptions		
	Noise	Object	Point Spread Function
Kundur & Hatzinakos	G	+ve, support	$\|\Sigma h\| < \infty, \|\Sigma h^{-1}\| < \infty$
You & Kaveh	G	+ve, piecewise smoothness	+ve, light, piecewise smoothness
Chan & Wong	G	+ve, piecewise smoothness	+ve, light, sym, piecewise smoothness
Yang et al.	G	+ve, support	+ve, support
Fish et al.	P	+ve, light	+ve, light, semiblind
McCallum	G	+ve, support	+ve, support
Schulz ML+	P	prior, +ve, light	+ve, light, multi
Schulz phase param	P	+ve, light	+ve, light, incoherent imaging, multi, param
Lane	G	loose +ve and support	loose +ve and support
Jefferies & Christou	G	+ve, support, Fourier modulus	+ve, support, bandlimit, multi
Thiébaut & Conan	G	+ve and support by param	+ve and light by param, incoherent imaging
Sheppard et al.	P	+ve by param	+ve, light and bandlimit by param, multi
Miura et al.	P	+ve, support, KT phase, light	+ve, multi, phase by Zernike polynomials
Conan et al.*	P	PSD, +ve	PSD
Fusco et al.*	G	number of stars	PSD
Mugnier et al.*	G	edge-preservation	statistics
Christou et al.*	G	+ve	+ve, bandlimit
Proposed: Kolmogorov	P	+ve, light	statistics, +ve, light, multi
Proposed: bispectrum	P	+ve	+ve, multi, phase closure

Table 6.1: *Summary of the assumptions made on the noise model, the object, and the PSF of existing iterative blind deconvolution techniques. P = Poisson, G = Gaussian, +ve = positivity, light = conservation of light, param = parameterisation, multi = multiframe, sym = centrosymmetric, PSD = power spectral density, KT = Knox-Thompson, * = myopic deconvolution. The two proposed methods are included at the bottom of the table for comparison.*

here, the wavefront at the complex aperture is denoted by:

$$A_k(u,v) = W(u,v) \exp\{\chi_k(u,v) + i\phi_k(u,v)\} \tag{6.12}$$

where $\chi_k(u,v)$ and $\phi_k(u,v)$ are respectively the log-amplitude and the phase of the complex wavefront. $W(u,v)$ is the aperture weighting function and is assumed to be constant. It is therefore not included in the subsequent equations. On expanding Eq. (6.12), the expression becomes the familiar equation describing the complex aperture:

$$A_k(u,v) = S_k(u,v) \exp\{i\phi_k(u,v)\} \tag{6.13}$$

where $S_k(u,v) = \exp\{\chi_k(u,v)\}$ is the magnitude of the wavefront. The scintillation $S_k^2(u,v)$ is the intensity distortion which arises from the wavefront propagating through the atmosphere. The optical transfer function (OTF) is the autocorrelation of the generalised pupil function:

$$H_k(u,v) = S_k(u,v) \exp\{i\phi_k(u,v)\} \otimes S_k(u,v) \exp\{i\phi_k(u,v)\} \tag{6.14}$$

and has a Fourier transform relationship with the PSF:

$$H_k(u,v) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} h_k(x,y) \exp\{-i2\pi(ux + vy)\} dx dy. \tag{6.15}$$

Similarly, the generalised pupil function $A_k(u,v)$ forms a Fourier transform pair with the complex image $a_k(x,y)$ at the focal plane of the telescope where $a_k(x,y) \rightleftharpoons A_k(u,v)$. The relationship between the PSF and the complex image $a_k(x,y)$ can therefore be expressed as:

$$h_k(x,y) = |a_k(x,y)|^2. \tag{6.16}$$

It is also possible to express the complex wavefront by its real and imaginary parts, where

$$A_k(u,v) = \Re\{A_k(u,v)\} + i\Im\{A_k(u,v)\}. \tag{6.17}$$

The corresponding complex image $a_k(x,y)$ at the image plane of the telescope is therefore related to the PSF by

$$h_k(x,y) = |\Re\{a_k(x,y)\}|^2 + |\Im\{a_k(x,y)\}|^2. \tag{6.18}$$

There are two alternatives for incorporating the prior information into the framework: by using the statistics of the scintillation and the phase, or by the statistics of the real and imaginary parts of the complex aperture.

Under the assumption of weak turbulence, the log-amplitude $\chi_k(u,v)$ and the phase distortions $\phi_k(u,v)$ are correlated. The cross-correlation is expressed as [100, 166]:

$$B_{\chi\phi}(\rho) = 2\pi^2 k^2 \times 0.033 \int_0^\infty dh \ C_n^2(h) \int_0^\infty dp \ p^{-8/3} \times J_0(p\rho) \sin(\frac{h}{k}p^2) \qquad (6.19)$$

where $C_n^2(h)$ is the vertical profile of the refractive-index structure characteristic, h is the height of the turbulence, J_0 is the zeroth-order Bessel function of the first kind, $k = 2\pi/\lambda$ is the wave number, λ is the wavelength of the light, and p and ρ are the horizontal distances in the spatial and frequency domains respectively. As a result of this cross-correlation, the scintillation and the phase are not independent and it is difficult to incorporate prior information of these quantities.

On the other hand, the real and imaginary parts of the complex wavefront are uncorrelated. As a result, the prior information on the complex wavefront can be incorporated via these quantities. Two further simplifications are possible according to Roddier [140]. Firstly the covariance of a complex field at ground level is unaltered by propagation and is equivalent to that at height h. Secondly, because the scintillation is small compared to the phase distortion, the former can be assumed to be constant. As a consequence, the covariance of the real and imaginary parts of the complex wavefront at the aperture is:

$$C(u_1, v_1, u_2, v_2) = \exp\{-\frac{1}{2}D_\phi(u_1, v_1, u_2, v_2)\} \qquad (6.20)$$

where D_ϕ is the phase structure function. This is the same as Eq. (4.49) and is reproduced here for completeness. It should be pointed out that the simplifications above have made the solution easier to obtain, without much sacrifice for its accuracy. The simplifications are therefore justified.

The implementation of this method requires the definition of an error metric and their gradients. The error to be minimised is:

$$E = E_{\text{con-Poisson}} + \gamma_h E_h. \qquad (6.21)$$

Since a direct relationship exists between the PSF $h_k(x,y)$ and the real and imaginary parts of the complex aperture $A_k(u,v)$, Eq. (6.21) can be written as:

$$E = E_{\text{con-Poisson}} + \gamma_A E_{\Re\{A\}} + \gamma_A E_{\Im\{A\}}. \qquad (6.22)$$

The necessary gradient calculation for updating $\hat{f}(x,y)$ is shown in Eq. (6.9). By using a positive starting estimate for $\hat{f}(x,y)$ in a Richardson-Lucy framework, its subsequent

estimates are also guaranteed to be positive. In this algorithm, the constraint of light conservation is also applied to the object.

The gradients necessary for updating the real and imaginary parts of the complex aperture are:

$$
\begin{aligned}
\frac{\partial E_{\text{con-Poisson}}}{\partial \Re\{\hat{A}_k(u,v)\}} &= \Re\left\{\mathfrak{F}\left\{\frac{\partial E_{\text{con-Poisson}}}{\partial \hat{a}_k(x,y)}\right\}\right\} \\
\frac{\partial E_{\text{con-Poisson}}}{\partial \Im\{\hat{A}_k(u,v)\}} &= \Im\left\{\mathfrak{F}\left\{\frac{\partial E_{\text{con-Poisson}}}{\partial \hat{a}_k(x,y)}\right\}\right\}
\end{aligned}
\tag{6.23}
$$

where

$$
\frac{\partial E_{\text{con-Poisson}}}{\partial \hat{a}_k(x,y)} = 2\left(\frac{d_k(x,y) - \hat{g}_k(x,y)}{\hat{g}_k(x,y)}\hat{f}^*(x,y)\right)\hat{a}_k(x,y).
\tag{6.24}
$$

The prior information regarding the complex wavefront is expressed as a function of covariances and is incorporated in the gradient calculation. The log-likelihood of the prior for the real part of the complex field is:

$$
\log[P(\Re\{\hat{A}_k(u,v)\})] = \frac{1}{2}\Re\{\hat{\mathbf{A}}_k\}^T\mathbf{K}^{-1}\Re\{\hat{\mathbf{A}}_k\}
\tag{6.25}
$$

where $\Re\{\hat{\mathbf{A}}_k\}$ is the lexicographically ordered vector of the real part of the complex wavefront, and $\mathbf{K}$ is the two-dimensional covariance matrix constructed using the covariance C. The structure of $\mathbf{K}$ is shown in Fig. 4.16(c). To maximise this likelihood, the derivative is taken to give:

$$
\frac{\partial \log[P(\Re\{\hat{A}_k(u,v)\})]}{\partial \Re\{\hat{A}_k(u,v)\}} = U[\Re\{\hat{\mathbf{A}}_k\}^T\mathbf{K}^{-1}]
\tag{6.26}
$$

where $U[\cdot]$ is the operator which reorders the vector into a two-dimensional quantity. The prior for the imaginary part of the complex aperture is obtained in a similar manner.

6.4.1   Feasibility investigation

The simulations performed to test the feasibility of the algorithm are as follows. The object is a binary star with a brightness ratio of 1.25 and a separation of 3.61 pixels. Assuming a telescope with a diameter of 1m and that the wavelength of the light is 650nm, this separation corresponds to 0.59 arcsec. The phase distortions introduced at the turbulent

No. photons	Proposed			Schulz's method		
	Success rate	Mean	Variance	Success rate	Mean	Variance
50	10 %	1.129	0.206	14 %	1.014	0.113
500	73 %	1.279	0.060	54 %	1.295	0.068
5000	98 %	1.295	0.004	90 %	1.277	0.034

Table 6.2: *Results of blind deconvolution with Kolmogorov statistics prior compared to Schulz's phase parameterisation method with an ensemble of 100 trials.*

layer are modelled using the method of Harding *et al.* [76] where $D/r_0 = 4$. The phase screens are of size 16×16. The data is assumed to have Poisson statistics, where the total number of photons in each frame is set to either 50, 500 or 5000. 2 frames of data are used in each of the 100 trials. The starting estimates of the phase are obtained by extracting the tip and tilt components from the speckle images and forming corresponding PSF estimates. The initial estimate of the object is uniform. The results of reconstructing the object using the proposed method, and those obtained using Schulz's phase parameterisation algorithm [154] are tabulated in Table 6.2. The success rate indicates whether a binary star with a correct separation is reconstructed, and the corresponding means and variances are calculated using only the successful trials.

For the photon levels of 5000 and 500, the proposed algorithm performs better than Schulz's algorithm in terms of both success rate and robustness. For the photon level of 50, both algorithms perform badly, although Schulz's algorithm has a higher success rate than the proposed method. It should be pointed out that reconstructions are relatively insensitive to the weighting of the prior γ_A at high photon levels (5000 and 500). As the number of photons decreases, more emphasis on the prior information is needed to obtain a solution. For the data with 5000 photons per frame, $\gamma_A = 1 \times 10^{-4}$. When there are only 50 photons per frame, γ_A has to be increased up to 0.2.

6.4.2 Summary and future work

In this section, the feasibility of using Kolmogorov statistics as prior information for the PSF in a blind deconvolution solution is investigated. From the simulations performed, it is shown that this approach is feasible and produces promising results at relatively high

light levels. However, a more thorough investigation is necessary in order to completely characterise the performance of this method. The factors to take into consideration include the use of more data frames, varying the severity of the atmospheric turbulence to reflect the characteristics of real data, and dealing with the effect of scintillation on the data. The ultimate goal is of course to test the algorithm on real data.

6.5   The technique of bispectrum for blind deconvolution

One of the first methods proposed for astronomical imaging which used the concept of phase closure was the speckle interferometry technique suggested by Labeyrie [103] in 1970. The modulus of the object spectrum is retrieved by:

$$|F(u,v)| = \left(\frac{\langle |D(u,v)|^2 \rangle}{\langle |H(u,v)|^2 \rangle + \text{constant}} \right)^{1/2}. \tag{6.27}$$

This approach is superior to techniques which use only one frame of data because frequencies up to the diffraction limit are preserved when the second moment is used in the calculations [147]. Although the atmospheric phase is eliminated and phase closure is achieved by this technique, the phase information of the object is not preserved. Consequently the object distribution in the spatial domain cannot be recovered.

In 1983, Lohmann et al. [113] proposed a speckle masking, or bispectrum technique to reconstruct the object. The correlation of three instead of two quantities means that the wavefronts with Kolmogorov statistics average to zero over a large ensemble. As a result, phase closure is achieved and phase preservation is possible. Consequently, both the phase and the magnitude of the object spectrum can be recovered.

This technique is divided into two main sections: triple correlation theory, and object reconstruction. Since only the former is used in the proposed blind deconvolution scheme, this latter is not discussed here. For details on the complete bispectrum technique, the reader is referred to Ref. [113].

6.5.1 Bispectrum, or triple correlation theory

The two-dimensional intensity distribution of the k^{th} speckle image in the absence of noise is defined as:

$$d_k(\mathbf{x}) = f(\mathbf{x}) \odot h_k(\mathbf{x}) \qquad (6.28)$$

where $\mathbf{x}$ is a two-dimensional vector in the spatial domain and is adopted in this section to simplify the notation. $f(\mathbf{x})$ is the object of interest, and $h_k(\mathbf{x})$ is the combined PSF of the atmosphere and the telescope.

The triple correlation of a speckle image is defined as:

$$
\begin{aligned}
\langle d^{(3)}(\mathbf{x}, \mathbf{x}') \rangle &= \langle [d(\mathbf{x})d(\mathbf{x} + \mathbf{x}')] \otimes d(\mathbf{x}) \rangle \\
&= \langle \int d(\mathbf{x}'')d(\mathbf{x}'' + \mathbf{x}')d(\mathbf{x}'' + \mathbf{x})d\mathbf{x}'' \rangle
\end{aligned}
\qquad (6.29)
$$

where the angular brackets denote an ensemble average over many statistically independent frames, and the superscript (3) denotes a triple correlation quantity. Essentially, this equation means that each speckle image $d_k(\mathbf{x})$ is multiplied by a shifted version of itself $d_k(\mathbf{x} + \mathbf{x}')$ to produce a product mask. This mask is then cross-correlated with $d_k(\mathbf{x})$ to form the triple correlation. Finally the ensemble average of the triple correlation is taken over many frames. In the Fourier domain, an equivalent quantity called the bispectrum or the triple product can be calculated:

$$D_k^{(3)}(\mathbf{u}, \mathbf{v}) = D_k(\mathbf{u})D_k(\mathbf{v})D_k(-\mathbf{u} - \mathbf{v}) \qquad (6.30)$$

where $\mathbf{u}$ and $\mathbf{v}$ are two-dimensional coordinates in the frequency domain. $D_k(\mathbf{u})$, $D_k(\mathbf{v})$ and $D_k(-\mathbf{u} - \mathbf{v})$ are identical Fourier transforms of $d_k(\mathbf{x})$ indexed differently. It should be pointed out that the bispectrum of a two-dimensional quantity is a four-dimensional function. It is because of this extension into four-dimensional space that phase information can survive.

In speckle masking, the relationship between the object and image bispectra is:

$$\langle D^{(3)}(\mathbf{u}, \mathbf{v}) \rangle = F^{(3)}(\mathbf{u}, \mathbf{v})\langle H^{(3)}(\mathbf{u}, \mathbf{v}) \rangle. \qquad (6.31)$$

It has been shown [113] that the masking transfer function $\langle H^{(3)} \rangle$ is real and greater than zero up to the telescope cutoff frequency. Due to this property, the phase of the complex bispectrum of the object is identical to the phase of the average bispectrum of the speckle

images. This is a direct result of the theory of phase closure presented in the following section.

The bispectrum technique proposed by Lohmann *et al.* [113] firstly calculates the object bispectrum, from which the object spectrum is then unravelled. The modulus of the object spectrum is calculated as in speckle interferometry [103], while a recursive relationship is used to estimate the object spectrum phase. These quantities are then combined, whose inverse Fourier transform yields the object distribution in the focal plane. Although this has been proven to be an effective method, the fact that the object bispectrum estimation is not updated means that an error in this quantity can manifest as an error in the object reconstruction.

In the proposed method, the bispectrum phase of the PSF is used as a penalty term. This quantity is updated at every iteration, reducing the possibility of error propagation as in the original method. Moreover, the formulation which incoporates a data consistency measurement as well as this penalty term ensures that a balance is achieved between the fidelity to the measured data and the prior information.

6.5.2 Phase closure

The property that the average phase of the bispectrum of the PSF is zero is due to the phase closure characteristic. Essentially, phase closure can be defined as the simultaneous preservation of the object phase and elimination of the undesirable phase component which is the atmospheric phase in this case. Based on the work of F. Roddier [141], let the atmosphere be modelled by a distribution of seeing cells of diameter r_0 in frequency space. The wavefront phase is assumed to be constant across each seeing cell, and the complex amplitude in the cells is assumed to be uncorrelated.

Consider two seeing cells in the telescope pupil, separated by a vector $\lambda \mathbf{u}$, with λ being the optical wavelength. If a point source is imaged through the pupil with two apertures corresponding to the two seeing cells, a fringe pattern with a narrow frequency bandwidth is produced. The major component $D(\mathbf{u})$ at the frequency $\mathbf{u}$ is produced by contributions of all pairs of points with a separation $\lambda \mathbf{u}$ (Fig. 6.3(a)). For frequencies less than r_0/λ, the average of the major component is non-zero. However, for frequencies greater than r_0/λ,
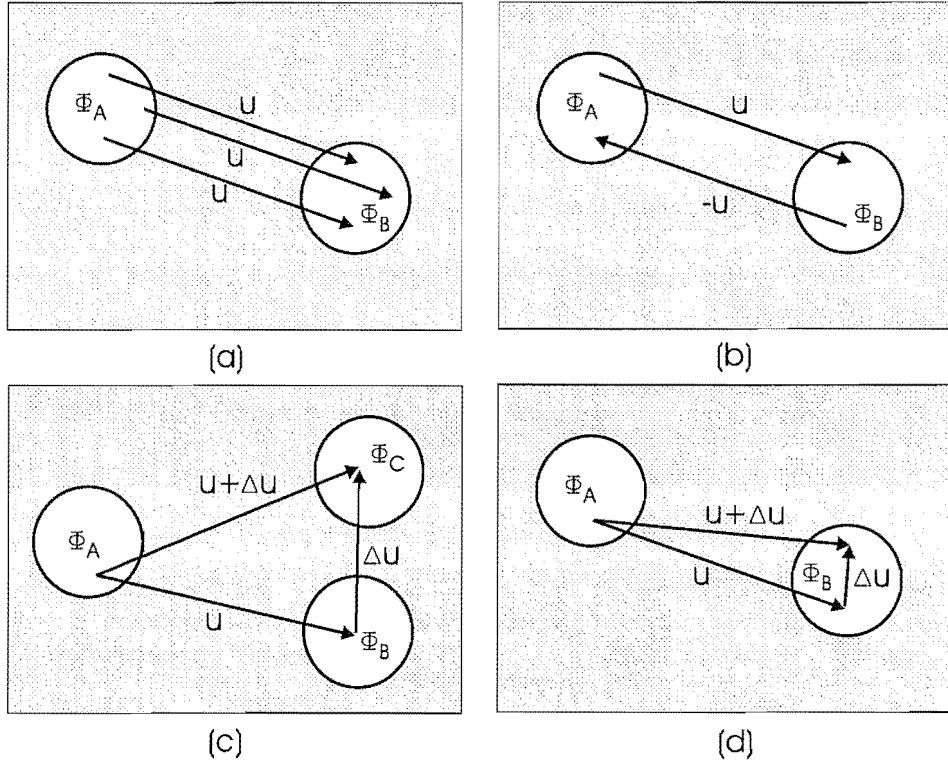
Figure 6.3: *The illustration of phase closure. (a) Imaging a point source through two sub-apertures. (b) Extending the idea to auto-correlation. Phase closure is achieved here. (c) Phase closure is achieved in the bispectrum technique. (d) Phase closure is achieved in the Knox-Thompson method for certain frequencies.*

the average of the major component tends to zero because the phase difference $\Phi_A - \Phi_B$, mod 2π, between the two apertures is an evenly-distributed zero-mean process. As a result, the Fourier component performs a random walk in the complex plane and averages to zero. The phase and the magnitude in this case are:

$$\arg\{D(\mathbf{u})\} = \phi(\mathbf{u}) + (\Phi_A - \Phi_B)$$
$$\langle D(\mathbf{u})\rangle \neq 0, \qquad u < r_0/\lambda$$
$$\langle D(\mathbf{u})\rangle = 0, \qquad u > r_0/\lambda \qquad (6.32)$$

where $\phi(\mathbf{u})$ is the Fourier phase at frequency $\mathbf{u}$.

If the same analysis is applied to the autocorrelation technique (Fig. 6.3(b)), the phase and

the magnitude are:

$$\arg\{D(\mathbf{u})D(-\mathbf{u})\} = \phi(\mathbf{u}) + (\Phi_A - \Phi_B) + \phi(-\mathbf{u}) + (\Phi_B - \Phi_A) = 0$$

$$\langle D(\mathbf{u})D(-\mathbf{u})\rangle = \langle |D(\mathbf{u})|^2\rangle \neq 0. \tag{6.33}$$

Although phase closure is achieved, the phase information of the object is not preserved. This is the principle employed in speckle interferometry, and as a result the object phase cannot be reconstructed.

In speckle masking, the number of seeing cells is increased to three (Fig. 6.3(c)). The phase and magnitude are:

$$\arg\{D(\mathbf{u})D(\Delta\mathbf{u})D(-\mathbf{u} - \Delta\mathbf{u})\} = \phi(\mathbf{u}) + (\Phi_A - \Phi_B) +$$

$$+ \phi(\Delta\mathbf{u}) + (\Phi_B - \Phi_C) + \phi(-\mathbf{u} - \Delta\mathbf{u}) + (\Phi_C - \Phi_A) \tag{6.34}$$

$$= \phi(\mathbf{u}) - \phi(\mathbf{u} + \Delta\mathbf{u}) + \phi(\Delta\mathbf{u}) \tag{6.35}$$

respectively. Here, phase closure is achieved as the effect of the atmospheric phase is eliminated, and both the object phase and magnitude information is preserved for frequencies up to the cutoff frequency of the telescope.

However, the calculation of the full bispectrum is computationally intensive. For a complex wavefront of size $N \times N$, the bispectrum is a quantity of size N^4. When multiple frames are used, which is usually in the order of 10^2, it is apparent that the computational requirement of this technique becomes impractical.

A simplification of the bispectrum technique is the Knox-Thompson (KT) method [99] where only two seeing cells are used. As a consequence, the vector $\Delta\mathbf{u}$ in Fig. 6.3(c) must be restricted to lie within the seeing cell as shown in Fig. 6.3(d) to ensure that phase closure is achieved. This corresponds to the constraint of the frequency u being less than r_0/λ. The phase is therefore:

$$\arg\{D(\mathbf{u})D(-\mathbf{u} - \Delta\mathbf{u})\} = \phi(\mathbf{u}) + (\Phi_A - \Phi_B)$$

$$+ \phi(-\mathbf{u} - \Delta\mathbf{u}) + (\Phi_B - \Phi_A) \tag{6.36}$$

$$= \phi(\mathbf{u}) - \phi(\mathbf{u} + \Delta\mathbf{u}) \tag{6.37}$$

while the magnitude $\langle D(\mathbf{u})D(-\mathbf{u} - \Delta\mathbf{u})\rangle$ is non-zero.

The implementation of the KT simplification is achieved by letting each of the two-dimensional frequency variables $\mathbf{u}$ and $\mathbf{v}$ be written as (u_1, u_2) and (v_1, v_2) respectively. Then instead of using all the combinations of these four variables to calculate the bispectrum $D_k^{(3)}(u_1, v_1, u_2, v_2)$, only the two planes which contribute most significantly to this quantity are used. These two planes correspond to $u_2 = 0, v_2 = 1$ and $u_2 = 1, v_2 = 0$.

The difference between Eqs. (6.36) and (6.37), and Eqs. (6.34) and (6.35) is that in the KT technique, the term $\phi(\Delta\mathbf{u})$ is absent. This is directly related to the fact that only a double instead of a triple product is used. To illustrate its significance, let the system of two seeing cells be extended to three and let u be made greater that r_0/λ. The phase of the system according to an autocorrelation operation is now:

$$\arg\{D(\mathbf{u})D(-\mathbf{u} - \Delta\mathbf{u})\} = \phi(\mathbf{u}) + (\Phi_A - \Phi_B)$$
$$+\phi(-\mathbf{u} - \Delta\mathbf{u}) + (\Phi_C - \Phi_A) \tag{6.38}$$

$$= \phi(\mathbf{u}) - \phi(\mathbf{u} + \Delta\mathbf{u}) - \Phi_B + \Phi_C \tag{6.39}$$

and clearly the atmospheric phase contribution is not closed. Therefore, the main difference between the bispectrum and the KT techniques is that the latter uses an autocorrelation calculation while the former uses a triple product. As a result, the KT method can only estimate correctly the lower frequency components due to the constraint imposed on $\Delta\mathbf{u}$, whereas the bispectrum method can estimate frequencies up to the diffraction limit.

6.5.3 Bispectrum usage in object reconstruction

Since Lohmann *et al.* proposed the application of bispectrum in astronomy, many algorithms had been devised to reconstruct the object distribution using this information. In the past ten years, the least-squares approach was popular. This section briefly outlines these techniques.

Haniff [72] proposed a method which minimises the modulo 2π phase difference between the measured and the estimated object bispectrum phases. The measured phase is that reconstructed from speckle images, while the estimated phase is that calculated using the spectrum of the object estimate. This provided an advantage over algorithms which required bispectrum phase-unwrapping. However, the modulus of the object spectrum required reference star measurements to be calculated.

The reconstruction of an object using bispectrum information, without the need of a reference star, was first proposed by Glindemann *et al.* [59]. In this paper, the special case of a binary star was considered because it can be parameterised by two variables (brightness ratio and separation) in both the object and bispectrum domains. This technique was later extended for use with an arbitrary object [58]. As with Haniff [72], the error to be minimised was the object bispectrum phase difference between the measured and estimated phases. Since the object is directly related to the estimated phase, the former can be reconstructed directly via the error metric, without employing the procedure of Lohmann *et al.* A positivity constraint was also incorporated according to a proposal by Lane [105]. Miura *et al.* [124,125] proposed two similar techniques, minimising respectively the difference between the estimated KT and bispectrum phases with the corresponding measured phases. In many of these techniques, not all the subplanes of the bispectrum are used in the reconstruction due to heavy computational requirements. Usually, only a small number of subplanes with high SNR are used.

However, it is only recently that the use of the KT phase has been incorporated into the blind deconvolution problem [123]. Similar to the previous techniques, the object is reconstructed by minimising the error between the measured KT phase and the estimated one. At the same time, the PSF is parameterised using Zernike polynomials and the associated error metric is the negative of the log-likelihood based on the Poisson noise model. Since it is unlikely that the local minima in both error metrics coincide, this scheme is more robust compared to single-error metric schemes where stagnation in local minima is a problem.

In all the algorithms above, the bispectrum or KT phase is reconstructed first. It is then used as the reference with which the estimated quantities are compared. Although the KT phase and especially the bispectrum phase are robust to noise, the calculation of such quantities means that the data has been processed. The error in reconstructing the bispectrum (or KT) phase is therefore reflected directly in the object reconstruction. This data processing, although justified for certain circumstances, such as in Chapter 8 where computational time is a major concern, is not necessary in this application. In the proposed method, the statistical distribution of the KT phase is used as prior information. The iterative framework means that the object and the KT phase are updated simultaneously, removing the reliance of the solution on an accurate KT reconstruction initially.

It should be noted that in the application of the bispectrum method to real data, the

telescope aberrations should be calibrated and compensated for. This can be achieved by measuring the effect of imaging an artificial point source through the optical system. However, this procedure is not performed here since simulated data is used.

6.5.4 Using the phase closure property as a prior

The application of the phase closure property to blind deconvolution is investigated in this section. Since the phase of the average of the PSF bispectrum should be zero, the entire phase component in the bispectrum domain can be attributed to the object. The KT approximation is used here to simplify the implementation.

Let the KT phase of the PSF be $\kappa(u_1, v_1, u_2, v_2)$. Then the associated error metric for this prior can be defined as the deviation of this term from zero, represented mathematically as:

$$E = \alpha \sum_{u_1, v_1} \frac{|\kappa(u_1, v_1, u_1 + \Delta, v_1)|^2}{\sigma^2} \qquad (6.40)$$

for the 10-plane, where $u_2 = u_1 + \Delta$ and $v_2 = v_1$, and

$$E = \alpha \sum_{u_1, v_1} \frac{|\kappa(u_1, v_1, u_1, v_1 + \Delta)|^2}{\sigma^2} \qquad (6.41)$$

for the 01-plane, where $u_2 = u_1$ and $v_2 = v_1 + \Delta$. The value of Δ is 1, and α is the weighting of the prior. The variance σ^2 represents the distribution of phase information [7] and has been used to weight error metrics in other algorithms using bispectrum information [58, 72, 125]. Since this prior information is used to constrain the solution of the PSF, this must be incorporated in the gradient calculation for updating the PSF. In other words, we wish to quantify the effect on the error when the PSF is perturbed. This calculation requires the application of the chain rule, and is more easily performed in the frequency domain:

$$\frac{\partial E(u, v)}{\partial H(u, v)} = \frac{\partial E(u, v)}{\partial \kappa(u_1, v_1, u_2, v_2)} \times \frac{\partial \kappa(u_1, v_1, u_2, v_2)}{\partial H^{(3)}(u_1, v_1, u_2, v_2)} \times \frac{\partial H^{(3)}(u_1, v_1, u_2, v_2)}{\partial H(u, v)}. \qquad (6.42)$$

It should be noted that u and v are generic one-dimensional frequency variables. The differentiation of each of the three terms is shown below.

Term 1:

$$\frac{\partial E(u, v)}{\partial \kappa(u_1, v_1, u_2, v_2)} = 2\alpha \frac{\kappa(u_1, v_1, u_2, v_2)}{\sigma^2} \qquad (6.43)$$

Term 2:

$$\frac{\partial \kappa(u_1, v_1, u_2, v_2)}{\partial H^{(3)}(u_1, v_1, u_2, v_2)} = \frac{1}{iH^{(3)}(u_1, v_1, u_2, v_2)} \tag{6.44}$$

Term 3:

$$\frac{\partial H^{(3)}(u_1, v_1, u_2, v_2)}{\partial H(u, v)} = H(u_1, v_1) \left[\frac{\partial H(u_2, v_2) H^*(u_1 + u_2, v_1 + v_2)}{\partial H(u, v)} \right]$$
$$+ H(u_2, v_2) H^*(u_1 + u_2, v_1 + v_2) \frac{\partial H(u_1, v_1)}{\partial H(u, v)}. \tag{6.45}$$

Applying the chain rule to the first term of this equation gives:

$$\frac{\partial H^{(3)}(u_1, v_1, u_2, v_2)}{\partial H(u, v)} = H(u_1, v_1) \left[H(u_2, v_2) \frac{\partial H^*(u_1 + u_2, v_1 + v_2)}{\partial H(u, v)} + \right.$$
$$\left. H^*(u_1 + u_2, v_1 + v_2) \frac{\partial H(u_2, v_2)}{\partial H(u, v)} \right] + H(u_2, v_2) H^*(u_1 + u_2, v_1 + v_2) \frac{\partial H(u_1, v_1)}{\partial H(u, v)}. \tag{6.46}$$

The implementation of the first two terms is straightforward. However, the third term is difficult to deal with in its present form. Since the KT approximation is used in the algorithm, it is easier to perform a substitution first before evaluating the differentiation at different possibilities of (u, v). For plane-10, $u_2 = u_1 + \Delta$ and $v_2 = v_1$, and the quantity $H^{(3)}(u_1, v_1, u_2, v_2)$ is written as:

$$H^{(3)}(u_1, v_1, u_1 + \Delta, v_1) = H(u_1, v_1) H(u_1 + \Delta, v_1) H(-2u_1 - \Delta, -2v_1). \tag{6.47}$$

Differentiating this with respect to $H(u, v)$ and evaluating the expression at $(u, v) = (u_1, v_1)$, $(u, v) = (u_1 + \Delta, v_1)$ and $(u, v) = (-2u_1 - \Delta, -2v_1)$ give:

$$\frac{\partial H^{(3)}}{\partial H} = H(u + \Delta, v) H^*(2u + \Delta, 2v) + H(u - \Delta, v) H^*(2u - \Delta, 2v)$$
$$+ H^*\left(\frac{u+\Delta}{2}, \frac{v}{2}\right) H^*\left(\frac{u-\Delta}{2}, \frac{v}{2}\right). \tag{6.48}$$

Similarly, for plane-01, $u_2 = u_1$ and $v_2 = v_1 + \Delta$, and the quantity $H^{(3)}(u_1, v_1, u_2, v_2)$ is:

$$H^{(3)}(u_1, v_1, u_1, v_1 + \Delta) = H(u_1, v_1) H(u_1, v_1 + \Delta) H(-2u_1, -2v_1 - \Delta). \tag{6.49}$$

Following the same differentiation and evaluation process, the result is:

$$\frac{\partial H^{(3)}}{\partial H} = H(u, v + \Delta) H^*(2u, 2v + \Delta) + H(u, v - \Delta) H^*(2u, 2v - \Delta)$$
$$+ H^*\left(\frac{u}{2}, \frac{v+\Delta}{2}\right) H^*\left(\frac{u}{2}, \frac{v-\Delta}{2}\right). \tag{6.50}$$

After the gradient calculations have been performed in the frequency domain, the product $\partial E/\partial H$ is taken to the spatial domain using the inverse Fourier transform. This is then incorporated with the gradient calculation based on the Richardson-Lucy (RL) framework to update the PSF. It should be noted that the additive form of the RL method is used.

6.5.5   Feasibility investigation

The simulations to test the feasibility of using the bispectrum/KT prior for blind deconvolution are as follows. The binary object has a brightness ratio of 1.25 and a separation of 3.61 pixels. Using a 1-metre telescope imaging at 650nm means that this separation corresponds to 0.59 arcsec. 500 phase screens of size 16×16 and $D/r_0 = 4$ are generated using the method of Harding et al. [76]. The data is assumed to have Poisson noise statistics, and the photon level of each data frame is set to either 5000, 2000 or 500. The noiseless case is also investigated. A total of 100 trials are performed for each set of parameters. The first 200 phase screens are used for trials with 2 frames of data each, while all the phase screens are used for the case with 5 frames per trial.

The results of the proposed method are compared to those obtained using Schulz's phase parameterisation and penalised maximum likelihood technique (ML+). The results are tabulated in Tables 6.3, 6.4 and 6.5. The success rate describes the percentage of the reconstructions where a binary star is formed. The other parameters are calculated using the successful trials. For Schulz's penalised maximum likelihood method and the proposed method, the weighting of the prior information, denoted by β and α respectively, are optimised to produce the best results.

In general, the success rate for all methods decreases as the amount of Poisson noise increases. The performance with 5 frames of data is consistently better than that with 2 frames of data as expected. Out of the three methods, Schulz's phase parameterisation method produced the most consistent results and the highest success rate. Although the proposed method is comparable to Schulz's ML+ method under the conditions where either no noise or 5000 photons per frame are present, its performance deteriorates under low light conditions, and it is preferrable to use Schulz's ML+ method.

No. photons	Success rate	Ratio mean	Ratio variance	Separation mean	Separation variance
2 frames					
Noiseless	97	1.25	0.0018	3.59	9.33×10^{-4}
5000	96	1.27	0.0102	3.60	0.0010
2000	87	1.32	0.0286	3.61	7.24×10^{-4}
500	41	1.32	0.0018	3.59	0.0070
5 frames					
Noiseless	100	1.24	8.69×10^{-4}	3.59	3.89×10^{-4}
5000	100	1.26	0.0051	3.59	5.72×10^{-4}
2000	94	1.33	0.0175	3.60	8.02×10^{-4}
500	70	1.44	0.103	3.60	0.0015

Table 6.3: *Results using Schulz's phase parameterisation method. The binary star has an intensity ratio of 1.25 and a separation of 3.61 pixels.*

No. photons	Success rate	Ratio mean	Ratio variance	Separation mean	Separation variance	β
2 frames						
Noiseless	90	1.26	0.0076	3.58	0.0024	0.2
5000	84	1.34	0.0530	3.59	0.0044	0.2
2000	53	1.44	0.177	3.57	0.0158	0.2
500	27	1.33	0.240	3.53	0.0356	0.8
5 frames						
Noiseless	96	1.29	0.0387	3.57	0.0086	0.2
5000	88	1.37	0.0632	3.57	0.0037	0.2
2000	74	1.30	0.112	3.55	0.0137	0.2
500	43	1.17	0.131	3.60	0.0353	0.8

Table 6.4: *Results using Schulz's penalised maximum likelihood method. β takes on a value between 0 and 1, and is used to determine the weighting of the prior term. The binary star has an intensity ratio of 1.25 and a separation of 3.61 pixels.*

No. photons	Success rate	Ratio mean	Ratio variance	Separation mean	Separation variance	α
2 frames						
Noiseless	98	1.26	4.39×10^{-4}	3.60	2.92×10^{-4}	0.001
5000	81	1.42	0.0153	3.61	1.50×10^{-4}	0.001
2000	30	1.49	0.0968	3.59	0.0134	0.1
500	1	0.71	-	3.14	-	0.1
5 frames						
Noiseless	100	1.25	7.69×10^{-5}	3.60	5.46×10^{-5}	0.001
5000	92	1.47	0.0104	3.60	1.67×10^{-4}	0.01
2000	46	1.52	0.137	3.61	0.0043	0.1
500	3	1.10	0.361	3.60	0.0142	0.1

Table 6.5: *Results of the proposed method using the phase closure property as prior information. α is the weighting of the prior term. The binary star has an intensity ratio of 1.25 and a separation of 3.61 pixels.*

6.5.6 Summary and future work

This section investigated the feasibility of applying a constraint on the PSF using the Knox-Thompson (a subset of the bispectrum) phase in a blind deconvolution framework. On comparing this method with Schulz's phase parameterisation and penalised maximum likelihood algorithms, the proposed method has worse performance for blind deconvolution. One of the main reasons for this is that the KT prior is a statistical prior that relies on ensemble statistics, and the minimum number of frames of data required to make this prior valid is not known. From the simulation results, using 5 frames of data is obviously insufficient. Nevertheless, the performance of using 5 frames is better than the performance using 2 frames, illustrated by the increased success rate in Table 6.5. This indicates that as the number of frames increases, the performance of the prior improves. Although simulations have not been performed on more frames of data, from the preliminary results, it is shown that the KT prior can be potentially useful if a larger ensemble of frames is used.

6.6 Summary

In this chapter, a comprehensive review on iterative blind deconvolution was presented. The general MAP framework adopted by most algorithms was discussed, and the common

forms of prior information used in these algorithms were listed.

The feasibility of two forms of prior information on the PSF was investigated: the application of Kolmogorov statistics, and the use of the phase closure property. The former has produced some promising results, and a further investigation is worthwhile. Although the latter does not produce comparable results to existing algorithms when a small number of frames is processed, the performance does improve with a greater ensemble, indicating that the statistical prior requires more data frames to be able to constrain the solution effectively.

The blind deconvolution algorithms presented in this chapter deal with the case where the speckle images are the only data captured. If more information is available, such as simultaneous wavefront sensing data, it can be used to provide a stronger constraint on the solution. This is discussed in the following chapters.

Chapter 7

Centroid estimation from multiple undersampled images

As discussed in Chapter 6, the goal in astronomical imaging is to reconstruct the object from the measured images at the telescope. Since the dominant effect of atmospheric turbulence on the object is a phase distortion, this distortion must be estimated for the reconstruction of the object.

While the technique of blind deconvolution is used to estimate the point spread function and the object simultaneously when only speckle images are captured, the availability of wavefront sensing data allows the phase to be estimated differently. In this case, the astronomical imaging problem can be viewed as a two-level inverse problem involving the phase estimation from wavefront sensing data and ultimately the reconstruction of the object.

To understand the advantage of using wavefront sensing data in phase estimation, consider the case where the phase is to be estimated from the speckle images. Theoretically, the mean tip and tilt components of this distortion can be estimated from the centroid of a speckle image. Although this removes most of the energy introduced by the phase distortion, the resolution required is not achieved because the higher order Zernike coefficients cannot be estimated. Moreover, because the speckle image is captured by a photon-counting camera, the resultant image is very noisy (Fig. 4.18(a)) and the error in the tip and tilt estimation is large. The estimation of the phase from a speckle image is therefore underdetermined

and ill-conditioned, and is a very difficult problem.

In order to improve the conditioning of the problem, more measurements are required
to estimate the unknowns. This is achieved by subdividing the aperture and is physically
realised by a Shack-Hartmann wavefront sensor (SH-WFS). A SH-WFS is a two-dimensional
array of lenslets across the aperture. Each lenslet produces an image, which is centroided to
produce the local tip and tilt components. These local measurements are then combined to
estimate the overall wavefront phase distortion. By increasing the number of measurements,
the problem is better conditioned and a solution can be found more easily.

In this chapter, the primary focus is on the inverse problem of reconstructing the phase from
wavefront sensing measurements. Specifically, the accuracy of the centroid calculation is
examined as it directly affects the quality of the phase reconstruction. Since the commonly
used method which calculates the centre of gravity is not optimal [174], a new method
for centroid calculation is proposed and is realised via a multiframe blind deconvolution
framework.

The theory of reconstructing the wavefront from centroid measurements is not required in
the understanding of this chapter and is therefore deferred until Chapter 8 where it is more
appropriate.

7.1   Centroid estimation using a Shack-Hartmann sensor

Phase estimation using a Shack-Hartmann wavefront sensor (SH-WFS) is a common ap-
proach as the SH-WFS can be used for both compact and extended objects. In a Shack-
Hartmann sensor the aperture is subdivided into an array of lenslets (Fig. 7.1(a)), producing
an array of spots that are usually captured on a CCD camera in an open-loop system. Fig.
7.1(b) shows a typical data frame captured by a SH sensor. Not all the spots are formed by
fully illuminated lenslets due to the geometry of the aperture. The displacement of a spot
from the point corresponding to the centre of the lenslet is used to produce the average
slope of the wavefront over the respective lenslet. These local slopes are then combined to
reconstruct the phase across the entire aperture.

For a single lenslet, the mean slope is directly proportional to the displacement of the spot

(a)   (b)

(c)
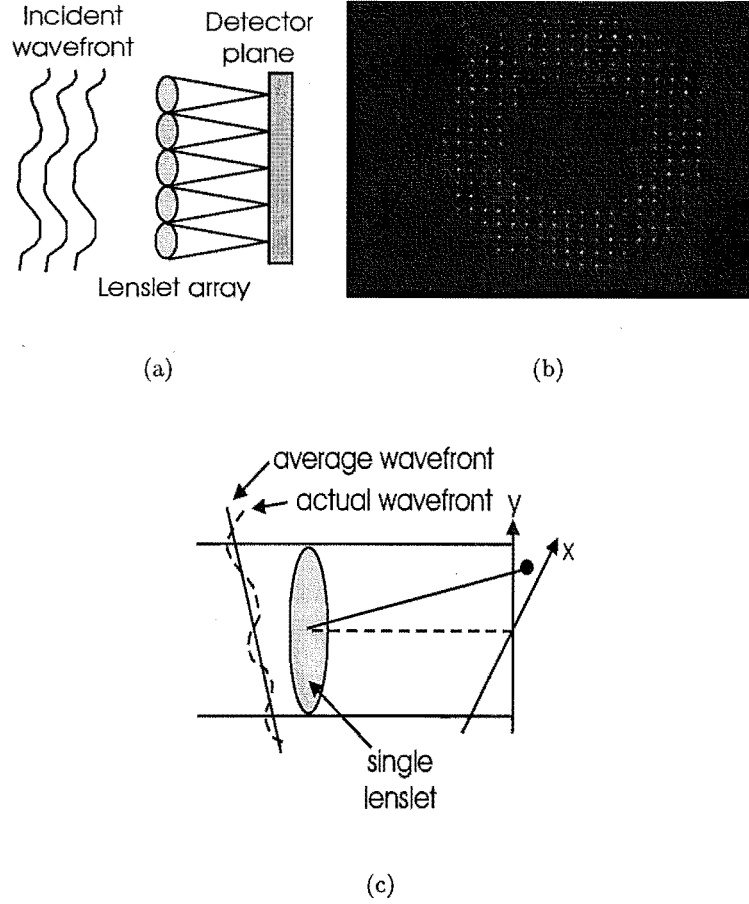
Figure 7.1: (a) An illustration of the Shack-Hartmann lenslet array. (b) A typical frame of wavefront sensing data captured at the Observatoire de Lyon. (c) The relationship between the slope of the wavefront and the displacement of the spot on the detector plane over a single lenslet.

in the focal plane (Fig. 7.1(c)). The estimation of a spot's displacement is performed by calculating the centre of gravity (or centroid) over a fixed window. For a data frame $d(x,y)$, the centroids in the $x-$ and $y-$ directions, $\bar{x}$ and $\bar{y}$, are calculated by [101]:

$$\bar{x} = \frac{\sum xd(x,y)}{\sum d(x,y)}$$
$$\bar{y} = \frac{\sum yd(x,y)}{\sum d(x,y)}. \tag{7.1}$$

Although this is the most common centroid estimator, it has many associated problems; these are discussed in the next subsection.

```
0.0860  0.0818  0.0309  0.0695  0.0980

0.0854  0.0660  0.0838  0.0621  0.0271

0.0594  0.2968  0.3068  0.0795  0.0252

0.0497  0.1897  0.2541  0.0957  0.0876

0.0900  0.0341  0.0703  0.0523  0.0737
```

Figure 7.2: *Sample data frame to illustrate the problems with the centroiding operation.*

7.1.1   Problems with the traditional centroid estimator

There are three principal problems in the traditional estimation of the centroid. The first is a bias introduced by the finite pixel size of the CCD cell, and the finite window size used. As an example, consider the estimation of the centroid of the data shown in Fig. 7.2. The centre of the array has the largest numerical value of 0.3068 and if this is used as the centre of a 3×3 window, indicated by a solid line in Fig. 7.2, the centroid is $(3.2284, 2.7803)$ corresponding to the row and column values respectively. If a 5×5 window is used (area enclosed by the dashed line), the centroid is $(3.1062, 2.8261)$. If the 3×3 window is instead centred around 0.2968 (shaded area) the result is $(3.1856, 2.3235)$. The inconsistency of the centroid position illustrates the problem caused by finite pixel and window sizes. These problems can be ameliorated by using a very large window with very small pixels to completely encompass and finely sample the spot.

It should be pointed out that this first problem is unique to the use of Shack-Hartmann wavefront sensing data in an open-loop. In a closed-loop adaptive optics system, the motion of the spot position between consecutive frames is small; it is thus simpler to determine the location of an appropriate centroiding mask. As a consequence, for closed-loop operation it is usual to use one quad-cell per sub-aperture, and the spot is approximately centred on the intersection of the four elements of the quad-cell by the adaptive optics. However, in an open-loop situation, the motion introduced by uncompensated atmospheric turbulence is much larger and it is necessary to employ larger arrays of detectors to locate the centroid.

The second problem is readout noise in the detector. Since the effect of the noise increases with a decreasing pixel size, it is desirable to use large pixels. As a result of the conflict-

ing requirements to overcome the first two problems, a compromise must be reached for determining the size and number of pixels used. This is particularly relevant for open-loop operation, where the spot position is ill-defined.

The final problem arises when compensating for the bias introduced by read noise. The CCD has a mean level reading which must be subtracted before the centroid is calculated; otherwise the resulting centroid estimate is biased towards zero. After this post-processing step, it is possible for the images to have negative values. The presence of negative values again affects the centroid calculation. In fact, it has been shown [174] that even in the absence of read noise, centroiding is sub-optimal and a superior performance can be obtained by a maximum *a posteriori* (MAP) approach.

In view of the practical and theoretical problems of centroiding, and the conflicting requirements for overcoming these problems, a model-fitting procedure is proposed as opposed to the traditional centroiding operation. This means that if a model of the spot is available, it can be used to identify the location of a spot in a noisy data frame to subpixel accuracy. Unlike traditional centroiding, this method does not have any of the associated problems discussed above.

A model-fitting procedure is used instead of cross-correlation because the data frames are undersampled. Since some high-frequency components have aliased into the low-frequency spectrum, Parseval's theorem no longer holds and the energy in the data spot and the model spot cannot be normalised against each other. In this case, the operation of cross-correlation produces incorrect results. This becomes apparent if we examine an object $f(x, y)$ and the model $m(x, y)$. The maximisation of the cross-correlation between these functions is equivalent to minimising the mean-squared error (MSE) $E\{(m(x, y) - f(x, y))^2\}$. The expansion of this expression leads to

$$E\{m(x, y)^2 - 2m(x, y)f(x, y) + f(x, y)^2\}. \qquad (7.2)$$

While the first term is a constant, and the second term represents the cross-correlation, the last term is not a constant when the data is undersampled. Since the minimisation of the MSE over a varying function is not meaningful, the cross-correlation operation cannot be applied.

There are also other approaches to the problem of centroid estimation, one of which utilises a tilt transition probability function estimated from reference star measurements as a form

(a)   (b)   (c)
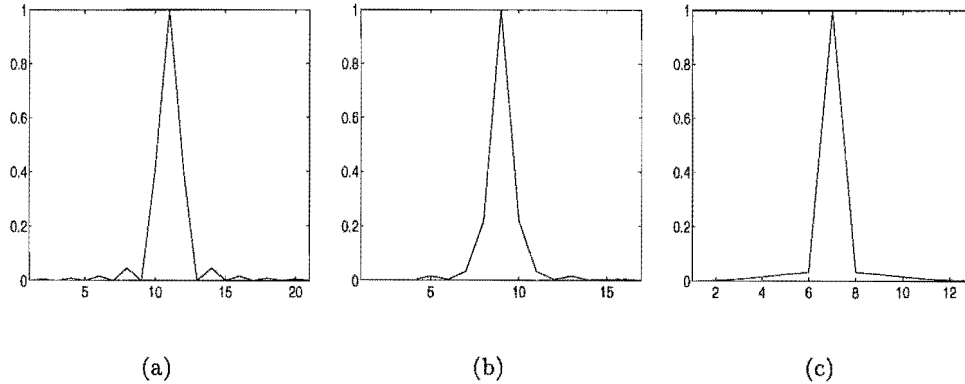
Figure 7.3: *The aliasing of a spot. (a) Cross-section of a diffraction-limited spot. (b) The spot sampled at 80% of the Nyquist rate. (c) The spot sampled at 60% of the Nyquist rate.*

of prior information to predict the tilt between consecutive frames of data [16]. This has been shown to produce better results than the traditional centroid estimator, but requires reference star measurements, which are not needed in the proposed framework. Sallberg *et al.* [151] used a MAP framework for centroid estimation, incorporating prior information such as the light level, the relative spot size, and statistical knowledge of the slopes over the subapertures. The proposed approach differs from this in the sense that the spot size is retrieved directly from the data rather than estimated, and the undersampling and non-uniform sensitivity of the CCD are taken into account.

Another consequence of the data frames being undersampled is that the spot is severely aliased. The aliasing of the spot is illustrated in Fig. 7.3. Fig. 7.3(a) shows the cross-section of a Nyquist sampled spot using a rectangular aperture. It is only shown up to 30% of the diffraction limit for clarity. Fig. 7.3(b) illustrates the effect of sampling at 80% of the Nyquist rate. The first sidelobe is no longer visible, but some high frequency information is preserved. Once again, it is shown up to 30% of the cutoff frequency. Fig. 7.3(c) is the corresponding spot sampled at 60% of the Nyquist rate. It is apparent that the diffraction-limited spot structure is lost. As a result, the model sought must be found at a higher resolution initially. In the model-fitting procedure, this high-resolution model is then downsampled according to the optics and the CCD sampling.

7.1.2 Obtaining a high-resolution model

One way of obtaining a high-resolution model is to use multiple frames of low-resolution data that are subpixelly displaced with respect to each other, and interlace them to produce the high-resolution model. This concept was first realised algorithmically by Huang and Tsay [82], who demonstrated the ability to reconstruct one high-resolution object from several downsampled noisy images based on the spatial aliasing effect. Other analyses on first and higher order aliasing elimination have also been performed, demonstrating the feasibility of such an approach [22,109].

There now exist many techniques which utilise the concept of uncontrolled microscanning where random subpixel shifts between data frames are exploited for resolution improvement [2,57,73,153]. Gillette et al. [57] used a procedure where the shifts between the data frames are estimated first, followed by the reconstruction of the high-resolution object. Schuler and Scribner [153] estimated the registrations based on the concept of optical flow, and restored the object using a Wiener filter. Alam et al. [2] also employed optical flow in registration estimation, but restored the object with a modified EM algorithm. Although this removes the need for recursive registration estimation, the errors in estimating the shifts result in errors in the object reconstruction. As a result, this approach is not adopted in this thesis.

Gerwe et al. [55] and Green and Hunt [65] utilised a slightly different multiframe approach called supersampling. The key point is that the reconstructions are upsampled onto a finer grid than the measured data, allowing for a 'greater degree of superresolution' [65]. However, the data frames are not registered with respect to each other, and it is therefore not appropriate for the application to the problem at hand.

One of the best methods to date which employ uncontrolled microscanning is that proposed by Hardie et al. [73], who used an iterative MAP framework for jointly estimating the registrations and the high-resolution object. This removes the danger of error propagtion present in some of the other algorithms. However, one drawback of this algorithm is its reliance on the knowledge of the PSF of the optics, which is not always available. In this chapter, an improvement to this existing algorithm is proposed and is applied to wavefront sensing data. The PSF of the optics, as well as the CCD structure are also estimated. This removes the need for making an assumption about the PSF, which can also lead to errors in the reconstruction.

In Hardie's original method, the displacements between the low-resolution images are provided by the uncontrolled motion of the camera mounted on an aeroplane. In astronomical imaging, the dominant effect of atmospheric turbulence over individual lenslets is the random motion of the image. In both situations, the low-resolution data frames are randomly displaced with respect to each other. This motion can effectively be used to increase the sampling rate of the data.

The technique discussed in this chapter was developed when dealing with astronomical data; however, it can also be used with other images. The derivation is based on a statistical assumption where the high-resolution object is assumed to be continuous. There exists another family of deterministic algorithms (stemmed from surveying techniques) where different assumptions are employed in the reconstruction [50]. In this thesis, the discussion is restricted to the former set of techniques.

7.2   Improving centroid accuracy: blind deconvolution

The blind deconvolution procedure is derived from Hardie's method where a sequence of images is processed via a MAP framework to simultaneously estimate the underlying object and the registrations between the data frames. When applied to wavefront sensing data, the registrations are equivalent to centroid positions, while the object is analogous to the model we are looking for. In this investigation, the existing algorithm is extended to estimate the point spread function of the optics and the CCD structure.

The observation model is defined as:

$$
\begin{aligned}
d_k(x,y) &= s_k\{f(x,y) \odot h^c(x,y) \odot h^o(x,y)\} + n_k(x,y) \\
&= g_k(x,y) + n_k(x,y).
\end{aligned}
\tag{7.3}
$$

where $f(x,y)$ is the original undegraded object of size $M \times M$, $h^c(x,y)$ is the PSF of the CCD cell structure, and $h^o(x,y)$ is the PSF of the optics. These are referred to as the CCD-PSF and optics-PSF respectively, and are unknown quantities. The notation $s_k\{\cdot\}$ is the subsampling operation for the k^{th} frame and models the undersampling of the CCD. $g_k(x,y)$ is the k^{th} noiseless low-resolution frame of size $N \times N$, and $n_k(x,y)$ represents the k^{th} realisation of the independent and identically distributed (i.i.d.) Gaussian noise samples. $d_k(x,y)$ is the k^{th} measured noisy low-resolution frame. A key point in this model

is that the image is modelled at a resolution higher than the data captured, enabling the algorithm to deal with data sampled beneath the Nyquist rate. It is also assumed that each low-resolution frame can be expressed as a linear weighted sum of the corresponding high-resolution pixels.

Similar to the statistical interpolation scheme in Chapter 5, the ratio R between the size of the high- and low-resolution images is M/N and is an integer greater than 1. However, unlike statistical interpolation, the value of R also determines the allowable step size for the registration process, and is directly linked to the location of the centroid positions of the spot. The larger the value of R, the smaller the step size can be. As a result, R cannot be determined simply by the sampling rate of the data, but is predominantly decided by the desired precision for the centroid estimation.

Although R can be made arbitrarily large, there is an upper limit for a reasonable value for it, due to two reasons. Firstly, the noise on the spot introduces uncertainty to the centroid position, posing as a fundamental limit to the accuracy. Secondly, due to the implementation of a systematic search over an area of R^2 pixels, an increase in R also increases the computational complexity. All these factors should be taken into account when selecting the value of R.

In the blind deconvolution formulation described by Eq. (7.3), the object $f(x, y)$ usually represents a scene. Initially, the problem can be simplified if $f(x, y)$ is a point source. This reduces the equation to:

$$d_k(x, y) = s_k\{h^c(x, y) \odot h^o(x, y)\} + n_k(x, y). \qquad (7.4)$$

Therefore the PSFs and the registrations are to be estimated simultaneously when only $d_k(x, y)$ is measured. This corresponds to the inversion of Eq. (7.4) and can be solved via a MAP framework, defined mathematically as [97]:

$$\max_{\hat{h}^c, \hat{h}^o, \hat{s}_k} P(\hat{h}^c(x, y), \hat{h}^o(x, y), \hat{s}_k | d_k(x, y)). \qquad (7.5)$$

This involves the minimisation of the sum of three error functions: the consistency error E_{con}, and the errors corresponding to the optics-PSF prior (E_o) and the CCD-PSF prior (E_c). These are represented as:

$$E = E_{\mathrm{con}} + E_o + E_c \qquad (7.6)$$

where

$$E_{\text{con}} = \sum_k \sum_{x,y} [\hat{s}_k \{\hat{h}^c(x,y) \odot \hat{h}^o(x,y)\} - d_k(x,y)]^2$$

$$E_o = \lambda_o \sum_{x,y} [\hat{h}^o(x,y) \odot c(x,y)]^2$$

$$E_c = \lambda_c \sum_{x,y} \hat{h}^c(x,y) \ln(\hat{h}^c(x,y)). \tag{7.7}$$

The parameters λ_o and λ_c are the weightings of the optics-PSF prior and the CCD-PSF prior respectively. They determine the balance between the fidelity of the reconstruction to the measured data, and the consistency with the prior information available regarding the unknowns. The kernel function $c(x,y)$ is the Laplacian and is the prior used in the original algorithm by Hardie *et al.* As discussed in Section 3.4, this prior is essentially a high-pass filter and ensures continuity in the reconstruction. A maximum entropy (ME) prior, discussed in Chapter 3, is used for the CCD-PSF which also inherently enforces positivity.

If the atmospheric turbulence is assumed to follow Kolmogorov statistics, then the centroid positions should obey a Gaussian distribution [43, 62]. This is a form of prior information for the registration process, and is used in limiting the search window for the registration parameters. However, no other explicit prior is used as in the case of the PSFs.

It should be pointed out that priors used in the algorithm are reasonable given the physical situation; however, other priors can also be used.

The gradient of E_{con} with respect to the optics-PSF $\hat{h}^o(x,y)$ is

$$\frac{\partial E_{\text{con}}}{\partial \hat{h}^o(x,y)} = \sum_k 2[\hat{s}_k \{\hat{h}^o(x,y) \odot \hat{h}^c(x,y)\} - d_k(x,y)] \otimes \hat{h}^c(x,y) \tag{7.8}$$

where $\otimes$ denotes the correlation operation. Its minimisation with respect to the CCD-PSF $\hat{h}^c(x,y)$ is obtained from Eq. (7.8) by reversing the roles of $\hat{h}^o(x,y)$ and $\hat{h}^c(x,y)$. The gradients for updating the prior terms are:

$$\frac{\partial E_o}{\partial \hat{h}^o(x,y)} = 2\lambda_o [\hat{h}^o(x,y) \odot c(x,y)] \otimes c(x,y)$$

$$\frac{\partial E_c}{\partial \hat{h}^c(x,y)} = \lambda_c [1 + \ln(\hat{h}^c(x,y))]. \tag{7.9}$$

The method of steepest descent is used to update these variables, the implementation of which is included in Appendix A. Positivity of the optics-PSF is enforced using the method

of Johnston and Lane [94], while the ME prior inherently enforces positivity on the CCD-PSF.

Due to the discrete model used, the minimisation of the error metric with respect to the registration parameters is achieved by a systematic search instead of a gradient calculation. The search is performed on a per-frame basis, and a shift of up to five low-resolution pixels is allowed. The set of registrations that produces the lowest error corresponds to the best estimate of the registrations for the next iteration.

7.3   Simulation results

In this section, simulations using a point source as an object are performed to demonstrate the feasibility of the technique, as well as retrieving models of the optics-PSF and CCD structure. The procedure is then extended to the blind deconvolution of a more complicated astronomical object, and is demonstrated on real data in the next section.

The data used for the simulation is generated to approximate the real data. A 20×20 Shack-Hartmann lenslet array is used, and the telescope is assumed to have a diameter of 1-metre. The ideal optics-PSF was generated assuming a constant magnitude at the aperture and is shown in Fig. 7.4(a). The effect of the atmosphere within the lenslets was modelled by 1000 phase screens with Kolmogorov statistics [76] and $D/r_0 = 1$. This value of D/r_0 is chosen as it represents the traditional region of operation for a Shack-Hartmann wavefront sensor. The variance of the mean slope per lenslet has a value of 0.27 when expressed in terms of pixels, which is again representative of the real data.

The data frames were generated from 9×9 phase screens, and were sampled at 60% of the Nyquist rate, resulting in simulated data in an 11×11 pixel array. The centre portion of size 5×5 is used for blind deconvolution to be compatible with the dimensions of the real data used. Fig. 7.4(b) shows the CCD-PSF, which is assumed to be a Gaussian-shaped function. Although the use of a Gaussian can cause ambiguities in a blind deconvolution solution as discussed in Section 6.3.1, this is chosen here because this function best models the sensitivity of a CCD cell. Moreover, as shown in Section 7.4, the CCD-PSF retrieved from experimental data is approximately Gaussian. Fig. 7.4(c) illustrates the centre portion of a typical data frame. It can be seen that it is severely aliased.
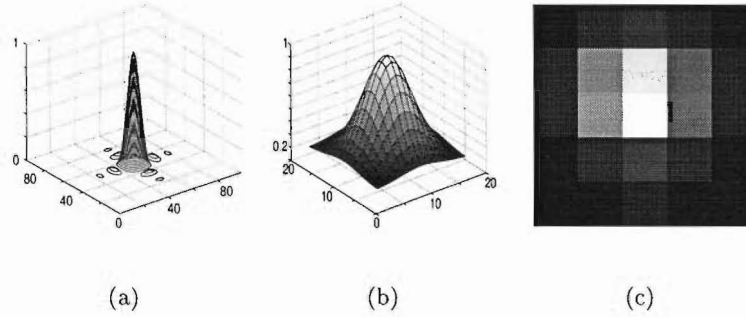
(a)   (b)   (c)

Figure 7.4: *Simulation data. (a) Optics-PSF. (b) CCD-PSF. (c) Centre portion of a typical simulated data frame with no noise. Note that the sampling rate in (a) and (b) is 9 times that of (c). The data frame has a range between 0 and 1.*

The initial optics-PSF estimate was generated by averaging the data frames and then convolving with a uniform array. The convolution operation is used to provide a smoother starting estimate. The initial CCD-PSF estimate is uniform. The initial displacements were set to zero. The weightings of $\lambda_o = 0.01$ and $\lambda_c = 0.001$ were found to be appropriate for this simulation. A value of $R = 9$ is chosen here, meaning that the centroid estimation is accurate to one-ninth of a low-resolution pixel. It should be pointed out that since the data is sampled at 60% of the Nyquist rate, a value of $R = 9$ means that the reconstruction is allowed to have more than 5 times the frequency content of a Nyquist sampled image. This is however not true in this case because the object to be recovered is a diffraction-limited spot, the frequency content of which is limited by the diffraction limit. Sampling more finely than the Nyquist rate does not improve the resolution of a diffraction-limited spot. Even in the case where the object is not a diffraction-limited pattern, the theoretical resolution improvement is not achievable in reality due to the noise inherent to the captured data.

It is assumed that the wavefront sensing data was captured with sufficient light, and the dominant noise component is read noise. As a consequence, in the simulations, the noiseless case and the cases where 30dB and 15dB white Gaussian read noise is present are investigated. This assumption reflects the noise on the real data captured.

Figs. 7.5(a) and 7.5(b) show the reconstructed optics-PSF and CCD-PSF in the noiseless case; Figs. 7.5(c) and 7.5(d) show the reconstructions in the presence of 30dB noise, and Figs. 7.5(e) and 7.5(f) show the reconstructions when 15dB noise is present. The optics-

PSF and the CCD-PSF reconstructed well in all cases, although some degradation of the sidelobe estimate is evident in the optics-PSF.

(a)	(b)
(c)	(d)
(e)	(f)
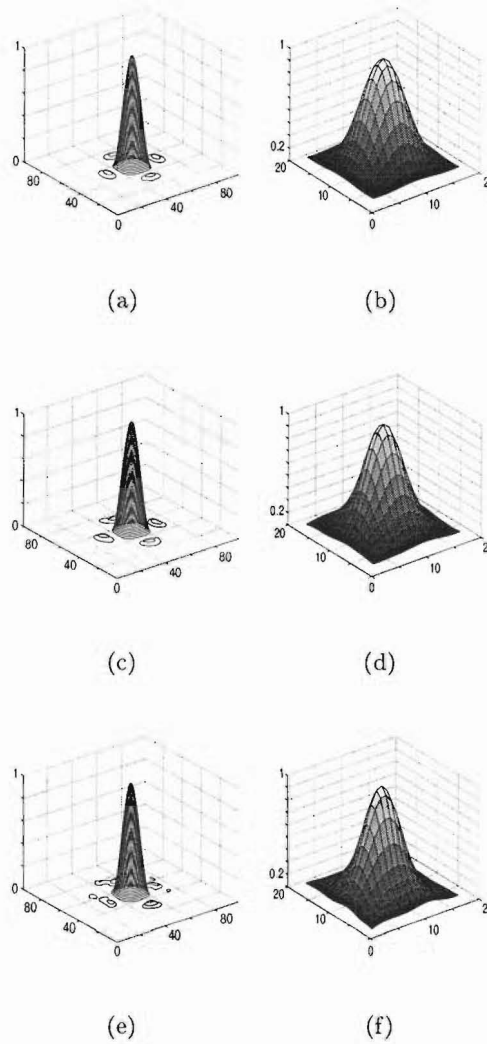
Figure 7.5: *Simulation results. (a) Noiseless case: optics-PSF. (b) Noiseless case: CCD-PSF. (c) 30dB noise: optics-PSF. (d) 30dB noise: CCD-PSF. (e) 15dB noise: optics-PSF. (f) 15dB noise: CCD-PSF.*

The centroid positions are calculated using two methods: firstly by traditional centroiding, where window sizes of 3×3 and 5×5 are both tested, and secondly by blind deconvolution. The per-frame mean-squared error (MSE) values and the variances of the error distribution are tabulated in Table 7.1. All the methods produced unbiased zero-mean distributions for

the centroid positions.

Several observations can be made from these results. When no noise is present, using a 5×5 window for centroiding has better performance than using a 3×3 window. However, the converse is true in the presence of noise, illustrating the detrimental effect of noise with a larger window size. In the noiseless case, blind deconvolution did not produce the best result. This is because the blind deconvolution algorithm estimates the displacements in discrete steps, and this can result in a residual error of up to half a step. As stated earlier, this problem can be ameliorated by using a smaller step which corresponds to increasing R, with a tradeoff of higher computational complexity. In the noisy cases where this error is not significant, blind deconvolution has superior performance to the centroiding operation.

Although blind deconvolution is superior to performing centroiding on the data, it is computationally intensive. For a 20×20 lenslet array, there are up to 400 spots that require blind deconvolution and centroid estimation. As a result, the blind deconvolution approach is not practical to be used as a routine procedure. An alternative approach is to use the optics-PSF and CCD-PSF estimated from a set of data, and produce a combined high-resolution model by convolving these quantities. This is expressed as:

$$\hat{h}^{\text{overall}}(x, y) = \hat{h}^o(x, y) \odot \hat{h}^c(x, y). \tag{7.10}$$

This model can then be used to estimate the centroid positions of the remainder of the spots.

The validity of this approach is tested by simulations. The reconstructions of the optics-PSF and CCD-PSF in all cases (Fig. 7.5) are convolved to produce overall high-resolution models. These models were then approximated by Gaussian functions, and the centroid positions, still in steps of one-ninth of a low-resolution pixel, are calculated assuming these Gaussian models. The results are tabulated in Table 7.2. The performance is equivalent to blind deconvolution at a significantly reduced computational requirement.

It should be pointed out that the data frames used in this simulation correspond to the images formed by a fully illuminated lenslet in the lenslet array. When applied to real data, the blind deconvolution procedure enables more complicated models to be developed for partially illuminated lenslets. For a given WFS configuration, the blind deconvolution procedure needs to be performed only once; this means that once the model is retrieved, it can be used for all the frames captured with this particular WFS set-up.

The traditional centroiding operation has been shown to be equivalent to fitting a Gaussian of variable width to the data [174]. However, this can lead to substantial errors especially when dealing with extended objects, which is in fact a common occurrence. By using a fixed Gaussian model, it is ensured that the spot position is identified correctly especially for the case of an extended object. Fig. 7.6(a) illustrates the case for a point source. The solid line is the measured data, and the dash-dotted line is the result of fitting a variable width Gaussian to this data. The location of the spot is slightly biased towards the right-hand side where it is dominated by noise. By using a fixed width Gaussian and performing model-fitting, the correct spot location can be found, shown by a dotted line. Fig. 7.6(b) shows the case for a binary object where the model is made up of two Gaussians. It is apparent that the variable width Gaussian introduces substantial error, and the method of model-fitting gives superior results.

Finally, since the algorithm reconstructs the object spectrum beyond the limit set by the data sampling, it is important to consider the frequency content of the reconstructions. The ensemble average spectra of the initial and reconstructed optics-PSFs are shown in Fig. 7.7 as solid and dash-dotted lines respectively. It can be seen that the frequency content of the reconstructed optics-PSF has been boosted in the region which is normally aliased by the inadequate sampling. It should be pointed out that the high-frequency content of the initial optics-PSF is caused by the blockiness of the data frames, and does not represent meaningful frequency content.

This section has proven the feasibility of the blind deconvolution framework, and has shown that the application of multiple frames of low-resolution data can result in a reconstruction of the object of an improved resolution. The need for the knowledge of the PSF is also removed. So far, only a point source has been used in the simulation which does not demonstrate the full potential of the method. In the following section, this technique is demonstrated on real data including an extended astronomical object.

7.4 Results using experimental data

The experimental data was collected with the Shack-Hartmann wavefront sensor of the SPID instrument [3] attached to a 1-metre telescope. The telescope is located at the Observatoire de Lyon, which suffers from severe turbulence due to its low altitude. Fig. 7.8 shows the

Noise	Centroiding 3×3		Centroiding 5×5		Blind deconvolution	
	MSE	Var	MSE	Var	MSE	Var
Noiseless	1.000×10^{-3}	1.048×10^{-9}	1.252×10^{-4}	1.235×10^{-10}	2.593×10^{-4}	2.593×10^{-10}
30dB	7.400×10^{-3}	7.441×10^{-9}	4.250×10^{-2}	4.255×10^{-8}	5.100×10^{-3}	5.147×10^{-9}
15dB	3.380×10^{-2}	3.380×10^{-8}	2.500×10^{-1}	2.501×10^{-7}	2.260×10^{-2}	2.145×10^{-8}

Table 7.1: Results demonstrating the accuracy of the centroiding methods. Traditional centroiding are performed using a 3×3 and a 5×5 window. These results are compared to the blind deconvolution method. MSE = mean-squared error per frame in wavefront sensing camera pixels. Var = variance of the error distribution.

Noise	Model-fitting	
	MSE	Var
Noiseless	0	0
30dB	5.000×10^{-3}	5.004×10^{-9}
15dB	2.180×10^{-2}	2.172×10^{-8}

Table 7.2: Results of calculating centroids with model-fitting. MSE = mean-squared error per frame in wavefront sensing camera pixels. Var = variance of the error distribution.

(a)   (b)
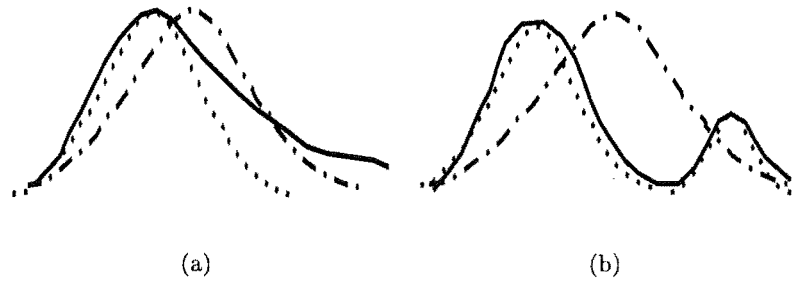
Figure 7.6: *The demonstration of the difference between using a variable and a fixed width Gaussian in locating a spot for (a) a point source and (b) a binary object. In each case, the solid line represents the measured data. The dash-dotted and dotted lines correspond to fitting a variable and a fixed width Gaussian to the data respectively.*

Normalised frequency   0.25 fs
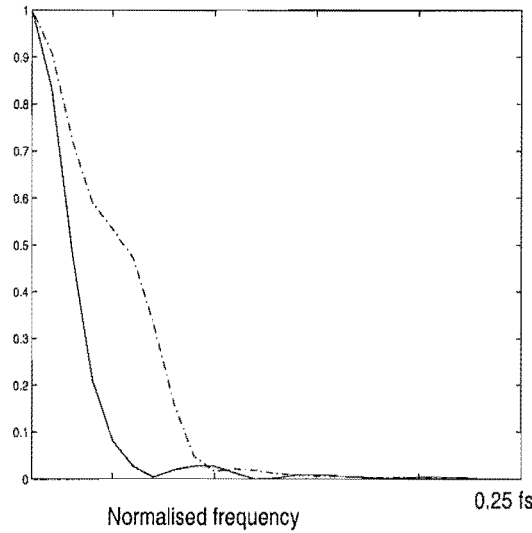
Figure 7.7: *The frequency content of the optics-PSF reconstruction has been boosted. Solid: spectrum of first optics-PSF estimate. Dash-dotted: spectrum of final optics-PSF estimate.*

cross-section of the optical transfer function (OTF) of the telescope. There is a significant attenuation of the middle frequency range due to the large secondary mirror.
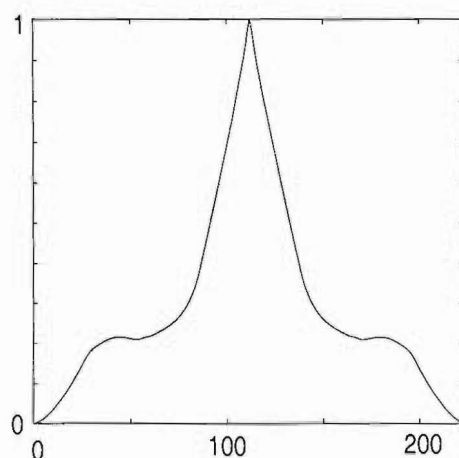
Figure 7.8: *Cross-section of the optical transfer function of the telescope.*
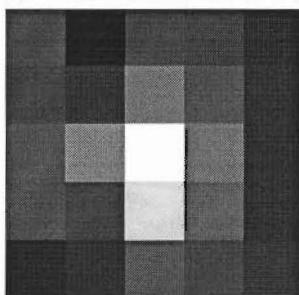
Figure 7.9: *A typical spot captured by the wavefront sensing camera. A grayscale is used where a white pixel has the highest value, and a black pixel has the smallest value.*

7.4.1   Reconstructing a point source

The object of observation (Pollux) is an unresolvable star of magnitude 1.2 imaged at 690nm and the exposure time of 2ms was chosen to ensure the turbulence was effectively frozen. The 20×20 array of lenslets was used to form wavefront sensing images. A typical image of a spot is shown in Fig. 7.9.

As with the simulated case, the initial estimation of the optics-PSF is the average of the data frames smoothed by a uniform array, while the first CCD-PSF estimate is uniform. The initial displacements are set to zero. 1000 frames are used in the reconstruction. In

(a)   (b)   (c)   (d)
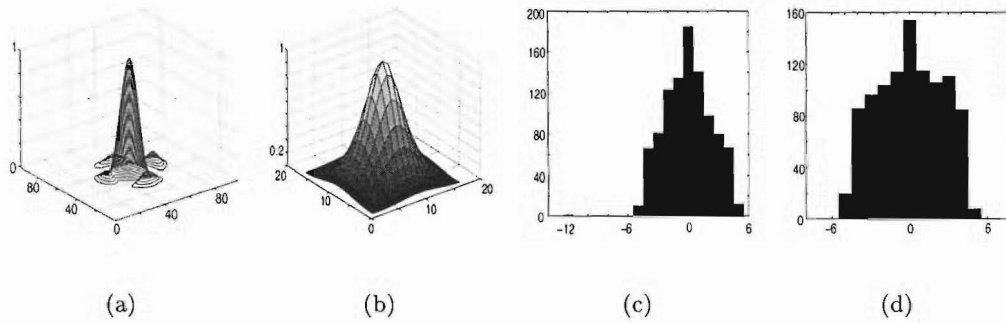
Figure 7.10: *Results using wavefront sensing data of Pollux captured at the Observatoire de Lyon. (a) Optics-PSF. (b) CCD-PSF. (c,d) Histogram of the row and column displacements. Each integer value along the horizontal axis represents one-ninth of a pixel of the CCD camera.*

order to reduce the computational load, each spot is 'centred' initially so that the pixel with the highest value is the middle pixel in the array. It should be pointed out that in many cases, there are two or more pixels that were of very similar values. The value of R is once again set to 9, while the weightings λ_o and λ_c are set to 0.05 and 0.005 respectively. These values are found to be appropriate for this set of data.

Fig. 7.10(a) shows the reconstruction of the optics-PSF where a diffraction pattern of a rectangular aperture is clearly visible. This is comparable to the ideal optics-PSF shown in Fig. 7.4(a). The fact that the sidelobes are higher is attributable to the higher order aberrations inherent in the WFS used to capture the data. Fig. 7.10(b) shows the reconstruction of the CCD-PSF. The magnitude is not uniform as assumed in many algorithms: there is a peak at the centre, with a decrease in magnitude towards the edges. Fig. 7.10(c) and 7.10(d) show the histograms of the row and column displacements. Each integer value along the horizontal axis represents one-ninth of a pixel of the CCD camera. These histograms show approximately Gaussian distributions, which are reasonable considering that the atmospheric turbulence can be modelled by Kolmogorov statistics, whose centroid positions are known to follow a Gaussian distribution [43, 62].

In the case presented here, the object was a point source, and the optics-PSF and CCD-PSF are reconstructed as the diffraction-limited pattern of a square aperture and a Gaussian-shaped function respectively. Although this substantiates the feasibility of the blind deconvolution algorithm on real point source data, it is often more common (and interesting) to

observe extended objects. The extension of the blind deconvolution procedure to the binary star Castor is now examined.

7.4.2   Reconstructing a binary star

The wavefront sensing images of the binary star were captured using a 14×14 Shack-Hartmann lenslet array with an exposure time of 3.2ms. In the blind deconvolution procedure, the initial estimate of the optics-PSF $h^o(x, y)$ is a diffraction-limited pattern of a square aperture. The sampling rate is chosen to match that of this particular SH array configuration, where the images are sampled at 62.5% of the Nyquist rate. The initial estimate of the CCD-PSF $h^c(x, y)$ is the same as that previously retrieved with the Pollux data. These two quantities are then convolved to provide the overall-PSF of the system as in Eq. (7.10), and Eq. (7.3) can be rewritten as:

$$d_k(x, y) = s_k\{\hat{f}(x, y) \odot \hat{h}^{\text{overall}}(x, y)\} + n_k(x, y). \tag{7.11}$$

Both $\hat{f}(x, y)$ and $\hat{h}^{\text{overall}}(x, y)$ were then updated in the full blind deconvolution procedure where the former represents the stellar object of interest. Because the first estimate of the overall-PSF is good, there was little change to it throughout the iterative procedure.

Fig. 7.11(a) shows a typical frame of the WFS data. Fig. 7.11(b) shows the reconstruction of $f(x, y)$ which is clearly a binary object. The intensity of the brighter star is 2.4 times that of the dimmer star which is comparable to the widely accepted value of 2.1. The separation of the reconstructed binary star is found to be 145 pixels. When a conversion of 0.0248 arcseconds per pixel (calculated from the system parameters including the telescope diameter, the focal length and the wavelength of the light) is applied, the separation is found to be 3.6 arcseconds. Although it is reasonable compared to the generally accepted value of 3.1 arcseconds, the accuracy of the calculated value can be improved by calibrating the image capturing system with a known binary instead of using the system parameters for a theoretical calculation.

Fig. 7.11(c) and 7.11(d) display histograms of the row and column displacements. Each integer value along the horizontal axis represents one-ninth of a pixel of the CCD camera. The pre-processing on the data of Pollux, where the highest-valued pixel was shifted to be the centre pixel of each frame, was not performed on the Castor frames initially. As a result,

(a)   (b)   (c)   (d)
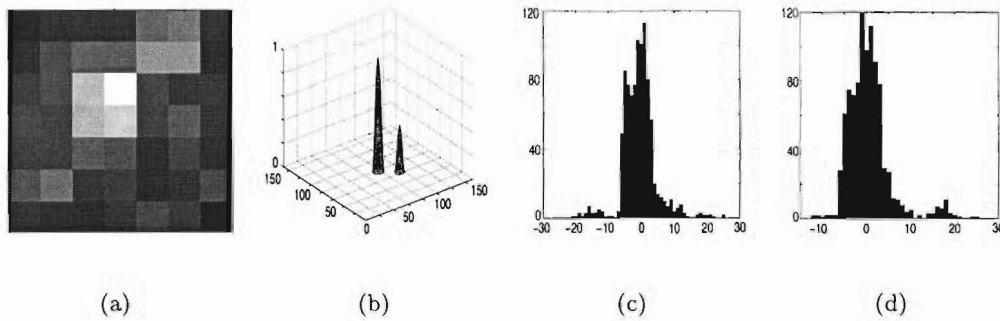
Figure 7.11: *Results using wavefront sensing data of Castor. (a) Typical data frame. (b) Reconstruction of the binary object. (c,d) Histogram of the row and column displacements. Each integer value along the horizontal axis represents one-ninth of a pixel of the CCD camera.*

the estimated displacements have a larger range than those of Pollux. It should be noted that whether the highest-valued pixel was shifted to be the centre pixel has not affected the reconstruction of the model in the blind deconvolution algorithm. Therefore, the fact that no explicit prior is used for the displacements is justified.

7.5 Demonstration of improvement by phase and object reconstruction

Ultimately the objective of improving the centroid estimates is to provide better wavefront sensing reconstructions. The wavefront sensing data used to demonstrate this was the same as that described in Section 7.4 (collected by observing Pollux and Castor), but in this case, images of the entire lenslet array are used. The centroids of the wavefront sensing data are calculated using two methods. Firstly, the model extracted from blind deconvolution in Section 7.4 is applied. This involves the downsampling of the high-resolution model retrieved with blind deconvolution, and identifying its occurrence in the noisy data frames to subpixel accuracy. The shifts are not restricted to a fixed step size here. Secondly, the method of traditional centroiding is used. In each case, the phase is reconstructed using 200 Zernike polynomials [147] and the object distribution is reconstructed. Here, only the end result is examined for comparison; the theory regarding the reconstruction of the phase from wavefront sensing data is detailed in Chapter 8.

(a)   (b)

Figure 7.12: *Reconstructions from wavefront sensing data of Pollux with a 20 × 20 Shack-Hartmann lenslet array. (a) Cross-section of the image using model-fitting. (b) Phase reconstruction using model-fitting. The circle denotes 24% of the diffraction limit.*

The phase reconstruction of Pollux using model-fitting has a Strehl ratio of 41%, compared to a Strehl ratio of 35% when traditonal centroiding is used. The cross-section of the reconstructed image using model-fitting is shown in Fig. 7.12(a), while the phase reconstruction is shown in Fig. 7.12(b) with the circle denoting 24% of the diffraction limit of the telescope. The reconstructions using traditional centroiding are not shown here because there is little visual difference from Fig. 7.12(a) and 7.12(b).

For a point source, there is only a slight improvement in performance, because using model-fitting to find the centroids provides only a small advantage. However, for an extended object like Castor, the improvement in performance is significant. Fig. 7.13(a) shows the object reconstruction using traditional centroiding. Comparing this to Fig. 7.13(b), where the object is reconstructed using model-fitting, it can be seen that the latter is clearly a binary object (with a separation of 3.5 arcseconds), while the former reconstruction is ambiguous. Fig. 7.13(c) and 7.13(d) show the phase reconstructions of the two methods respectively. The circle denotes 24% of the diffraction limit. The stripes in the centre portion in Fig. 7.13(d) are more uniform than those in Fig. 7.13(c), demonstrating a better phase reconstruction.
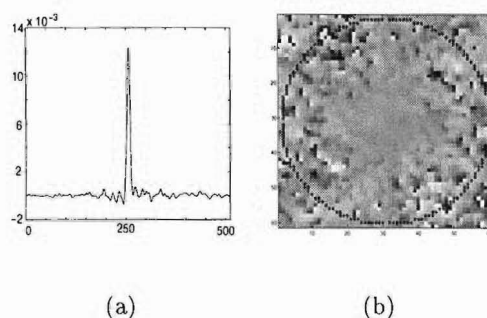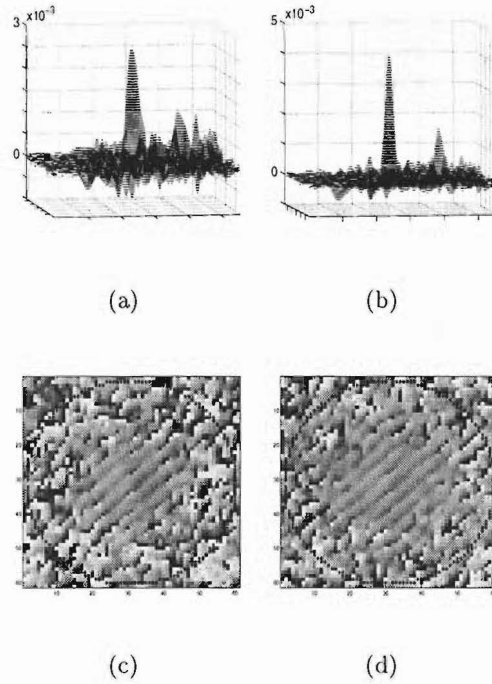
Figure 7.13: *Reconstructions from wavefront sensing data of Castor with a 14 × 14 Shack-Hartmann lenslet array. (a) Reconstruction of the image using traditional centroiding. (b) Reconstruction of the image using model-fitting. In (a) and (b), only the centre portions of size 280 × 280 have been shown for clarity. (c) Phase reconstruction using traditional centroiding. The circle denotes 24% of the diffraction limit. (d) Phase reconstruction using model-fitting.*

7.6 Summary

Object reconstruction from simultaneous wavefront sensing data and speckle images is a two-level inverse problem. The reconstruction of the object relies on the phase estimation, which in turn is dependent on the wavefront sensing measurements.

In this chapter, the second problem of phase reconstruction from wavefront sensing measurements was addressed. Specifically, the accuracy of the existing centroid estimator was examined, and an improved method was proposed. This method retrieves the model of the spot using a blind deconvolution framework, and performs model-fitting to locate the centroid of the spots. The results on both simulated and real data show that the model-fitting procedure produces superior results to those of the traditional centroiding method.

The blind deconvolution procedure, although computationally intensive, needs to be performed only once for a given wavefront sensor configuration. Once the model is retrieved, the model-fitting procedure can be carried out efficiently.

The multiframe blind deconvolution procedure proposed here is not limited to astronomical data processing. It is a general technique for resolution improvement, and can be applied to other images.

Chapter 8

Joint application of WFS data and speckle images in IBD

In Chapter 7, the problem of astronomical object reconstruction was viewed as a two-level inverse problem: firstly the phase is reconstructed from centroid measurements, and secondly the object is reconstructed with the estimated phase. This two-level structure is adopted by the methods of deconvolution from wavefront sensing and adaptive optics where the phase estimation precedes the object reconstruction. However, information regarding the phase is not only present in the wavefront sensing images, but is also intrinsic to the speckle images. Therefore if simultaneous wavefront sensing data and speckle images are available, the phase estimate can be made jointly using both sets of data, improving the accuracy and fidelity of the reconstruction.

The method proposed in this chapter is to pose the problem as one of blind deconvolution, and include the wavefront sensing data as a form of additional constraint. Whilst the estimation of the phase from intensity (speckle images) in a blind deconvolution (BD) framework is highly non-linear, the wavefront sensing data provides a sufficiently strong constraint in the proposed formulation that the initial solution is close to the true solution and as a result, the non-linearity is removed [175].

The concept of using simultaneous speckle images and wavefront sensing data is not new; a similar approach has been used by several authors [24, 27, 51, 126], whose algorithms

187

were primarily developed in response to the need for the post-processing of adaptive optics imagery. These are called myopic deconvolution methods. Strong constraints as well as good quality prior information such as the power spectral density (PSD) of the PSF, and the number of stars in a cluster are used in these algorithms.

The algorithm proposed by Mugnier et al. [126] is very similar to the algorithm proposed here. However, a direct comparison of the performance of the two algorithms is difficult without any quantitative parameters. Specifically, the issue of the number of photons required in the speckle images to guarantee a good solution needs to be addressed. As a comparison, Mugnier used 10 frames with 67500 photons per frame in the processing, whereas 300 frames with an average of 120 photons per frame are used here. While the number of frames used can be varied experimentally, the light level depends on the size of the telescope. This illustrates the importance of the ability to obtain a solution from images with very few photons. In this chapter, a suggestion as to how a quantitative comparison can be performed is discussed.

Another joint application of wavefront sensing data and speckle images was proposed by Roggemann et al. [145, 146]. Here, the entire Shack-Hartmann (SH) image is processed together with the entire speckle image formed by the same aberration. This formulation allows the SH sensor to operate above its usual range of $D/r_0 = 1$ per lenslet, and large aberrations can be estimated. This approach to the problem is beyond the scope of this thesis; here the operation of the SH sensor is restricted to its usual range of $D/r_0 = 1$ per lenslet.

This chapter is structured as follows. Section 8.1 presents the theory of reconstructing a phase from wavefront sensing measurements. The proposed method is discussed in Section 8.2. Section 8.3 examines the implementation issues when dealing with real data, and results on experimental data are included in Section 8.4. A suggestion for performance quantification is presented in Section 8.5, and conclusions are drawn in Section 8.6.

8.1   The processing of wavefront sensing data

One common approach of measuring the atmospheric phase distortion is to use a wavefront sensor (WFS). There are three main types of WFSs: the Shack-Hartmann (SH) WFS and

the shearing interferometer which are sensitive to the slope of the wavefront phase [147], and the curvature wavefront sensor [142]. Since the real data used in this thesis was captured by the SH-WFS, the other types of WFSs are not discussed.

The SH sensor comprises a square array of lenslets which segments the aperture into smaller sections called subapertures. A typical frame captured by a SH-WFS is shown in Fig. 8.1(a). Basically, the displacement of a spot from the point corresponding to the centre of a lenslet is used to produce the average or mean slope of the wavefront over the respective lenslet. These local slopes are then combined to reconstruct the phase across the entire aperture. Fig. 8.1(b) illustrates the relationship between the slope of the incoming wavefront over a subaperture and the displacement of the spot in the detector plane. The wavefront slope $\mathbf{s} = \Delta\phi/d$ is represented by:

$$\mathbf{s} = k\frac{|\mathbf{x}_s|}{f_l} \tag{8.1}$$

where $k = 2\pi/\lambda$ is the wavenumber, and f_l is the focal length of the subaperture lens. The spot location, $\mathbf{x}_s = (x_s, y_s)$, is found by computing the centroid (or centre of gravity) of the spot on the detector plane as in Eq. (7.1).

The application of this concept is illustrated in Fig. 8.2. A binary object in Fig. 8.2(a) is imaged through atmospheric turbulence described by a simulated phase screen with $D/r_0 = 8$ (Fig. 8.2(b)). The image measured by the wavefront sensor is shown in Fig. 8.2(c). By measuring the centroid of the spot in each lenslet, the mean slope over each subaperture is found, which are then combined to form the overall phase estimate (Fig. 8.2(d)). The object reconstruction obtained using this phase estimate is shown in Fig. 8.2(e). Here, it is assumed that a simultaneous speckle image is captured, so that a Wiener filter can be used to provide the reconstruction:

$$\hat{f}(x,y) = \mathfrak{F}^{-1}\left\{\frac{D(u,v)H^*(u,v)}{|H(u,v)|^2 + \text{const}}\right\}. \tag{8.2}$$

Here $D(u,v)$ and $H(u,v)$ are the spectra of the data and the PSF respectively. A diffraction-limited image of the object is shown in Fig. 8.2(f) for comparison. In the case where multiple frames of wavefront sensing data and simultaneous speckle images are available, this procedure can be extended to a vector Wiener filter, where

$$\hat{f}(x,y) = \mathfrak{F}^{-1}\left\{\frac{\langle D(u,v)H^*(u,v)\rangle}{\langle|H(u,v)|^2\rangle + \text{const}}\right\} \tag{8.3}$$

and $\langle\cdot\rangle$ denotes an ensemble quantity. This is essentially the concept of deconvolution from wavefront sensing (DWFS) which is discussed in detail in Chapter 9.
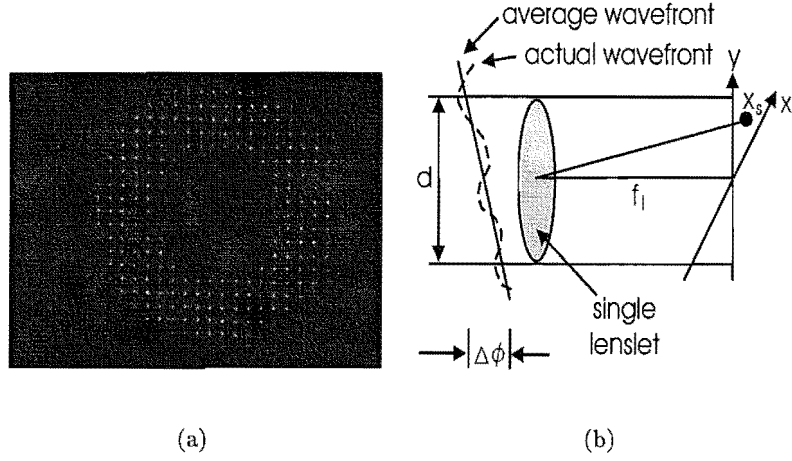
Figure 8.1: (a) A typical frame of wavefront sensing data. (b) The re-
lationship between the slope of the wavefront and the displacement of
the spot on the detector plane.

A more rigorous, mathematical approach to finding the wavefront slope over a subaperture
can be obtained by computing its average wavefront phase gradient [186]:

$$\mathbf{s}(j) = \int W_{sj}(u,v)\nabla\phi(u,v)dudv \tag{8.4}$$

where $\mathbf{s}(j)$ is the wavefront slope of the j^{th} subaperture, $\nabla\phi(u,v)$ is the spatial gradient of
$\phi(u,v)$, and $W_{sj}(u,v)$ is the weighting function of the j^{th} subaperture.

When noise is taken into account in the SH-WFS data, the slope measurement is:

$$
\begin{aligned}
\mathbf{s}_m(j) &= \mathbf{s}(j) + \mathbf{s}_n(j) \\
&= \int W_{sj}(u,v)\nabla\phi(u,v)dudv + \mathbf{s}_n(j) \\
&= -\int \nabla W_{sj}(u,v)\phi(u,v)dudv + \mathbf{s}_n(j)
\end{aligned}
\tag{8.5}
$$

where the last equation is obtained by integrating the first term by parts [178]. $\mathbf{s}_n(j)$ is
the random variable describing the combined effects of photon and read noise in the j^{th}
subaperture, and $\mathbf{s}_m(j)$ is the measurement of the corresponding subaperture. The phase
and the sensor noise are assumed to be independent.

(a)   (b)

(c)   (d)

(e)   (f)
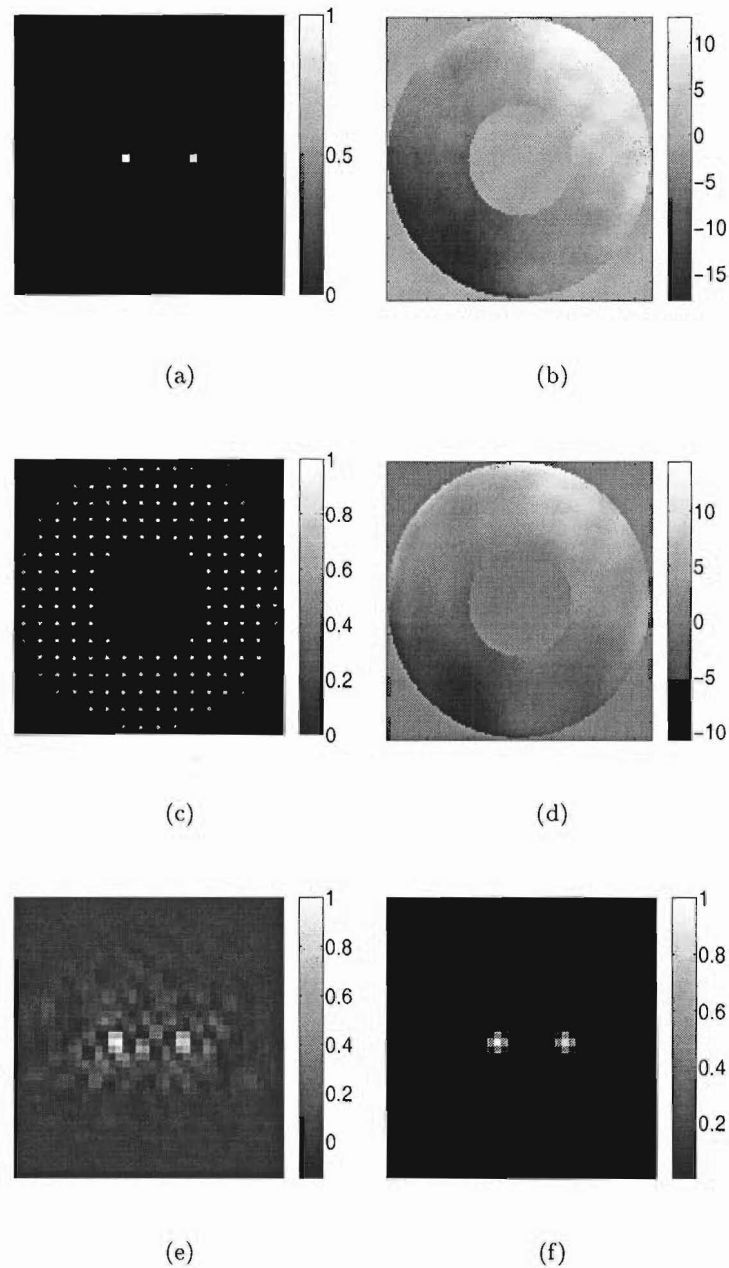
Figure 8.2: *A demonstration of the phase reconstruction from Shack-Hartmann wavefront sensing data. (a) The simulated binary star. (b) A phase screen with $D/r_0 = 8$. Note that an aperture with a central obscuration is imposed. (c) The wavefront sensing image. (d) Phase reconstruction. (e) Object reconstruction. (f) The diffraction-limited image of the object for comparison.*

8.1.1   Phase reconstruction

The reconstruction of a wavefront phase is most conveniently expressed as a weighted sum of a set of basis functions, an appropriate choice of which is the Zernike basis functions. A phase estimate $\hat{\phi}(u, v)$ can therefore be written as:

$$\hat{\phi}(u, v) = \sum_{i=1}^{N} c_i Z_i(u, v) \tag{8.6}$$

when the first N modes of the Zernike bases excluding piston are used. The weights c_i are obtained by

$$c_i = \sum_{j=1}^{2N} M_{ij} s_j, \tag{8.7}$$

or, in vector notation $\mathbf{c} = \mathbf{Ms}$. $\mathbf{c} = (c_1, c_2, \dots, c_N)$ is the vector of the weights, and $\mathbf{s} = (s_1, s_2, \dots, s_{2N})^T$ is the vector of the slope measurements. It should be pointed out that a wavefront sensor with N subapertures provides $2N$ measurements from the row and column displacements. The ordering of the slope measurements adopted in this thesis is:

$$
\begin{aligned}
s_1 &= s_m^x(1), \\
s_2 &= s_m^y(1), \\
s_3 &= s_m^x(2), \\
s_4 &= s_m^y(2), \\
&\quad\cdot \\
&\quad\cdot \\
s_{2N-1} &= s_m^x(N), \\
s_{2N} &= s_m^y(N).
\end{aligned} \tag{8.8}
$$

In other words, the row and column slope measurements are interlaced.

The matrix $\mathbf{M}$ is the reconstruction matrix which maps the slope measurements to the weights of the bases. A common choice of $\mathbf{M}$ is derived from the least squares solution, where the difference between a set of noise-free wavefront sensor slopes and the corresponding set of slopes made up from Zernike basis functions is minimised. This error is expressed as:

$$\Delta = \mathbf{s} - \Theta \mathbf{c} \tag{8.9}$$

where the elements of Θ, Θ_{ji}, are:

$$\Theta_{ji} = \frac{\partial s_j}{\partial c_i}. \tag{8.10}$$

Θ is called the interaction matrix [108] since it defines the relationship between the centroid measurements and the Zernike weights. In a least squares formulation, the squared error to be minimised is $\Delta^T \Delta$. By differentiating this with respect to c_i and equating it to zero, the optimal set of weights $\mathbf{c}^{\mathrm{opt}}$ is:

$$\mathbf{c}^{\mathrm{opt}} = (\Theta^T \Theta)^{-1} \Theta^T \mathbf{s} \tag{8.11}$$

and the direct result is that $\mathbf{M} = (\Theta^T \Theta)^{-1} \Theta^T$.

The least squares solution provides a good estimation of the weights for the Zernike basis functions, and this method is commonly used. However, the effects arising from non-zero correlations between WFS measurements are not accounted for [144]. In this case, prior information can be introduced into the solution in the form of a covariance matrix.

The covariance matrix of the Zernike polynomial coefficients, $\mathbf{K}_{zz}$, is given by:

$$\mathbf{K}_{zz} = \langle \mathbf{a}\mathbf{a}^{\mathbf{T}} \rangle \tag{8.12}$$

where $\mathbf{a}$ is a vector of coefficients for the Zernike modes. For the i^{th} row and the j^{th} column of the covariance matrix, the entry is [147]:

$$\begin{aligned}
a_i a_j &= 0.0072\Big(\frac{D}{r_0}\Big)^{5/3}(-1)^{(n_i+n_j-2m_i)/2}[(n_i+1)(n_j+1)]^{1/2}\pi^{8/3}\delta_{m_i m_j} \\
&\times \frac{\Gamma(14/3)\Gamma[(n_i+n_j-5/3)/2]}{\Gamma[(n_i-n_j+17/3)/2]\Gamma[(n_j-n_i+17/3)/2]\Gamma[(n_i+n_j+23/3)/2]}
\end{aligned} \tag{8.13}$$

for $i - j$ even. For $i - j$ odd:

$$a_i a_j = 0. \tag{8.14}$$

Following the same notation in Chapter 4, m_i and n_i in the above equations are the azimuthal and radial orders associated with the i^{th} Zernike polynomial respectively, and m_j and n_j are those of the j^{th} Zernike polynomial. The Gamma function $\Gamma[\cdot]$ is defined in Section 2.4.3. When the covariance of the Zernike polynomials is incorporated into the framework, the least squares solution is modified to a minimum variance solution:

$$\mathbf{M} = (\mathbf{K}_{zz}^{-1} + \Theta^T \mathbf{K}_{nn}^{-1}\Theta)^{-1}\Theta^T \mathbf{K}_{nn}^{-1} \tag{8.15}$$

where $\mathbf{K}_{nn}$ is an identity matrix due to the Gaussian nature of the noise. Both of these formulations are used in this thesis. The least squares solution is used for calibrating any scaling factors, while the prior information is incorporated in the reconstructions.

8.1.2   Inadequacies in using WFS data alone

Using the wavefront sensing data and speckle images in a Wiener filter to reconstruct the object is easy to implement. However, one problem is that since the PSF estimate is made solely from data measured in the wavefront sensor channel, any information regarding the PSF contained in the speckle image channel is not used. Another deficiency is that it is not possible to estimate the PSF exactly, which causes a bias problem [149]. This error in estimating the PSF is reflected in the object reconstruction.

One solution to this problem is to use both the wavefront sensing data and the speckle images to estimate the PSF. This can be achieved in a blind deconvolution framework, using the wavefront sensing data as a constraint on the solution, ensuring that the wavefront estimate is made jointly from the data that is present in the image and the wavefront sensing measurements. This also removes the common problem of the trivial solution in blind deconvolution, where the object is estimated to be a point source, and the PSF the data.

8.2   An improved method for object and PSF reconstruction

When solving an inverse problem, a common approach is to firstly obtain a good model of the system. The object estimate is then put through this degradataion model, and the estimated data is compared to the actual data observed. The discrepancy is used to update the estimate of the object. Prior information can also be included in the framework in order to regularise the solution. This is essentially the maximum *a posteriori* (MAP) framework.

In the proposed blind deconvolution problem, the two sets of data are the wavefront sensing images and the speckle images. In the previous section, the forward model of the wavefront sensing images was presented, where the phase $\phi(u,v)$ is expanded onto a set of Zernike bases, whose weights are linearly related to the centroid measurements of the wavefront sensing image. The forward model for the speckle images is denoted by the convolution operation as in Eq. (6.1); and the relationship between the PSF, the phase and the generalised pupil function is detailed in Section 6.4.

The solution to the associated inverse problem is illustrated diagrammatically in Fig. 8.3.
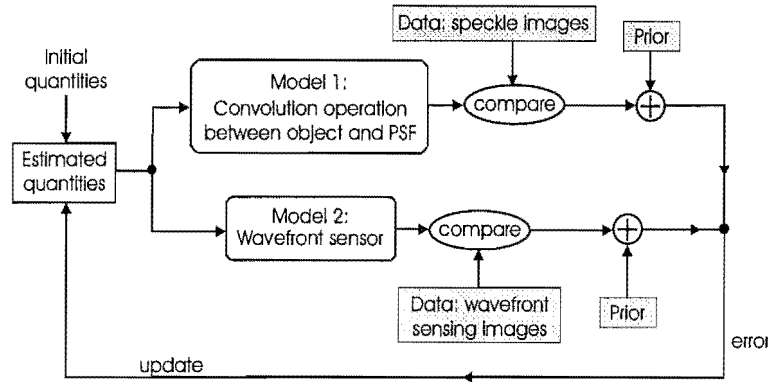
Figure 8.3: *A block diagram illustrating the inverse problem solution when both speckle images and wavefront sensing data are used in the reconstruction.*

Mathematically, the solution is posed as a problem in a MAP framework:

$$\max_{\hat{f}, \hat{h}_k} P(\hat{f}(x,y), \hat{h}_k(x,y)|d_k(x,y)). \qquad (8.16)$$

It should be pointed out that because the phase $\phi(u,v)$ has a direct relationship with the PSF, and the phase can be expanded with a set of Zernike weights, it is equivalent to write the MAP solution as:

$$\max_{\hat{f}, \hat{c}_k} P(\hat{f}(x,y), \hat{c}_k|d_k(x,y)). \qquad (8.17)$$

The application of Bayes' rule and other mathematical simplifications lead to the log-likelihood equation:

$$\max_{\hat{f}, \hat{c}_k} \log[P(\hat{f}(x,y), \hat{c}_k|d_k(x,y))] = \max_{\hat{f}, \hat{c}_k} \log[P(d_k(x,y)|\hat{f}(x,y), \hat{c}_k)]$$
$$+ \log[P(\hat{f}(x,y))]$$
$$+ \log[P(\hat{c}_k)]. \qquad (8.18)$$

The first term is the consistency between the estimates and the observed data, while the other terms are for the inclusion of prior information or regularisation.

The consistency term consists of two components:

$$E_{\mathrm{con}} = E_{\mathrm{speckle}} + \gamma_c E_{\mathrm{centroid}}, \qquad (8.19)$$

where

$$E_{\mathrm{speckle}} = \sum_k \sum_{x,y} [\hat{f}(x,y) \odot \hat{h}_k(x,y) - d_k(x,y)]^2 \qquad (8.20)$$

assuming white Gaussian noise, and

$$E_{\text{centroid}} = \sum_k [\hat{\mathbf{s}}_k - \mathbf{s}_k]^2. \tag{8.21}$$

The scalar γ_c determines the weighting between each of the consistency terms. E_{speckle} constrains the estimates of f and h to be consistent with the data, while E_{centroid} ensures that the centroid estimates are close to $\mathbf{s}_k$.

It should be pointed out that while in E_{speckle}, the quantity used for comparison is the measured data $d_k(x, y)$, it is not so in E_{centroid}. Instead, a good estimate of the centroids, $\mathbf{s}_k$, which can be obtained by the model-fitting procedure discussed in Chapter 7, is used for comparison. This is done to compromise between the complexity of the model and the speed of obtaining a solution. Optimally, one can update the phase estimate and project it through the wavefront sensor model to obtain estimates of the wavefront sensing images, and compare them to the measured images. However, this adds extra computational requirements to the algorithm, which is already computationally intensive (see Section 8.3.3). As a result, this compromise is reached between speed and accuracy.

The necessary gradient calculation for updating the object $\hat{f}(x, y)$ is:

$$\frac{\partial E_{\text{con}}}{\partial \hat{f}(x, y)} = \sum_k 2[\hat{f}(x, y) \odot \hat{h}_k(x, y) - d_k(x, y)] \otimes \hat{h}_k(x, y) \tag{8.22}$$

while the calculations for updating the weights $\hat{\mathbf{c}}_k$ are

$$\frac{\partial E_{\text{con}}}{\partial \hat{\mathbf{c}}_k} = \frac{\partial E_{\text{speckle}}}{\partial \hat{\mathbf{c}}_k} + \gamma_c \frac{\partial E_{\text{centroid}}}{\partial \hat{\mathbf{c}}_k}. \tag{8.23}$$

The first term $\partial E_{\text{speckle}}/\partial \hat{\mathbf{c}}_k$ is:

$$
\begin{aligned}
\frac{\partial E_{\text{speckle}}}{\partial \hat{\mathbf{c}}_k} &= \frac{\partial E_{\text{speckle}}}{\partial \hat{\phi}_k(u, v)} \times \frac{\partial \hat{\phi}_k(u, v)}{\partial \hat{\mathbf{c}}_k} \\
&= \text{Im}\left\{ \mathfrak{F}\left\{ \frac{\partial E_{\text{speckle}}}{\partial \hat{a}_k(x, y)} \right\} \times \hat{A}_k^*(u, v) \right\} \times \frac{\partial \hat{\phi}_k(u, v)}{\partial \hat{\mathbf{c}}_k} \tag{8.24}
\end{aligned}
$$

where [168]:

$$\frac{\partial E_{\text{speckle}}}{\partial \hat{a}_k(x, y)} = 2[\hat{f}(x, y) \odot \hat{h}_k(x, y) - d_k(x, y)] \otimes \hat{f}(x, y) \times 2\hat{a}_k(x, y). \tag{8.25}$$

The term $\partial \hat{\phi}_k(u, v)/\partial \hat{\mathbf{c}}_k$ is obtained by expanding the gradient $\partial E_{\text{speckle}}/\partial \hat{\phi}_k(u, v)$, which is a two-dimensional array the same size as $\hat{\phi}_k(u, v)$, onto the set of Zernike bases. The second term $\partial E_{\text{centroid}}/\partial \hat{\mathbf{c}}_k$ is:

$$\frac{\partial E_{\text{centroid}}}{\partial \hat{\mathbf{c}}_k} = 2(\hat{\mathbf{s}}_k - \mathbf{s}_k)^T \Theta. \tag{8.26}$$

Both the estimates of the object and the weights are updated using the method of steepest descent, the formulation of which is included in Appendix A.

8.2.1 Prior information for the object and the phase

The prior information enforced on the object estimate includes positivity, implemented using the method of Johnston and Lane [94], and Tikhonov-Miller (TM) regularisation. The corresponding error term is written as:

$$E_f = \gamma_f \sum_{x,y} \hat{f}(x,y)^2 \tag{8.27}$$

and its gradient with respect to $\hat{f}(x,y)$ is:

$$\frac{\partial E_f}{\partial \hat{f}(x,y)} = 2\gamma_f \hat{f}(x,y). \tag{8.28}$$

The prior for the Zernike weights is the covariance matrix of the Zernike coefficients. The calculation of the entries in the covariance matrix has been discussed in Section 8.1.1. The error term corresponding to this prior is:

$$E_c = \frac{1}{2}\gamma_z \sum_k \hat{c}_k^T K_{zz}^{-1} \hat{c}_k \tag{8.29}$$

where K_{zz} denotes the covariance matrix. Its gradient with respect to $\hat{c}_k$ is:

$$\frac{\partial E_c}{\partial \hat{c}_k} = \gamma_z \hat{c}_k^T K_{zz}^{-1}. \tag{8.30}$$

8.3 Implementation issues for experimental data

The use of experimental data has several implementation issues which are discussed in the following subsections. These include the centroiding operation, the calibration between the wavefront sensing and the speckle image cameras, and the speed and memory requirements.

8.3.1 Centroiding operation

The aim of the centroiding operation is to calculate the movement of a spot from the position corresponding to the centre of the lenslet. The accuracy of the traditional centroid

(a)   (b)
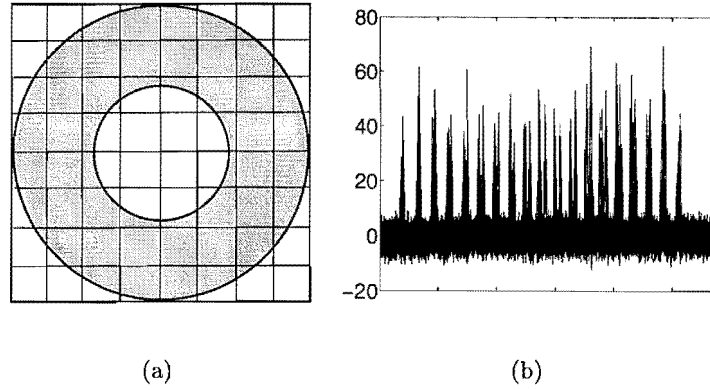
Figure 8.4: (a) An illustration of the aperture of a telescope with central obscuration. (b) The projection of a typical frame of wavefront sensing data captured at the Observatoire de Lyon. The corresponding two-dimensional image was shown in Fig. 8.1(a).

estimator has been discussed in Chapter 7. Here two implementation issues are addressed. Firstly, because the aperture of a telescope is usually circular with a central obscuration, it is inevitable that some of the subapertures are not fully illuminated. In a simulation, it is usual to include the subapertures that have more than 50% illumination in the reconstruction. Fig. 8.4(a) shows an example of an 8 × 8 lenslet array, where 40 out of the 64 subapertures can be used. However, when dealing with real data, a different criterion must be used for including or excluding spots. It is common to use a threshold, whose value can be determined by examining the data. Fig. 8.4(b) shows a projection of a real wavefront sensing data frame (Fig. 8.1(a)). The threshold in this case is set to 20, a value found by comparison with the noise floor.

Another problem encountered is that when centroiding each spot, it is necessary to find reference coordinates for the spots. In order to circumvent this problem in real data, a long exposure wavefront sensing image can be taken. The positions of the spots in this image approximate the centres of the lenslets. This is a valid operation because the movement of the spot produced by atmospheric turbulence follows a zero-mean Gaussian distribution when observed over a long period of time [43, 62]. Although fixed telescope aberrations do not affect the location of the centroids, they cause a systematic error which should be calibrated initially so that it can be subtracted from the estimated phase. This is discussed in the next subsection.

8.3.2   Calibration between the WFS and CP40 cameras

As discussed in Section 6.4, there exists a Fourier transform relationship between the PSF and the OTF. However, the relationship is not direct in the case of the data captured at the Observatoire de Lyon. This is because the wavefront sensing data and the speckle images are captured on cameras with the same number of pixels, yet one represents samples in the frequency domain, and the other in the spatial domain. This means that the sampling rates are not necessarily compatible. Moreover, the two cameras were not aligned with respect to each other during the observing run; as a result the wavefront sensing data frames are rotated with respect to the speckle images. It is therefore essential to calibrate and resample one set of data for the Fourier transform relationship to hold.

There are two alternatives for compensating for this problem. In the first method, most of the correction is performed in the spatial domain. Let the complex aperture be denoted by:

$$A(u,v) = S(u,v) \exp\{ic_p\phi(u,v)\} \tag{8.31}$$

where $S(u,v)$ is the aperture shown in Fig. 8.5(a). This is obtained by averaging an ensemble of images of the aperture on the telescope. Note the presence of the supports for the secondary mirror (or 'spider') in the figure. $\phi(u,v)$ represents a phase reconstructed from a frame of wavefront sensing data. However, in order to illustrate the process, a simple phase screen with a tip, tilt and defocus (Zernike modes 2, 3 and 4) is used for $\phi(u,v)$ in this section (Fig. 8.5(b)). c_p is the scaling for the phase screen and must be calibrated experimentally. For the data captured here, c_p is approximately 34.

Taking a direct Fourier transform of the autocorrelation of $A(u,v)$ gives an incorrect PSF as shown in Fig. 8.6(a). Only the centre 160×160 pixels are shown for clarity. The factors that must be compensated for include the rotation, the non-square pixels (the pixels are slightly wider in the vertical direction), and the sampling. By interpolating this PSF on a different grid, a corrected PSF is obtained as shown in Fig. 8.6(b). Note that this is rotated by approximately 90°, and appears slightly larger because of the difference in sampling rate. Similar to the scaling of the phase c_p, these parameters have been obtained from an optimisation procedure using the captured data.

However, this method is computationally intensive. This is especially significant in a recursive blind deconvolution framework, where each frame of PSF must be generated in

(a)   (b)
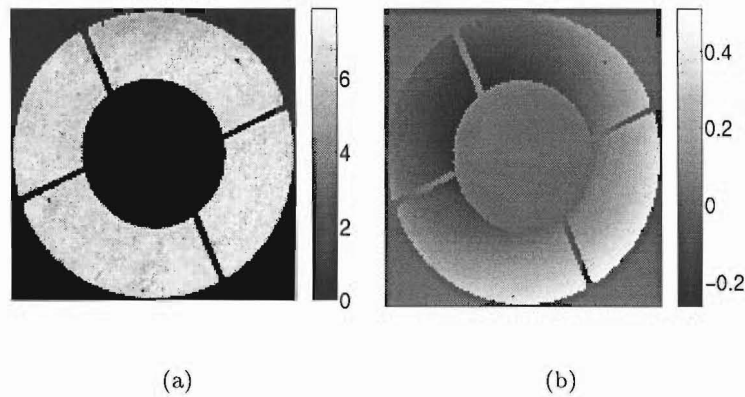
Figure 8.5: *(a) The telsecope aperture obtained by averaging an ensemble of images (131×131 pixels). (b) The phase screen used in this section for illustration purposes. It is made up of tip, tilt and defocus components.*

(a)   (b)
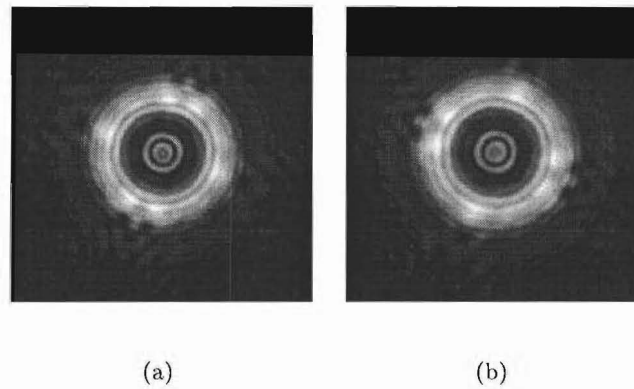
Figure 8.6: *Calibrating between the WFS and the image channels in the image plane. (a) The initial, uncompensated PSF. (b) The compensated PSF.*

each iteration. In order to improve the speed performance of the algorithm, an alternative method is devised.

The second method performs the correction in the frequency domain so that a direct Fourier transform can be taken to form the PSF. This is achieved by expanding the phase $\phi(u, v)$ over a compensated set of Zernike polynomials. This requires the generation of a new set of Zernike bases with the sampling, pixel size and rotation taken into account. The aperture must also be resampled to match the Zernike bases (Fig. 8.7). Although the new set of
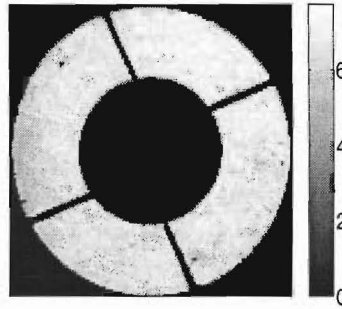
Figure 8.7: *The resampled and rotated aperture (121 × 123 pixels).*

Zernike bases has a non-square number of pixels, it is still defined on a unit circle, except that each basis is sampled slightly more finely in one direction. The resultant PSF is the same as that in Fig. 8.6(b) within rounding errors.

It is important to note that in this blind deconvolution framework, two sets of Zernike polynomials are used. The first set is the resampled set used for the representation of the phase in the aperture. The second set is the reference (with a square number of pixels) used for generating the interaction matrix Θ, which is then used to find the weights of the Zernike polynomials from the centroid measurements. A question one might ask is that, can the weights, reconstructed using the second set of Zernike polynomials, be used to generate a phase screen using the first set of polynomials? The answer is that this is not a problem or a contradiction, because the weights are invariant regardless of the basis functions used.

The final calibration that must be performed is the incorporation of the fixed telescope aberration into the reconstruction. This aberration is characterised by modes 4 to 15 of the Zernike bases, the coefficients of which are listed in Table 8.1. Once again, these parameters are optimised in an iterative procedure on the real data.

8.3.3 Speed and memory requirements

For any recursive algorithm, it is important to consider the speed and memory requirements for computer implementation. In this blind deconvolution problem, there are 300 data frames and 300 PSFs (each of size 512 × 512) to be processed. As a result, the speed and memory requirements must be examined.

Zernike mode	coefficient
Z_4 (defocus)	-0.0040
Z_5 (astigmatism)	0.0005
Z_6 (astigmatism)	-0.0060
Z_7 (coma)	0.0039
Z_8 (coma)	-0.0058
Z_9	-0.0039
Z_{10}	-0.0032
Z_{11} (spherical aberration)	-0.0016
Z_{12}	-0.0033
Z_{13}	0.0012
Z_{14}	-0.0036
Z_{15}	-0.0001

Table 8.1: *The characterisation of the fixed telescope aberrations using Zernike polynomials for the telescope at the Observatoire de Lyon.*

To reduce the memory requirement of the data frames, it is observed that each frame has only approximately 120 photons, embedded in a 512×512 array. It is therefore possible to store them as sparse arrays in Matlab [117] which reduces the memory storage significantly.

Since it is possible to initialise the phase estimates with reconstructions from the wavefront sensing data, their good quality means that there is little change to them during the course of the iteration. Therefore instead of calculating each quantity as required, as performed in a blind deconvolution framework, a simplification to a deconvolution problem is possible. By generating and saving the PSFs in memory, and accessing them slice by slice when needed reduces the memory requirement without much loss in the quality of the reconstruction. The speed is also improved by 26%. However, because the PSF data still amounts to 600MB, its access from hard disk is still relatively time-consuming.

In the next section, experimental results using both the blind and non-blind deconvolution approaches are presented.

8.4 Results using experimental data

The experimental data was collected with the SPID instrument [3] attached to a 1-metre telescope at the Observatoire de Lyon. The object of observation (Castor) is a binary star. The wavefront sensing data was captured using a 14×14 Shack-Hartmann lenslet array with an exposure time of 3.2ms, while the speckle images used in this processing were captured at a wavelength in the band $620 - 627$nm.

Out of the 2500 frames captured, every 5^{th} frame up to a total of 300 frames are used in the reconstruction. Consecutive frames are not used in order to reduce the temporal correlation between frames. The average number of photons per speckle image is 120, an example of which is shown in Fig. 4.18(a). This low light level renders it impossible to distinguish any features of the astronomical object. For the wavefront sensing data, the centroid estimates are calculated using the model-fitting procedure proposed in Chapter 7, thus providing a good quality estimate of the Zernike weights.

Two initial estimates for the object are tested: a zero start, and a uniform start. The weightings γ_c and γ_f are found experimentally and take on the values of $\gamma_c = 1 \times 10^{-5}$ and $\gamma_f = 1 \times 10^{-4}$ respectively. Although these values seem *ad hoc*, a good initial guess can be achieved by scaling them such that each of the error terms has approximately the same energy. The application of a prior for the weights in the form of the Kolmogorov covariance matrix is found to be unnecessary in this case because the starting estimates of the weights are good and consequently they are already sufficiently regularised.

The brightness ratio and the separation of the two stars of the binary are used to determine the quality of the reconstructions. To calculate these parameters, a support for each star is firstly identified in the reconstruction. The centre point of the support for each star is then used for calculating both the brightness ratio and the separation. Although this approach may introduce an error, it is the best approach given the noisy nature of the reconstructions. Figs. 8.8(a) and 8.8(b) show the reconstruction of Castor using a zero start and a uniform start in a blind deconvolution framework. The brightness ratio of the two stars are 2.4 and 2.3 in the respective reconstruction, which is in close agreement with the widely accepted value of 2.1. The separations between the two stars are both approximately 3.2 arcsec. Although this value is slightly lower than the reference value of 3.6 arcsec, they are still comparable. Moreover, as discussed in Section 7.4.2, this separation is dependent

on computed values from the optics which have margins of error. A more accurate method
is to calibrate the arcsecond-per-pixel value from a known binary. Another possible source
of error is that the reconstructions are still a convolution of the underlying object with
the telescope PSF [168]. This can be overcome by applying another blind deconvolution
procedure on the reconstructions in Fig. 8.8.

To improve the speed of the algorithm, a non-blind alternative is also investigated, with
the results shown in Fig. 8.8(c) and 8.8(d) for a zero start and uniform start respectively.
The brightness ratios and separations are both 2.1 and 3.1 arcsec respectively. These are
different from but compatible with the blind deconvolution simulations.

The brightness ratios obtained using blind deconvolution are higher than those retrieved
from the deconvolution method because complete convergence had not been achieved (due
to the heavy computational requirement). The error metric was however decreasing slowly
at this point. As the algorithms proceeded, the brightness ratios were approaching the ref-
erence value of 2.1, indicating that this value would eventually be reached if the algorithms
were allowed to proceed further. This observation is supported by the findings in Ref. [175]
which state that if the initial phase estimate is close to the true phase, which is the situation
at hand, then uniqueness of the solution is guaranteed.

In all the reconstructions, although a binary star is clearly visible, each star is not smooth.
This can be due to the insufficient regularisation of the object. One solution is to increase
the value of the weighting γ_f. This will give a smoother and slightly broader reconstruction.
Another approach is to use a different form of regularisation instead of the Tikhonov-Miller
regularisation used here. One possibility is the Laplacian kernel.

The blind deconvolution and the deconvolution methods achieved different but compatible
results. This shows that because the Zernike weights were reconstructed accurately initially,
there is only a small variation to their values throughout the course of the iteration. As a
result, a non-blind formulation can be used for the object reconstruction when the speed of
convergence is an issue. It should also be pointed out that the algorithm is robust to the
starting estimate of $\hat{f}(x,y)$.

Although no quantitative comparison can be made between the results achieved in this
section and those shown in the paper by Mugnier *et al.*, it must be pointed out that here, the

(a)   (b)

(c)   (d)
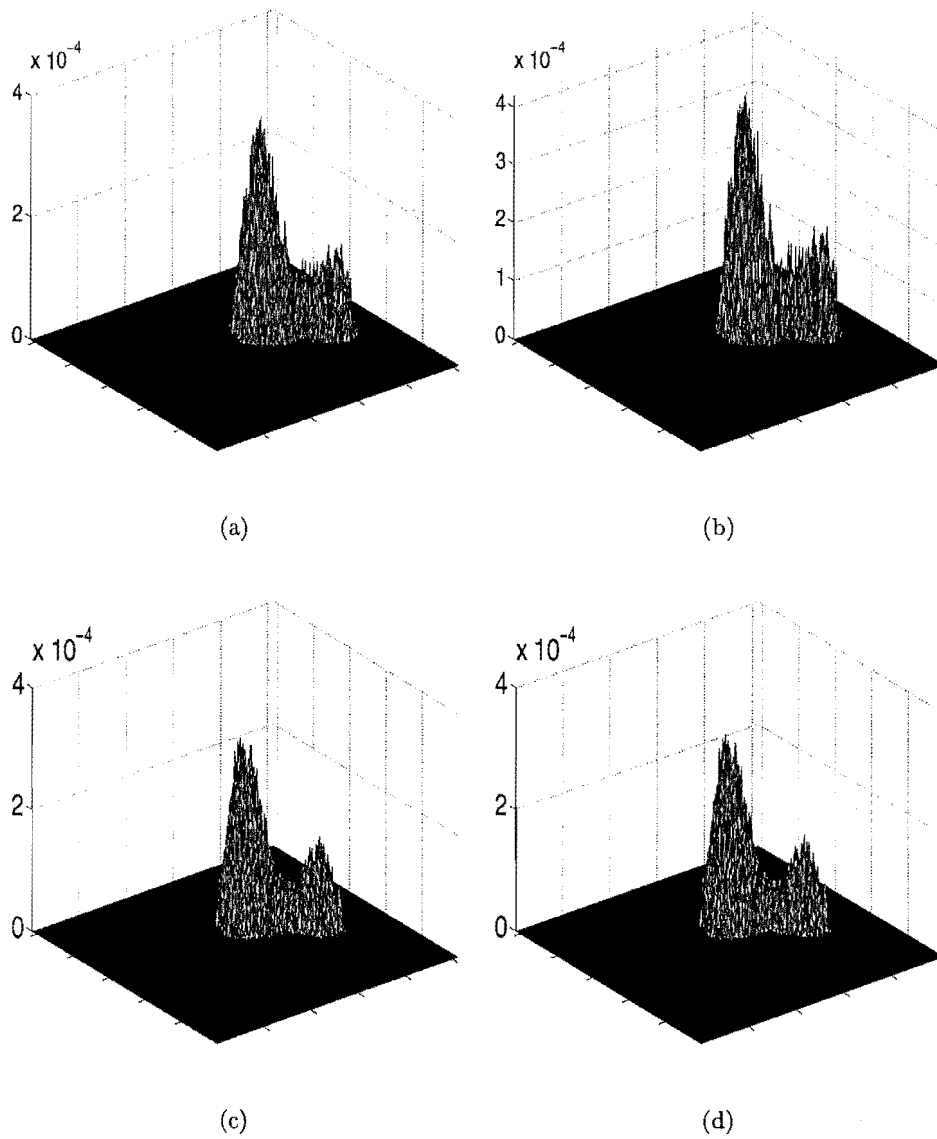
Figure 8.8: *The reconstruction of Castor with (a) a zero-start and (b)
a uniform start in a blind deconvolution framework. (c) and (d) are
the corresponding reconstructions in a non-blind framework. The non-
blind deconvolution is performed to speed up the iterative algorithm.
The frames are of size 512 × 512.*

300 frames of speckle images have an average of 120 photons per frame, while in Mugnier's paper, 10 speckle images with 67500 photons per frame were used. While the number of frames is easily controlled experimentally, the light level is dependent on the size of the telescope. It is therefore important to be able to reconstruct an object under low light levels.

8.4.1   Comparison with DWFS

The results from the proposed algorithm are also compared to the conventional deconvolution from wavefront sensing (DWFS) method. Since no reference star measurements were taken, the estimates of the PSFs were computed using wavefront sensing measurements of the binary star Castor. Fig. 8.9 shows the object reconstruction using the same 300 data frames. The binary star is not clearly visible, and the reconstruction is very noisy with negative values present in the background. The visual quality of results using the proposed method is clearly superior to that of DWFS.

Figure 8.9: *A comparison with the result of DWFS.*

8.5 Quantification of performance

The feasibility of using wavefront sensing data in a blind deconvolution framework has been demonstrated. The next step is to quantify the performance of the algorithm. Specifically, the amount of information gained in the phase reconstruction by the joint processing of speckle images and WFS data is of considerable interest. The analysis has not been performed here, but suggestions on how this can be achieved are proposed.

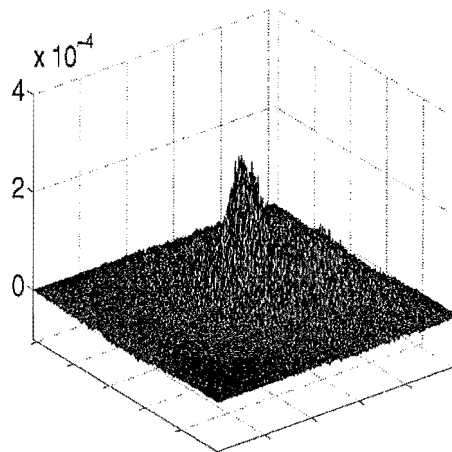Since blind deconvolution is recursive and involves two unknowns, it is difficult to characterise its performance analytically; instead simulations are performed. The parameters to consider include the photon levels, the number of frames, the SNR of the wavefront sensing data and the degree of undersampling in the WFS data. It is important to obtain a lower bound on the photon level required in the speckle images, and a range of values down to 1 photon per frame should be investigated. The undersampling factor is of particular interest regarding the SPID camera because the sampling rate can be as low as 60% of the Nyquist rate.

There are three metrics that can be used to assess the quality of the phase reconstruction. Firstly, the aperture averaged RMS residual phase error [147] for each phase reconstruction can be compared at different photon levels and different values of D/r_0. Secondly, the aperture averaged RMS residual phase error for individual Zernike modes can be examined. It is known that the Shack-Hartmann WFS is insensitive to certain Zernike modes [111]. This metric examines whether the estimation of the phase from simultaneous speckles and WFS frames can ameliorate this problem. Finally, the metric proposed in Eq. (8.19) can also be used. These metrics have not been applied to the investigation in this section because ensemble trials are needed to produce a meaningful result, and in its current form, the algorithm is too computationally intensive for this purpose.

The quality of the object reconstruction can be assessed by standard image metrics such as the mean-squared error if the ground truth is available as in the case of simulated data. For experimental data such as a binary star, the mean and variance of the separation and intensity ratio of an ensemble can be used.

8.6   Summary

In this chapter, a blind deconvolution framework which incorporates simultaneous wavefront sensing data as a constraint was investigated. This ensures that the estimate of the PSFs is made jointly from both sets of data, which is superior to the PSF estimation from WFS data alone in DWFS, and from the speckle images alone in blind deconvolution. The good initial estimate of the PSF is also an advantage in this algorithm.

This algorithm has been demonstrated on real data. Both a blind deconvolution and a deconvolution scheme were tested. The need for the latter arises due to the heavy computational requirement of the algorithm. These schemes produced different but comparable results, indicating that the deconvolution method can be used if the speed of convergence is an issue.

Several improvements can be made to the proposed algorithm. As discussed in Section 8.2, the measured wavefront sensing data is not used in the algorithm unaltered. Instead, the centroid values are calculated using the model-fitting procedure proposed in Chapter 7, and used as the reference values in the algorithm. This is performed to reduce the computational requirements. However, it is worthwhile to investigate the performance of the optimal solution which uses an additional consistency term:

$$E_{\text{con-WFS}} = \sum_k \sum_{x,y} [\hat{d}_{k,WFS}(x,y) - d_{k,WFS}(x,y)]^2 \qquad (8.32)$$

where $\hat{d}_{k,WFS}(x,y)$ is a function of both the projection of the phase estimate through the WFS model $P_{WFS}\{\hat{\phi}_k(u,v)\}$ and the object estimate $\hat{f}(x,y)$. This formulation is consistent with the way the speckle images are handled.

The existing computational requirement of the algorithm is large even though optimisation of the Matlab code has been performed. Using a different language such as C could improve the speed of the algorithm.

Chapter 9

Post-processing of partially compensated imagery

A Shack-Hartmann wavefront sensor (SH-WFS) is one of the main components in the techniques of deconvolution from wavefront sensing (DWFS) and adaptive optics (AO) for reconstructing the phase. Although the phase reconstruction is smoother than the true phase, a good overall object phase is nevertheless achievable, supported by the results in Chapters 7 and 8 (see Figs. 7.12 and 7.13). However, the small error in estimating the phase affects the accuracy of the magnitude of the object spectrum. Consequently, the object estimate is incorrect.

Roggemann *et al.* recognised this problem with the magnitude estimation in DWFS and proposed the use of reference star measurements in addition to the measurements of the object [149]. This was proposed because in general the reference information used in the denominator in Eq. (9.3) has a higher SNR than the object information used in the numerator [29,62,148], meaning that the dominant source of errors in DWFS is in the numerator. However, as long as the estimation of $H_{\mathrm{obj}}(u,v)$ in the numerator is not accurate, the magnitude of the object still cannot be reconstructed accurately.

In this chapter an alternative solution based on recursive inverse filtering is proposed. The first advantage is that the reliance on reference star measurements is removed. Secondly, the effect of the inaccuracies in the numerator and the denominator is addressed collectively

209

by the recursive inverse filter, offering the potential to overcome the problem of estimating the object magnitude.

The problem addressed in this chapter is also present in adaptive optics (AO) when the Shack-Hartmann WFS is used. As a result, the solution proposed here is also applicable to AO imagery with simple modifications. Although simulations are not performed on AO images specifically, it is believed that such post-processing will also improve the quality of the images.

This chapter is structured as follows. Section 9.1 discusses the problems with conventional DWFS, and Section 9.2 outlines the steps involved in simulating DWFS data. The proposed parameterised filter is based on the NASRIF algorithm [102] which is presented in Section 9.3. Its adaptation to the problem of DWFS is addressed in Section 9.4. Section 9.5 presents simulation results using both the NASRIF and the adapted approach. Results obtained using conventional DWFS are used as a comparison. Conclusions are drawn in Section 9.6.

9.1 Problems with conventional DWFS

Deconvolution from wavefront sensing (DWFS) is a hybrid imaging technique which combines elements of adaptive optics and post-detection image processing to obtain high resolution images through atmospheric turbulence. In addition to speckle images, the technique of DWFS also utilises simultaneous wavefront sensing images, and the vector Wiener filter [133, 150] is used to recover the object distribution. This technique was first proposed by Fried [46] and later extended by Primot et al. [135]. Fig. 9.1 shows a block diagram of the DWFS method [147], and illustrates how the light of the distorted wavefront enters the aperture of the telescope and is split by the beam splitter. A proportion of the light is used to form a speckle image, while the rest of the light is directed to a Shack-Hartmann wavefront sensor (SH-WFS) where the measurements are used to estimate the instantaneous point spread function (PSF) in a post-processing step. The mathematical model of the system is detailed in Sections 6.2, 6.4 and 8.1.

The solution to DWFS proposed by Primot et al. [135] is a vector Wiener filter over an
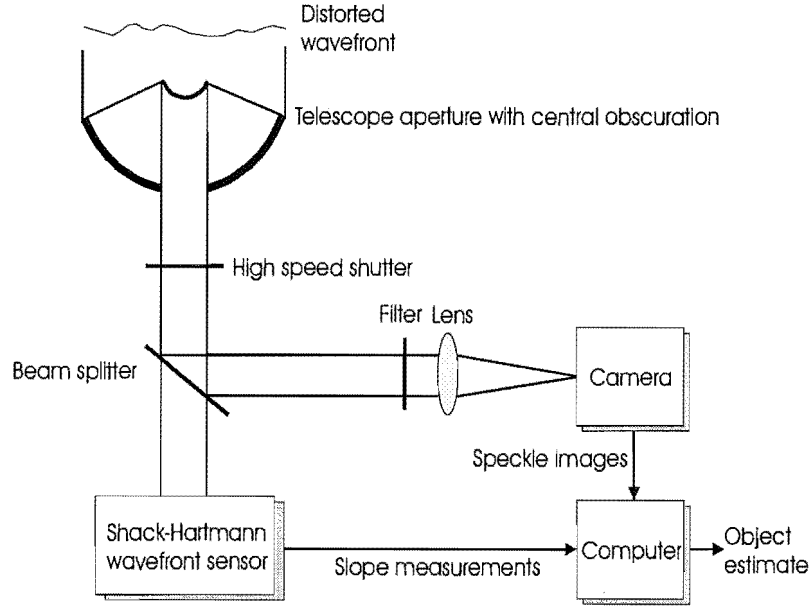
Figure 9.1: *Block diagram of the DWFS method.*

ensemble of images:

$$\hat{F}(u,v) = \frac{\langle D(u,v)\hat{H}^*(u,v)\rangle}{\langle|\hat{H}(u,v)|^2\rangle + c}. \tag{9.1}$$

This formulation is superior to applying the Wiener filter to individual frames before averaging because the division operation causes problems at the locations where the OTF is zero [135,180]. The SNR of the estimator is also improved by the averaging operation [147]. c can be a frequency-dependent function which attenuates the noisy high-frequency components of $\hat{F}(u,v)$ [147]; it is however set to a constant value here to simplify the algorithm.

A problem with this estimator arises from the different correlation properties of the true atmospheric phase $\phi_k(u,v)$ and the estimated phase $\hat{\phi}_k(u,v)$. Recall that the OTF $H(u,v)$ is an autocorrelation of the complex aperture (Eq. (6.14)). Therefore any error in estimating the phase manifests as an error in the OTF reconstruction. By rewriting Eq. (9.1) as

$$\hat{F}(u,v) = F(u,v)\frac{\langle H(u,v)\hat{H}^*(u,v)\rangle}{\langle|\hat{H}(u,v)|^2\rangle + c}, \tag{9.2}$$

it is apparent that if the reconstructed OTF $\hat{H}(u,v)$ does not equal the true OTF $H(u,v)$, the denominator of Eq. (9.1) is biased and does not provide object spectrum estimates $\hat{F}(u,v)$ with the proper amplitudes [149].
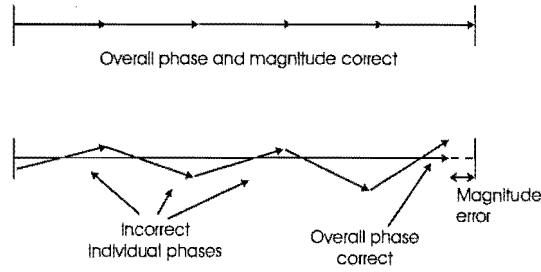
Figure 9.2: *An illustration of the effect of the phase error on the magnitude error in estimating $H(u, v)$.*
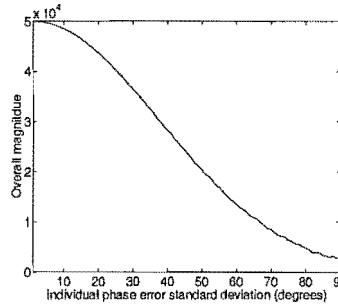
Figure 9.3: *The effect of individual phase errors on the overall magnitude error. In the absence of individual phase errors, the overall magnitude is 50000. This value is reduced significiantly as the individual phase errors increase.*

Another interpretation of how the magnitude problem arises is to view the individual elements of the ensemble in the numerator of Eq. (9.1) as vectors as shown in Fig. 9.2. In the absence of phase errors, the overall magnitude is estimated correctly. When errors are present in estimating the individual phases, averaging over a large ensemble eventually produces the correct overall phase, but not a correct magnitude. Fig. 9.3 quantifies the effect of the phase error on the magnitude error using 50000 frames. The phase error is modelled simply as a zero-mean Gaussian random variable with a standard deviation varying from 0 to 90 degrees. As the individual phase error increases, the magnitude error also increases. This error is significant and is directly linked to the core and halo structure in the reconstruction of a point source (Section 4.4.2).

Roggemann *et al.* [149] recognised this problem and proposed a solution which uses a reference star to obtain the quantities in the denominator. Eq. (9.1) can therefore be written

as:

$$\hat{F}(u,v) = \frac{\langle D(u,v)\hat{H}^*_{\mathrm{obj}}(u,v)\rangle}{\langle|\hat{H}_{\mathrm{ref}}(u,v)|^2\rangle + c}. \tag{9.3}$$

This was proposed because the reference information in general has a higher SNR than the object information in the numerator [29,62,148], thus the dominant source of errors in DWFS is in the numerator. However, as long as the estimation of $H_{\mathrm{obj}}(u,v)$ in the numerator is not accurate, the magnitude of the object still cannot be estimated correctly. This approach also imposes the need for measuring a reference star which introduces problems such as anisoplanatism (as discussed in Chapter 4) and a reduction in valuable observation time. Furthermore, it cannot be guaranteed that the atmospheric properties remain unchanged between the measurement of the reference star and the object of interest.

In view of the existing problems, the approach of estimating the object using a parameterised linear inverse filter in an iterative framework is investigated in this chapter. In this approach, the phase of the object is assumed to be reconstructed accurately, and the filter is used to correct for the inaccuracies of the magnitude as a result of both the denominator and the numerator collectively. This removes the need for a reference star. In fact, the recursive filter proves to be substantially independent of observed data and traditional atmospheric parameters (such as r_0) that the need for their measurement and estimation is removed.

The parameterised approach also offers the advantage of regularisation and significantly reduces the number of variables to be estimated. An existing algorithm which can act as a tool for the parameterisation is the NASRIF (Nonnegativity And Support constraint Recursive Inverse Filtering) algorithm first proposed by Kundur and Hatzinakos [102]. The linearity of this algorithm results in easy implementation and analysis.

Two variations of this approach are investigated. Firstly, the direct application of the NASRIF method is examined. This is essentially a pixel-based parameterisation of the reconstruction filter. Secondly, a global parameterisation is used to approximate the overall shape of the term $\langle|\hat{H}(u,v)|^2\rangle$. This approach has the advantage of strongly regularising the solution, a point which has not been dealt with in the original NASRIF algorithm. It should be emphasised that the global parameterised filter is used to provide a starting solution that is close to the true solution; the filter is then updated iteratively in the course of the algorithm to solve the problem.

9.2 Simulating DWFS

The motivation for simulating DWFS is to generate the quantities in the numerator of Eq. (9.1), and to facilitate the investigation of the shape of $\langle |\hat{H}(u,v)|^2 \rangle$ in the denominator for parameterisation. A simulation is preferrable to an analytic evaluation because the latter involves a four-dimensional integral which cannot be evaluated in closed form [147].

The simulation procedure is adapted from that in Ref. [147], a block diagram of which is shown in Fig. 9.4. A phase screen of a specified D/r_0 is generated using the method of Harding *et al.* [76]. The strength of the atmospheric turbulence used here includes $D/r_0 = 20, 15, 8$ and 5. The aperture (with or without the central obscuration, also referred to as circular and annular apertures) is then imposed on the phase screen, resulting in the complex wavefront.

For the generation of speckle images $d(x, y)$, the PSF is firstly generated from the complex wavefront, which is then convolved with the object and corrupted with Poisson noise. The noise levels are set to either 5000 or 100 photons per frame. While the former is optimistic for real data, the latter is representative of images captured at a telescope. The noiseless case is also investigated.

To generate wavefront sensing data for a reference star, the complex wavefront is put through a 16×16 Shack-Hartmann wavefront sensor. It is then corrupted with 15dB of Gaussian noise which reflects the noise level of real data. The noiseless case is also investigated for comparison. The centroids are calculated, which are used to reconstruct the phase $\hat{\phi}(u, v)$ using 200 Zernike polynomials. The estimate of the PSF $\hat{h}_{\text{ref}}(x, y)$ is then found from this reconstructed phase. The wavefront sensing data for the object is generated in a similar manner: the noiseless WFS data is convolved with the object before the noise corruption process. The reconstructed PSF is called $\hat{h}_{\text{obj}}(x, y)$. This process is repeated, and the relevant quantities are accumulated. A total of 2000 frames are used. The quantities generated using this method are used throughout this chapter.

The assumptions made in this simulation are as follows. Firstly, Nyquist sampling is employed. Secondly, it is assumed that Poisson noise is the dominant noise source in the speckle images, and Gaussian noise is the dominant noise in the WFS data. This is a valid assumption when compared to the characteristics of the real data processed in this
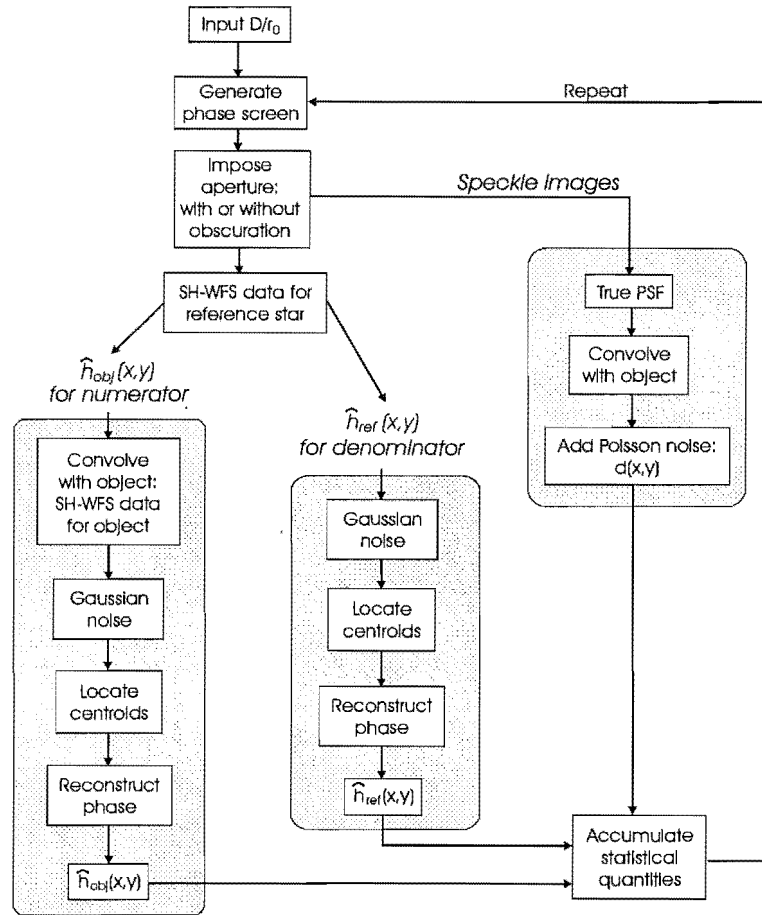
Figure 9.4: *Block diagram for simulating DWFS.*

---

thesis. For instance in Chapter 7, the WFS data has a sufficiently high light level that the assumption of Gaussian noise modelled the situation well.

Finally, it should be pointed out that the phase screens used to generate the reference star data are the same as those used for the object. This is not realistic for experimental data, but was done to ensure the best performance of the existing method of DWFS.

9.3 The NASRIF algorithm

The NASRIF algorithm is an image restoration procedure involving the filtering of the blurred image to obtain the original object. The principal advantages of this algorithm

are that explicit knowledge of the PSF is not required, and that the object estimate is a weighted linear sum of the data. This results in easy implementation and analysis. The convex cost function used also ensures the uniqueness of the solution provided that certain constraints are met.

This method enforces positivity and support constraints for the object, however there are few constraints on the PSF. In its originally proposed form the PSF and its inverse are only required to be absolutely summable. The inverse filter coefficients must be greater than (or equal to) zero, and must sum to unity. The object estimate $\hat{f}(x, y)$ is obtained by the convolution of the measurements $m(x, y)$ and the inverse filter $h_{\text{inv}}(x, y)$ to be estimated:

$$\hat{f}(x, y) = m(x, y) \odot h_{\text{inv}}(x, y) \tag{9.4}$$

where $h_{\text{inv}}(x, y)$ is an approximation to the inverse of the PSF $h^{-1}(x, y)$. To fulfil the constraints on the quantities, the error metric to be minimised is:

$$
\begin{aligned}
E \;=\; & \sum_{x,y} \hat{f}^2(x, y) \left[\frac{1 - \text{sgn}(\hat{f}(x, y))}{2} \right] S_f(x, y) + \sum_{x,y} \hat{f}^2(x, y)[1 - S_f(x, y)] \\
& + \; \gamma \left[\sum_{x,y} h_{\text{inv}}(x, y) - 1 \right]^2 \tag{9.5}
\end{aligned}
$$

where the sign function $\text{sgn}(\cdot)$ was defined in Section 2.2 and $S_f(x, y)$ is the support of the object. The three terms of the error metric penalise respectively the negative pixels within the support, the non-zero pixels outside the support, and the deviation of the sum of the filter coefficients from one. This error metric adopts a slightly different notation from that in Ref. [102], but they are mathematically equivalent.

To obtain an object estimate, the filter $h_{\text{inv}}(x, y)$ is updated using a conjugate gradient procedure until a pre-defined criterion is met. In the original paper, it was pointed out that because the problem is not regularised, visual inspection is sometimes required for terminating the algorithm. Although this is a valid form of regularisation, in a later paper by Ng et al. [128], eigenvalue truncation and total variation are used to improve the regularisation of the problem.

The major disadvantage of this algorithm is that the PSF is insufficiently constrained. Even if constraints on the PSF are available, they cannot be easily converted to constraints on the inverse of the PSF. The fact that the inverse filter $h_{\text{inv}}(x, y)$ does not have a specified size means that the quality and the computational complexity of the solution are dependent

on the filter size. The size of this optimal filter must be found by experimentation. The object estimate is also relatively sensitive to the support size. This can be overcome by a support-finding algorithm as in the original paper, however it is computationally intensive and a more deterministic approach would be preferrable.

The implementation of the NASRIF algorithm requires the calculation of the derivative of the error metric in Eq. (9.5) with respect to the filter $h_{\text{inv}}(x,y)$:

$$\frac{\partial E}{\partial \hat{h}_{\text{inv}}(x,y)} = 2\left\{\tfrac{1}{2}[1 - \text{sgn}(\hat{f}(x,y))]\hat{f}(x,y)\right\}S_f(x,y) \otimes m(x,y)$$
$$+2\hat{f}(x,y)[1 - S_f(x,y)] \otimes m(x,y)$$
$$+2\left\{[\textstyle\sum_{x,y}\hat{h}_{\text{inv}}(x,y)] - 1\right\}. \tag{9.6}$$

The centre portion of this quantity (of the same size as $h_{\text{inv}}(x,y)$) is used to update the filter using a conjugate gradient procedure, the implementation of which is detailed in Appendix A.

9.4   Applying NASRIF to DWFS

The NASRIF algorithm can be applied to DWFS by firstly rewriting Eq. (9.3) in the spatial domain:

$$\hat{f}(x,y) = \langle d(x,y) \otimes \hat{h}_{\text{obj}}(x,y)\rangle \odot \mathfrak{F}^{-1}\left\{\frac{1}{\langle|\hat{H}_{\text{ref}}(u,v)|^2\rangle + c}\right\}$$
$$= \langle d(x,y) \otimes \hat{h}_{\text{obj}}(x,y)\rangle \odot h_{\text{inv}}(x,y) \tag{9.7}$$

where $\mathfrak{F}^{-1}$ is the inverse Fourier transform operation. Here $m(x,y) = \langle d(x,y) \otimes \hat{h}_{\text{obj}}(x,y)\rangle$ which is equivalent to the numerator in the Fourier domain formulation. The quantity $1/\langle|\hat{H}_{\text{ref}}(u,v)|^2\rangle + c$ is called the Fourier filter. The inverse Fourier transform of the Fourier filter is called the spatial filter, denoted by $h_{\text{inv}}(x,y)$.

The significance is that since Eq. (9.7) takes the same form as Eq. (9.4), the spatial filter can be represented by an arbitrary inverse filter $h_{\text{inv}}(x,y)$, and a standard optimisation technique can be used to find $\hat{f}(x,y)$. This solution is essentially a pixel-based parameterisation approach to the problem of magnitude estimation in DWFS.

The essence of this approach is that the object phase is assumed to be reconstructed accurately by the numerator; the only source of error is in the magnitude estimation of the

object. By using a recursive inverse filter which is capable of compensating for the magnitude errors, a fully compensated estimate can be obtained.

Since the inverse filter of the NASRIF algorithm is usually of a much smaller size than the dimension of the quantity to be parameterised, the number of unknowns in the solution is reduced significantly. This is supported by the compactness of the spatial filter as shown in Section 9.4.1.

9.4.1   Global parameterisation

Whilst the NASRIF algorithm parameterises the solution to DWFS in terms of pixels, an alternative of the filter $h_{inv}(x, y)$ can be derived by globally parameterising the Fourier filter or the spatial filter. Once again, the parameters can be updated iteratively to give an estimate of the object. Here, because the entire Fourier (or spatial) filter is parameterised with a small number of parameters, the effective degrees of freedom of the filter is reduced and the solution is regularised.

To test whether the Fourier filter or the spatial filter can be parameterised, the quantities in the denominator of Eq. (9.3) are examined. Fig. 9.5 illustrates the Fourier filters for a telescope. Only the cross-section is shown here because these functions are essentially radially symmetric. The strength of the atmospheric turbulence investigated are $D/r_0 = 20, 15, 8$ and 5, corresponding to the solid, dashed, dash-dotted and dotted lines respectively. The top and bottom rows show the Fourier filters in the absence and presence of a central obscuration in the telescope aperture respectively. The effect of the presence of spiders (support for the secondary mirror) is assumed to be negligible and is not investigated here. The left-hand column illustrates the ideal case where no noise is present on the wavefront sensor. In the right-hand column, 15dB white Gaussian noise is present on the WFS. It can be seen that in the presence of an obscuration, the central frequency components are significantly attenuated. In general, the filter becomes broader as D/r_0 decreases. In all cases the filters are smooth functions and are comparatively easy to parameterise.

The corresponding spatial filters are shown in Fig. 9.6. Only the centre portion is shown because the filters are compact. Although the Fourier filter shapes are different for different parameters, the spatial filters are all essentially high-pass filters.

(a)   (b)

(c)   (d)
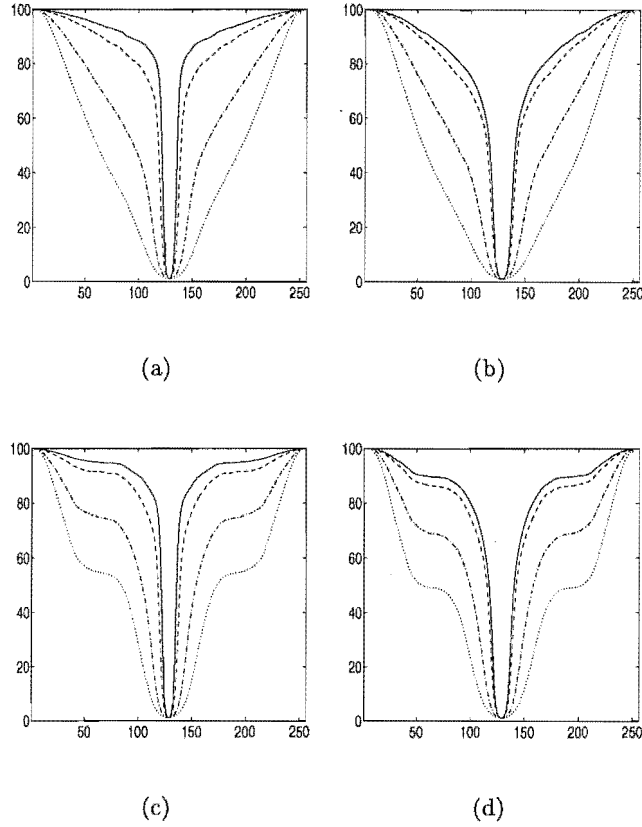
Figure 9.5: *Fourier filter shapes for $D/r_0 = 20, 15, 8$ and 5, represented by solid, dashed, dash-dotted and dotted lines. (a) No central obscuration, ideal case with no noise on the WFS. (b) No central obscuration, 15dB noise on WFS. (c) With central obscuration, ideal case with no noise on the WFS. (d) With central obscuration, 15dB noise on WFS.*

Since the filter in the absence of the central obscuration can be considered as a subset of the filter when the obscuration is present, the filter is parameterised based on the latter. The Fourier filter instead of the spatial filter is chosen to be parameterised because its shape appears to be more well-defined and mathematically tractible. Since the Fourier transform is a linear operation, the parameterisation can easily be transformed between the Fourier and the spatial domains.

The proposed parameterisation of the Fourier filter is a 10-parameter radially symmetric

(a)　　　　　　　　　(b)

(c)　　　　　　　　　(d)
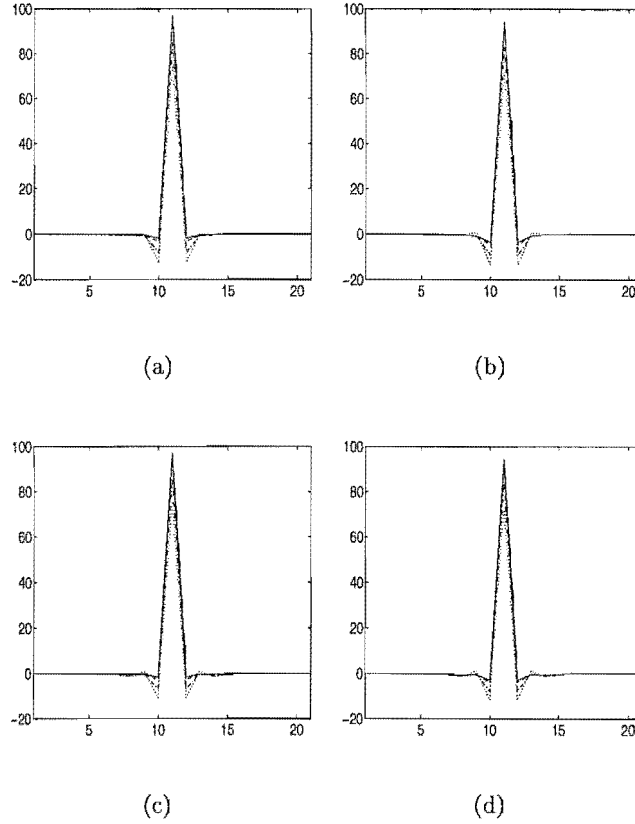
Figure 9.6: *Spatial filter shapes for $D/r_0 = 20, 15, 8$ and 5, represented by solid, dashed, dash-dotted and dotted lines. These correspond to the Fourier filters shown in Fig. 9.5. Only the centre portions are shown here. (a) No central obscuration, ideal case with no noise on the WFS. (b) No central obscuration, 15dB noise on WFS. (c) With central obscuration, ideal case with no noise on the WFS. (d) With central obscuration, 15dB noise on WFS.*

function:

$$\mathfrak{F}\{h_{\text{inv}}(r)\}(\rho) = K_1 - K_2 \exp\{-K_3(\rho^2)\} -$$

$$- K_4 \exp\{K_5(\rho^2)\} \odot [\text{circ}(K_6(\rho^2)) - K_7 \text{circ}(K_8(\rho^2))] - K_9 \exp\{-K_{10}(\rho^2)\} \quad (9.8)$$

where $\rho = \sqrt{u^2 + v^2}$ and $\text{circ}(\cdot)$ is the circle function defined in Section 2.2. The spatial filter model is the inverse Fourier transform of the Fourier filter model:

$$h_{\text{inv}}(r) = k_1 \delta - k_2 \exp\{-k_3(r^2)\} -$$

$$- k_4 \exp\{-k_5(r^2)\} \left(\frac{J_1(k_6 r)}{k_6 r} - k_7 \frac{J_1(k_8 r)}{k_8 r} \right) - k_9 \exp\{-k_{10}(r^2)\} \quad (9.9)$$

where J_1 is the Bessel function of the first kind, order 1, and $r = \sqrt{x^2 + y^2}$. It is important to note that the parameters K_1 to K_{10} in the Fourier filter are different but are uniquely related to the parameters k_1 to k_{10} in the spatial filter. Moreover, the parameterisation presented here is only one of many possibilities, and other models can be used. A question that could be raised here is whether the parameterisation should include parameters that describe atmospheric properties, such as r_0. In Fig. 9.5, it is shown that the shape of the Fourier filter does alter with a different r_0, but this is adequately described by the parameterisation in Eq. (9.9). Moreover, as verified with simulations in Section 9.5, the fact that atmospheric parameters are not used in the parameterisation does not affect the solution. This is advantageous because it removes the need to estimate such parameters and makes the processing more generic.

The following comparison examines whether the parameterisation is a good approximation to the actual functions. The case corresponding to $D/r_0 = 15$ is used because this gives approximately $D/r_0 = 1$ per lenslet for the 16×16 Shack-Hartmann array used here, and is representative of real seeing conditions. 15dB noise is assumed on the WFS.

In Fig. 9.7(a), the cross-section of the actual Fourier filter and its parameterisation are shown as dash-dotted and solid lines respectively. This is for the case where no central obscuration is present in the aperture. Fig. 9.7(b) shows the corresponding quantities for an aperture with a central obscuration. The parameters K_1 to K_{10} and c for each case are tabulated in Table 9.1. The corresponding spatial filters are shown in Figs. 9.7(c) and 9.7(d) respectively. It can be seen that there is good agreement between the parameterisations and the true functions, especially in the spatial domain. This means that the starting estimate of the object is close to the true solution. During the course of the algorithm, the filter is updated iteratively to find the solution.

Common to all parameterisation problems, there exists a tradeoff between the complexity of the parameterisation, and the actual degrees of freedom of the underlying function. It is generally accepted that there is a danger in overparameterising a function, where the number of parameters used is greater than the function's realistic degrees of freedom. In this application, the choice of ten parameters was found appropriate.

Parameters	No central obscuration	With central obscuration
K_1	100	100
K_2	62	75
K_3	0.006	0.004
K_4	0.005	0.016
K_5	0.0035	0.004
K_6	0.36	0.38
K_7	-0.7666	0.3
K_8	0.0781	0.2
K_9	30	15
K_{10}	2.4×10^{-4}	5×10^{-4}
c	0.01	0.01

Table 9.1: *Parameters for the Fourier filter for a telescope aperture in the absence and presence of a central obscuration. These correspond to the plots in Fig. 9.7(a) and (b).*

(a)   (b)

(c)   (d)
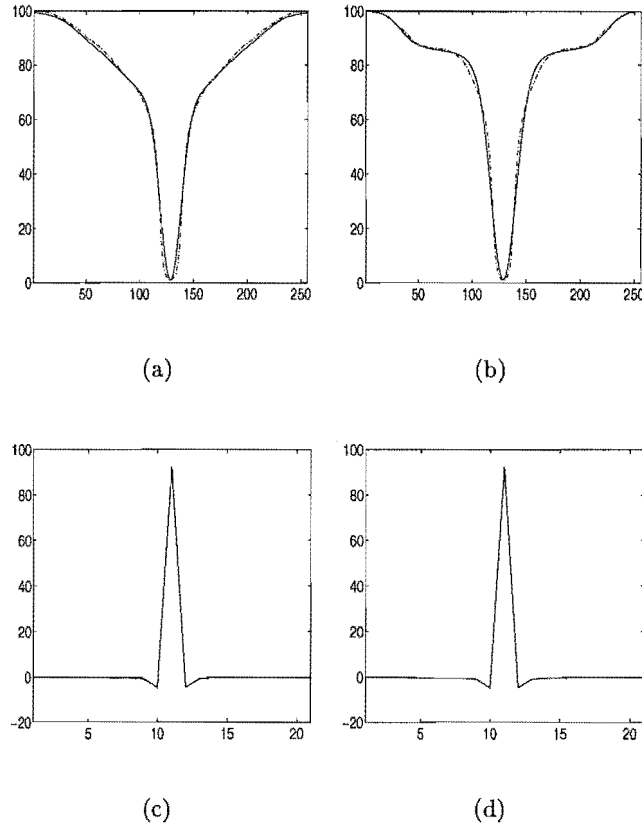
Figure 9.7: *Comparison between the parameterisations and the true functions for $D/r_0 = 15$. (a,b) Fourier filters (dash-dotted lines) and their parameterisations (solid lines) in the absence and presence of a central obscuration. (c,d) Corresponding spatial filters in (a) and (b).*

9.4.2 Gradient calculations

In order to optimise the filter coefficients of $h_{inv}(x, y)$ to obtain an object estimate $\hat{f}(x, y)$ for the global parameterisation, the gradients of the cost function with respect to the filter parameters are required. These are calculated via the chain rule:

$$\frac{\partial E}{\partial k_i} = \frac{\partial E}{\partial h_{inv}(x, y)} \frac{\partial h_{inv}(x, y)}{\partial k_i} \quad i = 1 \ldots 10. \tag{9.10}$$

The first term is shown in Eq. (9.6). The differentiation of the filter with respect to its 10 coefficients are as follows:

$$\frac{\partial h_{\text{inv}}(x,y)}{\partial k_1} = 1$$

$$\frac{\partial h_{\text{inv}}(x,y)}{\partial k_2} = \sum_r -\exp\{-k_3 r^2\}$$

$$\frac{\partial h_{\text{inv}}(x,y)}{\partial k_3} = \sum_r k_2 r^2 \exp\{-k_3 r^2\}$$

$$\frac{\partial h_{\text{inv}}(x,y)}{\partial k_4} = \sum_r -\exp\{-k_5 r^2\}\left(\frac{J_1(k_6 r)}{k_6 r} - k_7\frac{J_1(k_8 r)}{k_8 r}\right)$$

$$\frac{\partial h_{\text{inv}}(x,y)}{\partial k_5} = \sum_r k_4 r^2 \exp\{-k_5 r^2\}\left(\frac{J_1(k_6 r)}{k_6 r} - k_7\frac{J_1(k_8 r)}{k_8 r}\right)$$

$$\frac{\partial h_{\text{inv}}(x,y)}{\partial k_6} = \sum_r -k_4 \exp\{-k_5 r^2\}\left(-\frac{J_1(k_6 r)}{k_6 r} + \frac{J_0(k_6 r)}{2k_6} - \frac{J_2(k_6 r)}{2k_6}\right)$$

$$\frac{\partial h_{\text{inv}}(x,y)}{\partial k_7} = \sum_r k_4 \exp\{-k_5 r^2\}\frac{J_1(k_8 r)}{k_8 r}$$

$$\frac{\partial h_{\text{inv}}(x,y)}{\partial k_8} = \sum_r k_4 k_7 \exp\{-k_5 r^2\}\left(-\frac{J_1(k_8 r)}{k_8^2 r} + \frac{J_0(k_8 r)}{2k_8} - \frac{J_2(k_8 r)}{2k_8}\right)$$

$$\frac{\partial h_{\text{inv}}(x,y)}{\partial k_9} = \sum_r -\exp\{-k_{10} r^2\}$$

$$\frac{\partial h_{\text{inv}}(x,y)}{\partial k_{10}} = \sum_r k_9 r^2 \exp\{-k_{10} r^2\} \tag{9.11}$$

where J_0 and J_2 are the first order Bessel functions of the zeroth and second kinds respectively. The filter parameters are updated with the method of steepest descent [116] with a line search routine.

9.5 Simulations and results

This section presents the results to compare the DWFS technique to a parameterised solution, the latter of which is divided into global and pixel-based (NASRIF) parameterisations. 2000 frames are used in generating the relevant quantities according to the procedure in Section 9.2. Same as the other investigations in this chapter, the values of D/r_0 are 20, 15, 8 and 5. The three noise levels in the speckle images are the noiseless case, and the cases where each of the 2000 frames has 5000 or 100 photons. For the wavefront sensing data, the noiseless situation and the presence of 15dB white Gaussian noise are investigated. It

should be noted that the object and the reference star measurements are generated using the same complex fields. This presents the best-case scenario where the usual problems of using a reference star such as anisoplanatism [147] are not present, and allows the DWFS method to be compared under the most favourable circumstances.

9.5.1 Reconstructing a binary object

The binary object to be reconstructed has a separation of 10 pixels and a intensity ratio of 1.25. Assuming the mean wavelength of the incoming wavefront is 650nm and the telescope is of 1-metre diameter, each pixel in the focal plane corresponds to 0.067 arcsec in a Nyquist sampled system, and angular resolutions of 0.164 and 0.132 arcsec are achievable for a circular and an annular aperture respectively. According to Rayleigh's resolution criterion, the angular resolution of a circular aperture is $1.22\lambda/D$. However, the presence of a central obscuration acts as a highpass filter and increases the angular resolution [13]. For an annular aperture with a central obscuration of diameter $0.4D$ (which is the case simulated here), the resolving power is approximately $0.98\lambda/D$, i.e. about 20% higher than its circular counterpart.

Fig. 9.8 shows the numerator quantities $\langle d(x,y) \otimes \hat{h}_{obj}(x,y) \rangle$ for the different parameters in the absence of a central obscuration in the aperture. Only the centre portions are shown here for clarity. The respective quantities in the presence of the obscuration are not shown as they are very similar to those in Fig. 9.8. The four rows correspond to $D/r_0 = 20, 15, 8$ and 5; while the three columns represent the noiseless case and the cases with 5000 and 100 photons per frame. In all cases, the binary object is visible when no noise is present. In the presence of noise, the binary object is not visible for $D/r_0 = 20$ and 15. It can be seen that there is a broad halo over the object for all cases.

Simulations with a circular aperture

The results of applying the method of DWFS are tabulated in Table 9.2. Table 9.3 shows the global parameterisation results, while Table 9.4 has the results using the original NASRIF algorithm. These tables correspond to the case where no central obscuration is present on the aperture.
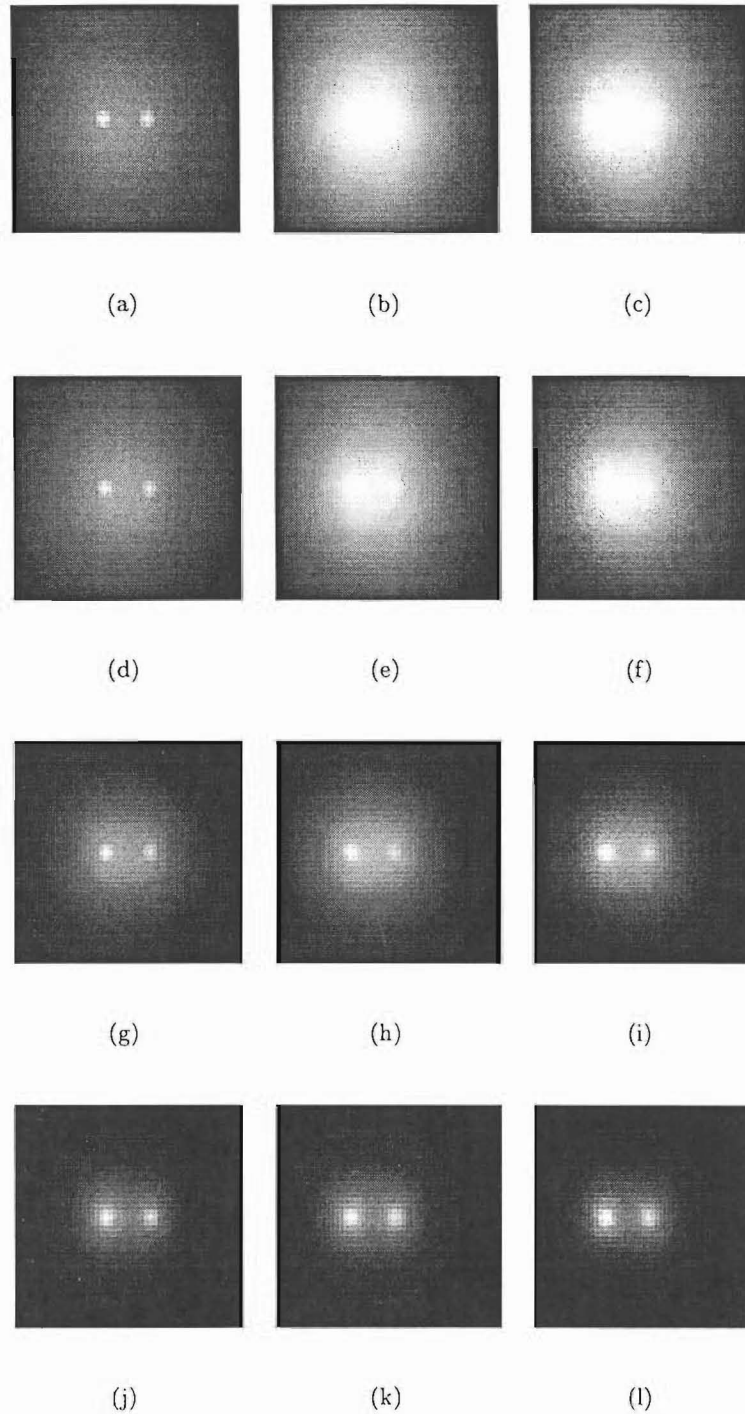
Figure 9.8: *The numerator quantities $\langle d(x,y) \otimes \hat{h}_{obj}(x,y) \rangle$ for (a) $D/r_0 = 20$, no speckle noise; (b) $D/r_0 = 20$, 5000 photons per speckle image; (c) $D/r_0 = 20$, 100 photons per speckle image; (d) $D/r_0 = 15$, no speckle noise; (e) $D/r_0 = 15$, 5000 photons per speckle image; (f) $D/r_0 = 15$, 100 photons per speckle image; (g) $D/r_0 = 8$, no speckle noise; (h) $D/r_0 = 8$, 5000 photons per speckle image; (i) $D/r_0 = 8$, 100 photons per speckle image; (j) $D/r_0 = 5$, no speckle noise; (k) $D/r_0 = 5$, 5000 photons per speckle image; (l) $D/r_0 = 5$, 100 photons per speckle image. Only the centre portions have been shown for clarity.*

Parameters	Noiseless	5000 photons	100 photons
$D/r_0 = 20$			
Separation	9.99	9.01	-
Intensity	1.23	1.06	-
$D/r_0 = 15$			
Separation	9.99	9.97	9.97
Intensity	1.24	1.18	1.16
$D/r_0 = 8$			
Separation	10.00	10.00	10.00
Intensity	1.25	1.26	1.22
$D/r_0 = 5$			
Separation	10.00	10.00	10.00
Intensity	1.25	1.25	1.24

Table 9.2: *Simulation results for the DWFS method. The object is a binary star with a separation of 10 pixels and an intensity ratio of 1.25. No central obscuration is present in the aperture.*

Fig. 9.9 shows an example of the reconstructions. Here $D/r_0 = 15$ and each of the frames has 100 photons. Fig. 9.9(a) shows the diffraction-limited object (in the absence of a central obscuration) for comparison. Only the centre portions have been shown for clarity. It is apparent that the visual difference between the methods is not significant.

On examining the tabulated results, it is shown that all the results are comparable to each other in terms of separation and intensity ratio, and none of the methods succeeded in reconstructing a binary star with $D/r_0 = 20$ and the photon noise level of 100 photons per frame. The fact that the methods based on parameterisation have comparable performance to the DWFS method is important because the need of a reference star measurement and its associated problems can be removed with little loss of accuracy in the solution. Moreover, the independence of the parameterisations from atmospheric parameters removes the need to estimate such parameters. This validates the initial assumption of not incorporating them into the model, which also has the advantage of reducing a possible source of modelling error.

The phase screens used for the generation of the reference star measurements are also used to generate the data for the binary star in order to optimise the performance of DWFS. In

Parameters	Noiseless	5000 photons	100 photons
$D/r_0 = 20$			
Separation	9.99	9.00	-
Intensity	1.23	1.04	-
Iterations	5	500	-
Total time (s)	9.06×10^3	7.48×10^5	-
$D/r_0 = 15$			
Separation	9.99	9.97	9.96
Intensity	1.23	1.14	1.12
Iterations	313	500	500
Total time (s)	4.18×10^5	7.46×10^5	1.06×10^6
$D/r_0 = 8$			
Separation	9.97	9.96	9.97
Intensity	1.22	1.21	1.19
Iterations	2	2	2
Total time (s)	2.38×10^3	2.50×10^3	3.47×10^3
$D/r_0 = 5$			
Separation	10.00	10.00	10.00
Intensity	1.25	1.25	1.24
Iterations	2	2	2
Total time (s)	5.04×10^3	3.57×10^3	3.56×10^3

Table 9.3: *Simulation results for the method of global parameterisation. The object is a binary star with a separation of 10 pixels and an intensity ratio of 1.25. A support of 31 × 36 is used. No central obscuration is present in the aperture.*

reality, the use of a reference star for DWFS introduces problems such as anisoplanatism and its performance would deteriorate. On the other hand, the parameterised methods do not rely on reference star measurements, and as a consequence the performance of these approaches is not affected.

The global parameterisation method produces more consistent results compared to the NASRIF algorithm. This is because the latter is sensitive to the filter and support sizes. Although good reconstructions have been obtained, an exhaustive search is required to find the optimal filter and support sizes which produce the best result. This has not been

Parameters	Noiseless	5000 photons	100 photons
$D/r_0 = 20$			
Separation	9.98	9.01	-
Intensity	1.05	1.01	-
Iterations	55	302	-
Total time (s)	4.40×10^3	2.37×10^4	-
Filter size	5	5	-
$D/r_0 = 15$			
Separation	9.99	9.02	9.95
Intensity	1.23	1.10	1.10
Iterations	5332	300	1500
Total time (s)	2.54×10^5	2.48×10^4	1.21×10^5
Filter size	5	9	5
$D/r_0 = 8$			
Separation	9.97	9.99	9.99
Intensity	1.16	1.26	1.22
Iterations	500	7263	13903
Total time (s)	4.00×10^4	3.01×10^5	5.61×10^5
Filter size	5	5	5
$D/r_0 = 5$			
Separation	10.00	9.99	9.99
Intensity	1.25	1.22	1.17
Iterations	19838	20000	20000
Total time (s)	5.20×10^5	8.26×10^5	7.19×10^5
Filter size	5	5	5

Table 9.4: *Simulation results for the original NASRIF method. The object is a binary star with a separation of 10 pixels and an intensity ratio of 1.25. The support used here is of size 31 × 36. No central obscuration is present in the aperture.*

performed in this set of simulations.

The variation in the per-iteration time for the parameterised methods is due to the line search process as well as the different speeds of the processors in different computers. Although each iteration of the NASRIF method takes less time than the global parameterisation method, many more iterations are required for convergence. However, it should be noted that the initial filter estimate of the global parameterisation is close to the true solution, whereas for the NASRIF algorithm, the best starting estimate of the filter is a uniform function so as not to bias the solution. This partially accounts for the greater number of iterations required for the NASRIF algorithm. For both methods, early termination of the algorithm is sometimes necessary, as the visual result has obviously converged, but the error metric is still decreasing gradually.

The convexity of the cost function in the NASRIF algorithm only applies when all the filter coefficients are greater than or equal to zero [102]. In the reconstructions here, this criterion is not met because the filters are highpass. As a result of this, the cost function is no longer unimodal, and stagnation of the algorithm can occur. However, this can be overcome algorithmically by using a partial conjugate gradient (instead of a conjugate gradient) procedure, where the gradient is reset every 10 iterations. This also speeds up the convergence of the algorithm.

An important point must be made regarding the support of the object estimate. Since pixels outside the support are forced to be zero, the support size determines the distribution of the noise. This means that if a tight support is used, the noise is concentrated within the support. In this case, it is better to trade off the tightness of the support with the noise level in the reconstruction.

Simulations with an annular aperture

The simulations have also been performed assuming that the telescope aperture has a central obscuration, resulting in an annular aperture. Fig. 9.10(a) shows the diffraction-limited image of the binary star. It can be seen that there is a ring around the core of each of the point sources. For clarity, the corresponding cross-section is shown in Fig. 9.10(b).
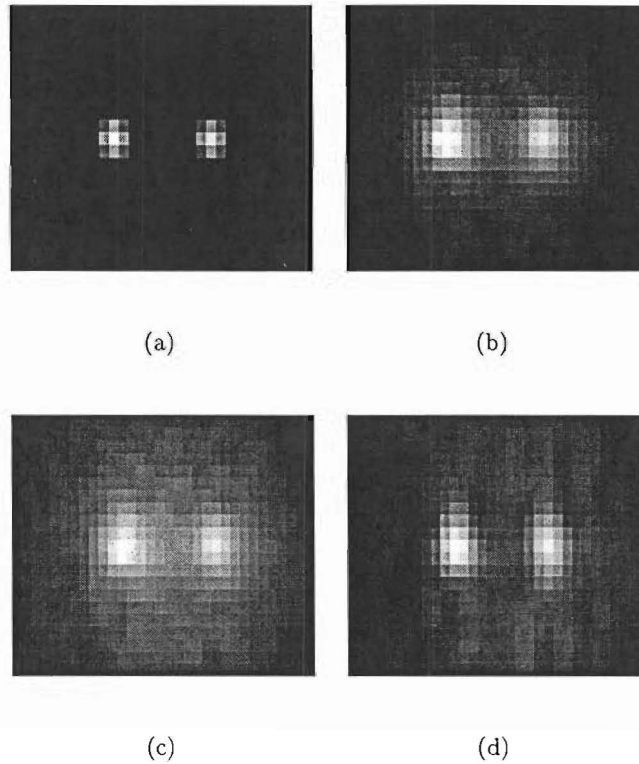
(a) (b)

(c) (d)

Figure 9.9: *Comparison between the true object and the object reconstructions using different methods. The true object is a binary with a separation of 10 pixels and intensity ratio 1.25. Here $D/r_0 = 15$ and each frame has 100 photons. (a) Diffraction-limited object. (b) Reconstruction with DWFS. (c) Reconstruction with global paramterisation. (d) Reconstruction with NASRIF, or pixel-based parameterised filter.*

The reconstructions using the three methods produced similar tabulated results to Tables 9.2, 9.3 and 9.4, and are not produced here. As discussed earlier, due to the increased resolving power with an annular aperture, the frequency content of the reconstructions is higher than that using the circular aperture. This is particularly noticeable for the noiseless case at all values of D/r_0 when using the global parameterisation method. These reconstructions are shown in Fig. 9.11. It should be pointed out that there is a halo over the object in Fig. 9.11(c), which could be attributed to the algorithm not having completely converged. The DWFS method also produced such reconstructions for all noiseless cases and for $D/r_0 = 5$. The NASRIF method managed to recover some high-frequency information for the case of $D/r_0 = 5$ at different noise levels, but the results are not as good as those of the other two methods. These reconstructions are shown in Fig. 9.12 with the corresponding

filters.

(a)   (b)
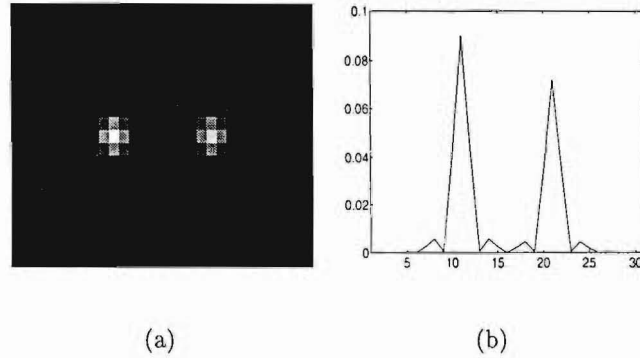
Figure 9.10: *(a) Diffraction-limited image of a binary star imaged through an annular aperture. Its cross-section is shown in (b) for clarity.*

The attainment of almost diffraction-limited reconstructions using the post-processing method of global parameterisation is significant since it offers an alternative to observing a reference star for DWFS. As discussed earlier, the DWFS data has been generated assuming the most favourable conditions; the performance of DWFS is therefore likely to deteriorate when these factors are taken into account. Valuable observation time is also saved as a result.

The NASRIF method did not manage to produce results comparable to the global parameterisation method is attributable to two factors. Firstly, the NASRIF method is sensitive to the filter and support sizes. Since these parameters were not optimised but were fixed for all reconstructions to reduce the computational complexity, the performance may have been sub-optimal. Secondly, the global parameterisation method is well-regularised and has been designed to suit the problem of DWFS specifically, whereas the NASRIF method is a generic method. It is therefore believed that if the filter and support sizes were optimised in the NASRIF method, its performance can be improved. The optimisation of the support size would also further improve the performance of the global parameterisation method.

(a) (b)

(c) (d)
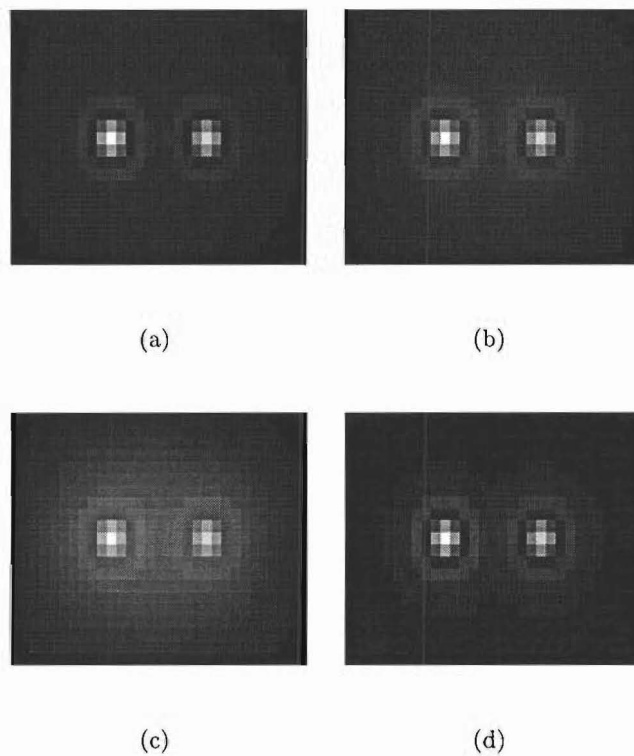
Figure 9.11: *The reconstruction of the binary star in the presence of an annular aperture. Near diffraction-limited features are visible using the global parameterisation method for the noiseless cases. The figures correspond to D/r_0 being (a) 20, (b) 15, (c) 8 and (d) 5.*

(a)

(b)

(c)
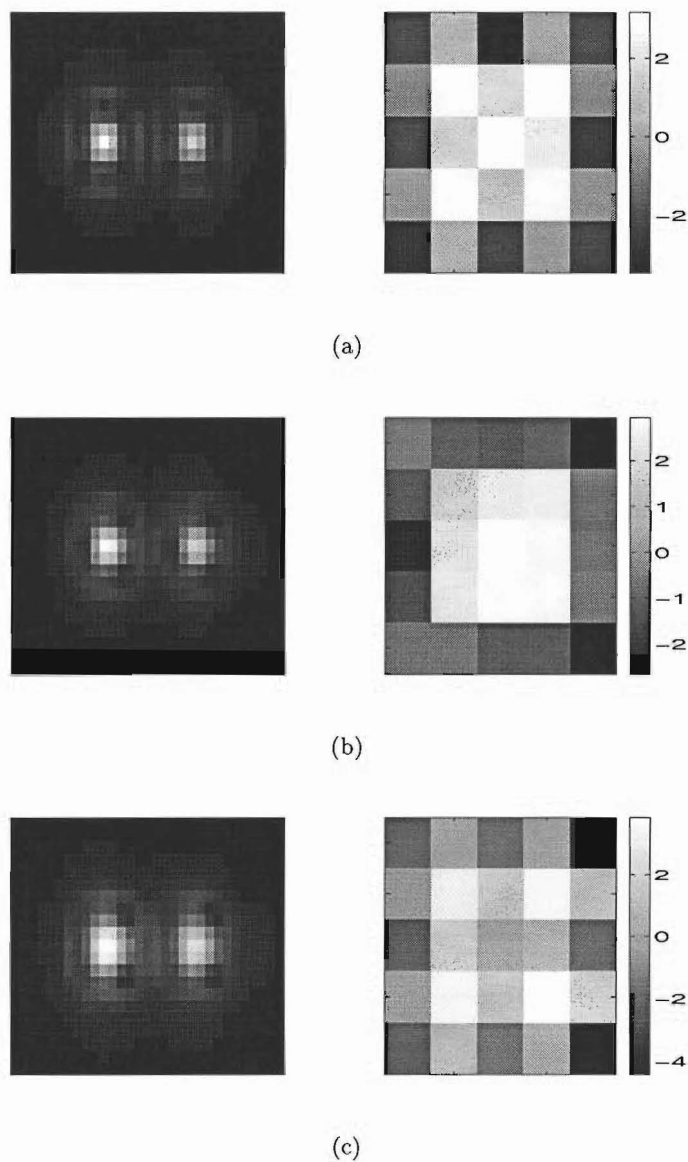
Figure 9.12: *The reconstruction of the binary star in the presence of an annular aperture. Some high-frequency information has been recovered using the pixel-based parameterisation (NASRIF method), but the diffraction limit is not reached. The reconstructions correspond to $D/r_0 = 5$ for (a) no Poisson noise, (b) 5000 photons per frame, and (c) 100 photons per frame. The corresponding filters are shown in the right-hand column.*

9.5.2 Reconstructing an extended object

The algorithms are also tested with a simulated extended object under the same noise conditions. Fig. 9.13(a) shows the original satellite object, while Fig. 9.13(b) shows the diffraction-limited image when using a circular aperture.

The results of applying the DWFS method, the global parameterisation method and the original NASRIF method are tabulated in Tables 9.5, 9.6 and 9.7 respectively. The normalised mean-squared error (MSE) is calculated with reference to the true object. The reconstructions are shown in Figs. 9.14, 9.15 and 9.16. Both the numerical and visual results must be interpreted together to provide a correct conclusion.

The global parameterisation produced the lowest MSE values. However, on examination of the reconstructions, the DWFS method has superior results for the cases of $D/r_0 = 20$ and 15. The lower MSE values for the global parameterisation reconstructions can be attributed to the fact that they are smoother and more spreadout, meaning that the error is distributed. On the other hand, the error for the DWFS reconstructions is more concentrated, resulting in a higher MSE value. For the cases where $D/r_0 = 8$ and 5, the two methods produced visually comparable reconstructions. The NASRIF method failed to match the performance of the other two algorithms. Once again, this is mainly due to its reliance on the knowledge of the optimal filter and support sizes, which are not available here.

(a)   (b)
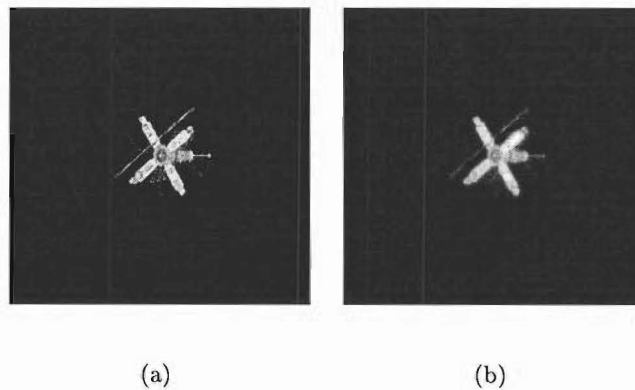
Figure 9.13: *The simulated extended object. (a) The original object and (b) the diffraction-limited image.*

Parameters	Noiseless	5000 photons	100 photons
	$D/r_0 = 20$		
MSE	5.68×10^{-4}	4.49×10^{-4}	4.54×10^{-4}
c	1×10^{-5}	1×10^{-2}	1×10^{-1}
	$D/r_0 = 15$		
MSE	5.73×10^{-4}	5.01×10^{-4}	5.62×10^{-4}
c	1×10^{-5}	1×10^{-2}	5×10^{-2}
	$D/r_0 = 8$		
MSE	5.11×10^{-4}	5.42×10^{-4}	5.66×10^{-4}
c	1×10^{-3}	1×10^{-2}	5×10^{-2}
	$D/r_0 = 5$		
MSE	5.67×10^{-4}	5.21×10^{-4}	5.14×10^{-4}
c	1×10^{-3}	1×10^{-2}	1×10^{-1}

Table 9.5: *Normalised mean-squared errors (MSE) for the DWFS method on the extended object. No central obscuration is present in the aperture.*

Parameters	Noiseless	5000 photons	100 photons
$D/r_0 = 20$			
MSE	4.01×10^{-4}	4.19×10^{-4}	4.81×10^{-4}
Iterations	2	2	3
Total time (s)	4.54×10^3	4.53×10^3	6.79×10^3
Support	106×121	106×121	106×121
$D/r_0 = 15$			
MSE	3.67×10^{-4}	3.79×10^{-4}	3.67×10^{-4}
Iterations	2	3	3
Total time (s)	4.03×10^3	7.16×10^3	6.94×10^3
Support	96×111	106×121	106×121
$D/r_0 = 8$			
MSE	3.09×10^{-4}	2.99×10^{-4}	4.13×10^{-4}
Iterations	3	4	5
Total time (s)	7.31×10^3	9.92×10^3	1.25×10^4
Support	106×121	106×121	106×121
$D/r_0 = 5$			
MSE	1.81×10^{-4}	2.78×10^{-4}	3.53×10^{-4}
Iterations	3	3	3
Total time (s)	6.47×10^3	6.11×10^3	6.15×10^3
Support	91×95	96×111	96×111

Table 9.6: *Normalised mean-squared errors (MSE) for the global parameterisation method on the extended object. No central obscuration is present in the aperture.*

Parameters	Noiseless	5000 photons	100 photons
	$D/r_0 = 20$		
MSE	9.14×10^{-4}	5.28×10^{-4}	4.84×10^{-4}
Iterations	15240	20000	20000
Total time (s)	5.18×10^5	7.56×10^5	6.29×10^5
	$D/r_0 = 15$		
MSE	3.60×10^{-4}	3.84×10^{-4}	4.10×10^{-4}
Iterations	11060	20000	20000
Total time (s)	8.14×10^4	7.62×10^5	8.73×10^5
	$D/r_0 = 8$		
MSE	2.48×10^{-4}	2.84×10^{-4}	3.53×10^{-4}
Iterations	11300	20000	20000
Total time (s)	4.81×10^5	9.07×10^5	7.10×10^5
	$D/r_0 = 5$		
MSE	1.62×10^{-4}	2.07×10^{-4}	2.52×10^{-4}
Iterations	48100	40000	38500
Total time (s)	1.69×10^6	2.76×10^6	1.61×10^6

Table 9.7: *Normalised mean-squared errors (MSE) for the NASRIF method on the extended object. No central obscuration is present in the aperture. The filter size is 9 × 9 and the support is 96 × 111 for all the reconstructions.*
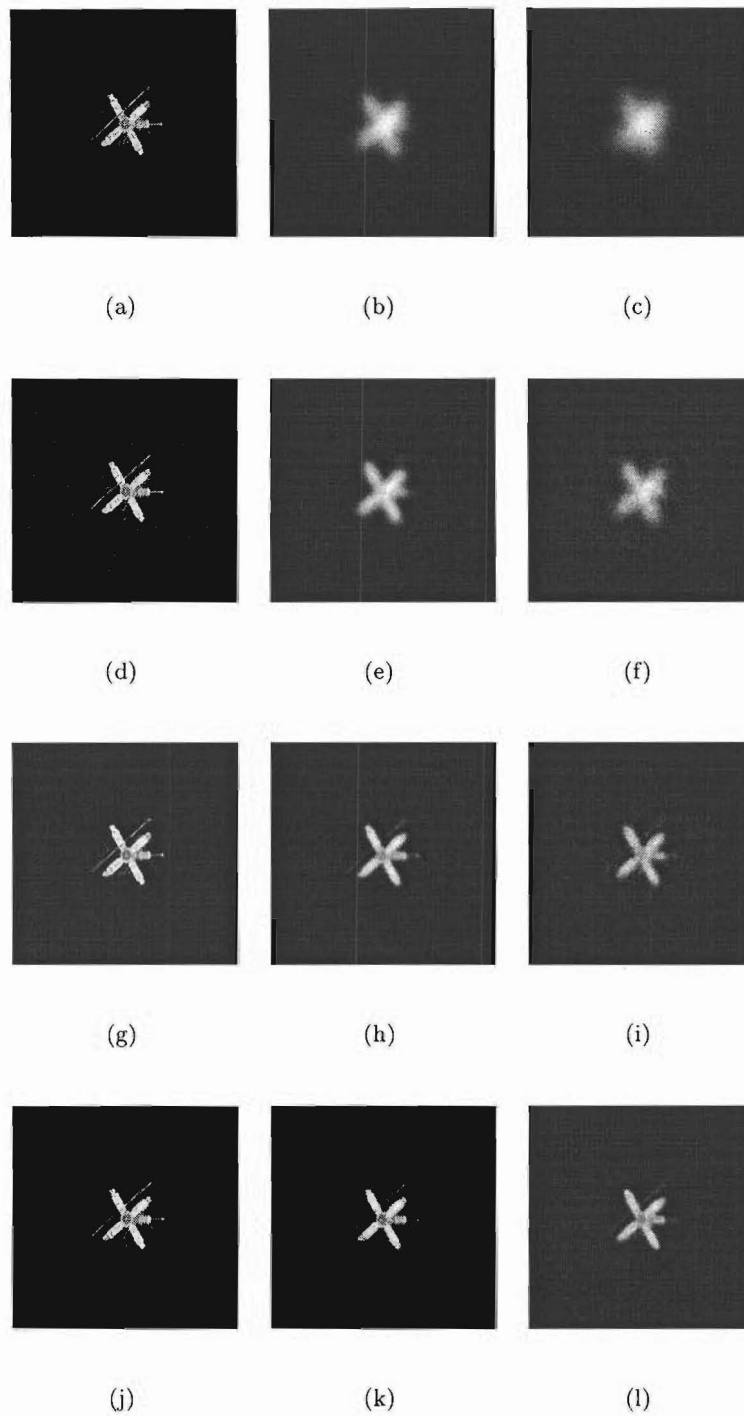
Figure 9.14: *Reconstructions of the satellite using the method of DWFS. A circular aperture is present at the telescope. The four rows correspond to $D/r_0 = 20, 15, 8$ and 5. The three columns correspond to the noiseless case and the cases when 5000 photons and 100 photons are present in each data frame.*
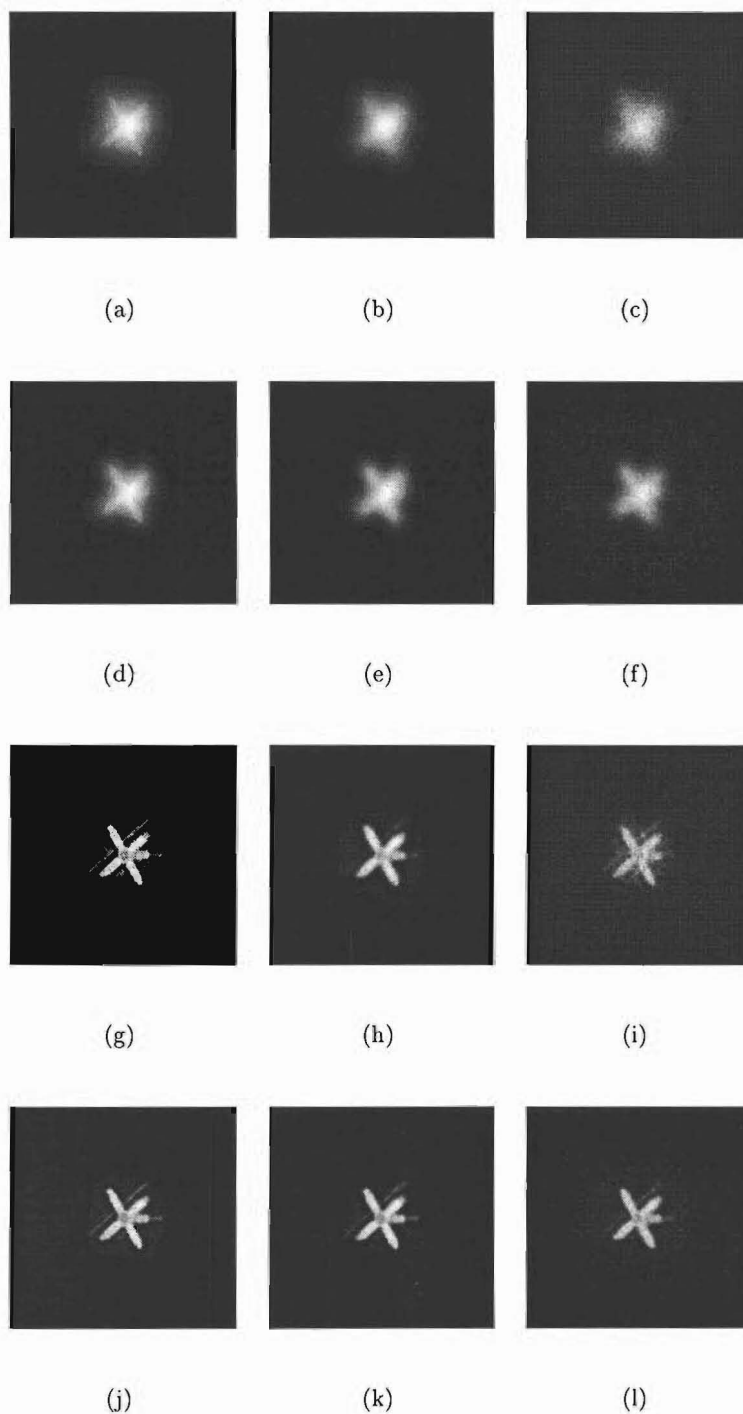
Figure 9.15: *Reconstructions of the satellite using the global parameter-isation method. A circular aperture is present at the telescope. The four rows correspond to $D/r_0 = 20, 15, 8$ and 5. The three columns correspond to the noiseless case and the cases when 5000 photons and 100 photons are present in each data frame.*

Figure 9.16: *Reconstructions of the satellite using the NASRIF method. A circular aperture is present at the telescope. The four rows correspond to $D/r_0 = 20, 15, 8$ and 5. The three columns correspond to the noiseless case and the cases when 5000 photons and 100 photons are present in each data frame.*
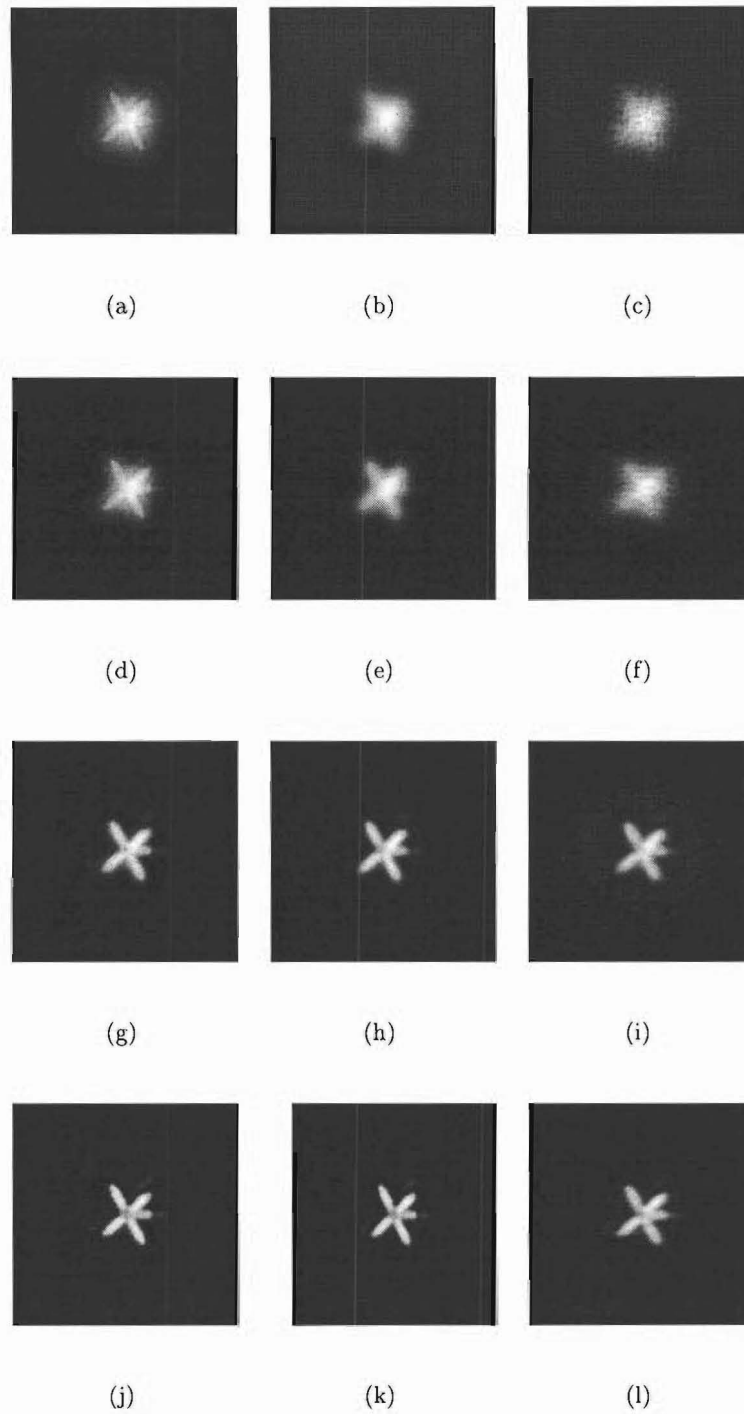
9.5.3 A comment on processing experimental data

The pixel-based parameterisation has been applied to experimental data of the binary star Castor. Preliminary results indicate that in order to obtain a solution of good quality, the quantity $m(x, y)$ should be initialised to the DWFS solution. It is also observed that this post-processing reduces the effect of noise on the reconstructions, and seems to improve the magnitude estimation of the object. However, further trials are required to fully characterise the performance of the algorithm on real data.

9.5.4 Other possible solutions

It should be pointed out that since the problem of solving DWFS by parameterisation is essentially a highpass filtering operation, it is also possible to apply other highpass filters (HPFs) to this problem. The simplest HPF is the Laplacian, denoted by:

$$\begin{pmatrix} 0 & -0.25 & 0 \\ -0.25 & 1 & -0.25 \\ 0 & -0.25 & 0 \end{pmatrix}.$$

This provides a very simple solution, but in general, the reconstruction will violate known constraints such as positivity and to a lesser extent the support. The reconstruction also suffers from ringing artifacts. Therefore, non-recursive *ad hoc* HPFs should be avoided if possible.

9.6 Summary

The techniques of DWFS and adaptive optics (AO) both utilise the Shack-Hartmann wavefront sensor for estimating the phase. However, the small errors in the reconstruction of the individual phases lead to a magnitude error in the object spectrum. Roggemann *et al.* [149] recognised this problem and proposed a technique to combine a set of reference star measurements together with the original measurements to produce the estimate of the object. However, the inaccuracy in the OTF in the numerator still prevents the accurate estimation of the magnitude of the object. Moreover, the use of a reference star suffers from problems such as anisoplanatism and reduces the amount of observation time.

In this chapter, a parameterised inverse filter was proposed as a solution to DWFS in order to remove the need for a reference star and potentially overcome the magnitude problem. This approach was considered because the object phase can be accurately reconstructed but not the magnitude. Since this problem arises from both the inaccuracies in the denominator and the numerator, it can be addressed collectively by the filter. Moreover, the filter's independence from observed data and atmospheric parameters means that they do not need to be measured or estimated.

Two parameterised filters were implemented and tested: the NASRIF algorithm, and a global parameterisation approach. Both of these approaches have produced comparable results to DWFS when applied to a compact object such as a binary star imaged through a circular aperture. When an annular aperture is present on the telescope, the resolving power is increased and almost diffraction-limited features are reconstructed by the methods of DWFS and global parameterisation. The comparable performance between the global parameterisation method and DWFS is significant. This is because the data for DWFS here has been generated under the most favourable conditions, and the algorithm's performance is likely to deteriorate when factors such as anisoplanatism are taken into account. This means that the global parameterisation method will probably be superior to DWFS when used with real data. The reliance on reference star measurements is also removed.

The DWFS method produced the best results at high levels of atmospheric turbulence for an extended object, but for lower levels of turbulence, the global parameterisation method produced comparable results without the reliance on a reference star. On the other hand, the NASRIF method failed to match the performance of either of the other two algorithms.

The parameterised filters can also be used to post-process adaptive optics (AO) imagery which suffers from the problems with using the Shack-Hartmann wavefront sensor. The parameterisation of the AO OTF draws on certain aspects of the existing myopic deconvolution techniques. For example, a mean PSF can be retrieved from the data accumulated by the AO control system during the acquisition [27, 51], which can then be used as an initial OTF parameterisation. Alternatively, a pixel-based parameterisation which does not require any knowledge of the PSF is also possible.

Chapter 10

Conclusions and future work

This chapter presents a conclusion to the research performed in this thesis. Possible improvement to the research, as well as ideas for future work are discussed.

10.1 Conclusions

The major focus of this thesis is to propose solutions for the inverse problems in image processing, with a particular application in astronomical imaging. Images acquired by any data-capturing instrument suffer from the distortions introduced by the optics, the sensor cells and the noise. The resolution of the captured images is therefore limited by these factors. A turbulent imaging medium also introduces degradations to the signals before they are captured. In astronomical imaging, it is the atmosphere which distorts the wavefronts of the light of the objects, severely limiting the resolution of the images captured by ground-based telescopes.

All the problems investigated in this thesis are inverse problems, and the key to the success of solving them is the correct modelling of the physical processes which give rise to the corresponding forward problems. However, because the actual system has an infinite amount of information, but the model used to approximate it can only have a finite amount of information, it is inevitable that information is lost. As a result, prior knowledge must be included in order to arrive at a meaningful solution.

The inverse problem of interpolation was investigated in Chapter 5. The ubiquity of the need for interpolation has resulted in many existing methods for this problem. However, most of these methods do not account for the sampling process of the instrument or take into account the statistics of the underlying object. The optimal Bayesian interpolator proposed in this chapter overcomes these problems. Existing interpolators were shown to be a subset of the framework of optimal interpolators under particular noise and sampling assumptions.

Since the effect of the atmosphere on the incoming wavefront of a star is predominantly a phase component, in order to reconstruct the object distribution, this phase distortion must be estimated or compensated for. When only speckle images are available, both the object and the phase must be estimated in a blind deconvolution framework. A review of the existing iterative blind deconvolution (IBD) algorithms in Chapter 6 indicated that strong priors such as phase parameterisation and statistical distributions provide a better constraint on the solution. Following this trend, the application of two new priors was investigated. Firstly, Kolmogorov statistics were used to constrain the atmospheric PSF. Simulations showed that this approach has the potential to out-perform Schulz's phase parameterisation algorithm, and a further investigation would be worthwhile. Secondly, the phase closure property used in the bispectrum technique was incorporated as a form of prior information in IBD. Preliminary results showed that this approach is not effective with a small ensemble, but enhanced performance was observed with an increased number of frames. Further simulations involving large ensembles should be performed to investigate this.

When Shack-Hartmann wavefront sensing data is available, it can be used to estimate the phase. The quality of this estimate depends intrinsically on the accuracy in locating the spots in the wavefront sensing data frames, and the popular method of calculating the centre of gravity is suboptimal and suffers from many problems. An improved model-fitting procedure based on blind deconvolution has been demonstrated on real data in Chapter 7.

In methods such as deconvolution from wavefront sensing (DWFS), the estimate of the phase is made solely from the wavefront sensing data, even though some information regarding the phase is present in the speckle images. Chapter 8 examined the use of both wavefront sensing data and speckle images for the estimation of the phase. The solution was posed as one of blind deconvolution, with the wavefront sensing data as a form of prior information.

Results on experimental data have demonstrated the success of this approach.

The hybrid method of DWFS suffers from a bias problem because the magnitude of the object spectrum cannot be accurately estimated. Roggemann *et al.* [149] proposed the use of reference star measurements to overcome this problem. However, this does not completely compensate for the inaccurate magnitude and relies on reference star measurements which suffer from other difficulties. In Chapter 9, a parameterised solution which has the potential to overcome these problems is proposed and investigated. It has been shown on simulated data to have comparable performance to DWFS, without the need of a reference star. Since adaptive optics (AO) imagery suffers from similar problems, the proposed method can also be used to post-process AO imagery.

10.2 Future work

In Chapter 6, the phase closure property employed in the bispectrum technique was used as prior information in blind deconvolution. Its performance was inferior to Schulz's phase parameterisation method when two frames were used for the reconstruction, but it did show improvement when the number of frames was increased to five. Therefore the first improvement to be made is to use a larger ensemble. Analysis into the minimum amount of information required for successful blind deconvolution should be performed. The contributing factors include the photon levels in addition to the ensemble size. Secondly, a comparison between the bispectrum technique and the blind deconvolution method is of considerable interest. This is because the former is a proven robust technique, while the latter is known to be severely ill-conditioned. The important question to be answered is whether incorporating the phase closure prior in a blind deconvolution framework provides all the necessary information to guarantee a solution.

The blind deconvolution algorithm in Chapter 7 used a maximum entropy (ME) prior for the PSF corresponding to the CCD sensitivity. Good results were produced when processing wavefront sensing data. Nevertheless, as discussed in Chapter 3, the ME prior controls the peak width and the baseline flatness, as well as enforcing positivity. It is therefore a strong constraint on the PSF. It is of interest to investigate the performance of other forms of priors in the blind deconvolution framework, such as centrosymmetricity and the smoothness constraint. A different data set (such as aerial photographs) should also be

used to validate the generality of the method.

It is known that phase reconstruction with the use of a Shack-Hartmann wavefront sensor (SH-WFS) suffers from the inability of the WFS to estimate certain Zernike modes. Since the phase information is available in both the WFS channel and the image channel, in Chapter 8 the joint application of both sets of data for phase estimation was investigated. Results using experimental data have proven the validity of this approach. The next step is to quantify whether the blind Zernike modes can be eliminated, as well as the information gain in terms of photon levels, number of frames, the SNR of the wavefront sensing data, and the degree of undersampling in the data. Since adaptive optics (AO) imagery involving the use of a Shack-Hartmann sensor suffers from similar problems, this technique should also be tested on AO imagery.

The expansion of a Kolmogorov phase screen using Zernike basis functions and Kahrunen-Loève basis functions for a Shack-Hartmann wavefront sensor has been compared [108]. It was shown that the latter is superior for phase reconstruction in terms of residual phase variance. Cannon [18] extended this work and derived a set of Kahrunen-Loève basis functions for an annular aperture. The next logical step is to investigate the use of this set of basis functions in a blind/myopic deconvolution framework, which should lead to a better phase reconstruction.

For very large telescopes (in the order of metres) at sites with good seeing, high-resolution imaging is possible. However, the use of a Shack-Hartmann (SH) sensor for imaging the reference star has problems with estimating some high frequency components due to the subdivision of the aperture. The pyramid wavefront sensor proposed in 1996 [136] is shown to perform better than a SH sensor in the vicinity of diffraction-limited imaging [34] as it is only limited by the diffraction effects of the telescope aperture. This calls for an investigation of myopic deconvolution using simultaneous pyramid sensing data and speckle images.

In a phase screen with Kolmogorov statistics, about 90% of the energy is concentrated in the second and third Zernike modes (tip and tilt), whose Zernike coefficients are large compared to the other modes. This large range of unknowns that must be estimated in the algorithm in Chapter 8 results in a slow convergence in a blind deconvolution framework. On the other hand, the convergence of the object seems to be achieved much faster. Other

algorithms which suffer from this problem of imbalance of convergence rate have benefitted from multiple iterations for the slower-converging unknown [12]. The application of this technique to the blind deconvolution problem in Chapter 8 should be investigated.

The most commonly used noise models are the Gaussian and the Poisson models. However, many physical situations cannot be described accurately by either of these models, but is a combination of the two processes. One example is a speckle image captured at a telescope. It is corrupted by both Poisson noise due to the low light level, as well as Gaussian noise due to the thermal activity of the CCD elements. As a result, the application of a combination of the Gaussian and Poisson noise models should be investigated. This is particularly relevant for blind deconvolution problems which are ill-conditioned and would benefit from an improved model of the system. There are two possible ways of achieving this: a weighted Gaussian model, which is a linearisation of the non-linear problem, and the exact model, where both the Gaussian and the Poisson PDFs are used [130]. A comparison of their performance, and an investigation of the tradeoff between the accuracy and the computational complexity of the models should be performed.

The presence of a central obscuration in a telescope is for practical purposes. However, it has inherently performed some highpass filtering on the incoming wavefront, resulting in a higher resolving power for the telescope compared to one with no central obscuration. It is of interest to investigate whether there are other practical configurations of the aperture that would result in a fundamental resolution improvement.

Appendix A

Optimisation approaches

In many unconstrained minimisation problem, an analytic solution is not possible due to the ill-posed nature and non-linearity of the problem. As a result, the problem must be solved iteratively by defining a cost function E to be minimised. To solve the problem, one starts at an initial estimate of the solution, and determines the direction and the step size in which the estimate has to move to minimise the cost function. A new estimate is obtained and this process is repeated until a minimum of the cost function is found. The important issues with iterative minimisation problems are global convergence, the speed at which it converges, and to a lesser extent the ease of implementation.

In this appendix, the methods of steepest descent, optimal steepest descent and conjugate gradient are presented. These methods are discussed in the context of two-dimensional image processing applications, although the theory is obviously extendable for other uses.

A.1 The method of steepest descent

The method of steepest descent is the simplest and most commonly used minimisation procedure. It is often used as a standard of reference with which other algorithms are compared.

Let the system be modelled as:

$$d(x,y) = f(x,y) \odot h(x,y) + n(x,y) \tag{A.1}$$

where $f(x,y)$ is the unknown object, $h(x,y)$ is the point spread function, $n(x,y)$ is the noise on the system and $d(x,y)$ the measured data. The associated cost function is $E(f(x,y))$. At the q^{th} iteration, $\hat{f}^q(x,y)$ is updated by:

$$\hat{f}^{q+1}(x,y) = \hat{f}^q(x,y) - \alpha^q \frac{\partial E(\hat{f}^q(x,y))}{\partial \hat{f}^q(x,y)} \tag{A.2}$$

where $\partial E(\hat{f}^q(x,y))/\partial \hat{f}^q(x,y)$ is the gradient. We denote it by $s^q(x,y)$. The step size α^q is a nonnegative scalar at the q^{th} iteration which minimises $E(\hat{f}^q(x,y) - \alpha^q s^q(x,y))$. The value of α^q can be found by standard line search techniques such as the bisection search, Golden section search or Brent's method [116, 134].

If the cost function to be minimised is quadratic it can be expressed as:

$$E(\mathbf{f}) = \frac{1}{2}\mathbf{f}^T\mathbf{Q}\mathbf{f} - \mathbf{f}^T\mathbf{b} \tag{A.3}$$

or

$$E(\mathbf{f}) = \frac{1}{2}\|\mathbf{H}\mathbf{f} - \mathbf{d}\|^2 \tag{A.4}$$

using lexicographical notation. Here, $\mathbf{Q} = \mathbf{H}^T\mathbf{H}$ and $\mathbf{b} = \mathbf{H}^T\mathbf{d}$. The minimisation of each of these error functions is equivalent to solving $\mathbf{Q}\mathbf{f} = \mathbf{b}$ for $\mathbf{f}$. In this case, the optimal step size for α^q can be calculated as [116]:

$$\alpha^q = \frac{(\mathbf{s}^q)^T\mathbf{s}^q}{(\mathbf{s}^q)^T\mathbf{Q}\mathbf{s}^q}. \tag{A.5}$$

This method is shown to converge globally [116]. However, a drawback is its slow convergence and as a result, other methods such as the conjugate gradient algorithm have been developed to overcome this problem.

A.2   The conjugate gradient algorithm

The conjugate gradient method is one of the most used minimisation routine because of its effectiveness in dealing with most objective functions, and its relative ease of implementation. The main difference between the conjugate gradient method and the method

of steepest descent is the selection of the direction of movement of the current estimate. At iteration q in the conjugate gradient method, the current negative gradient vector is evaluated, which is then added to a linear combination of the previous direction vectors to generate a new conjugate direction vector along which the current estimate should move. One major advantage of the conjugate gradient method is that for an n-dimensional function, convergence to the minimum is achieved in n or fewer steps.

The conjugate gradient algorithm is described here. Starting with $\mathbf{f}^0$ and $\mathbf{y}^0 = \mathbf{b} - \mathbf{Q}\mathbf{f}^0 = -\mathbf{s}_0$, the algorithm requires the following steps:

$$
\begin{aligned}
\mathbf{f}^{q+1} &= \mathbf{f}^q + \alpha^q \mathbf{y}^q \\
\alpha^q &= -\frac{(\mathbf{s}^q)^T \mathbf{y}^q}{(\mathbf{y}^q)^T \mathbf{Q} \mathbf{y}^q} \\
\mathbf{y}^{q+1} &= -\mathbf{s}^{q+1} + \beta^q \mathbf{y}^q \\
\beta^q &= \frac{(\mathbf{s}^{q+1})^T \mathbf{Q} \mathbf{y}^q}{(\mathbf{y}^q)^T \mathbf{Q} \mathbf{y}^q}
\end{aligned}
\tag{A.6}
$$

where $\mathbf{s}^q = \mathbf{Q}\mathbf{f}^q - \mathbf{b}$. Here $\mathbf{s}^q$ is the gradient, $\mathbf{y}^q$ is the search direction and β^q is the update of the search direction.

The conjugate gradient algorithm can be extended to solve non-quadratic problems by the use of line search methods. In this case, $\mathbf{Q}$ does not need to be calculated. One of these methods is the Fletcher-Reeves method, which uses a standard line search procedure (such as the golden section search) to calculate α^q while the value of β^q is obtained using

$$
\beta^q = \frac{(\mathbf{s}^{q+1})^T \mathbf{s}^{q+1}}{(\mathbf{s}^q)^T \mathbf{s}^q}.
\tag{A.7}
$$

Alternatively, the Polak-Ribiere method can be used to find β^q where

$$
\beta^q = \frac{(\mathbf{s}^{q+1} - \mathbf{s}^q)^T \mathbf{s}^{q+1}}{(\mathbf{s}^q)^T \mathbf{s}^q}.
\tag{A.8}
$$

Experimental results seem to favour the latter method.

A.3   Partial conjugate gradient and optimal steepest descent

There are situations in which a full conjugate gradient procedure is not suitable, and a partial conjugate gradient method is preferrable. This means that the gradient is reset

to be a function of only the current variables every n iterations. The need for this can arise from a multimodal cost function stagnating in a local minimum. If $n = 1$, then the conjugate gradient method is reduced to an optimal steepest descent method.

Bibliography

[1] J.G. Ables, "Maximum entropy spectral analysis", *Astron. Astrophy. Suppl. Series*, Vol. 15, No. 3, pp 383–393, June 1974.

[2] M.S. Alam, J.G. Bognar, S. Cain and B.J. Yasuda, "Fast registration and reconstruction of aliased low-resolution frames by use of a modified maximum-likelihood approach", *Appl. Opt.*, Vol. 37, No. 8, pp 1319–1328, March 1998.

[3] S. Altarac, M. Tallon, E. Thiébaut and R. Foy, "Speckle interferometry. Data acquisition and control for the SPID instrument", *IEEE Trans. Nucl. Sci.*, Vol. 45, No. 4, pp 2057–2061, August 1998.

[4] H.C. Andrews and B.R. Hunt, *Digital Image Restoration*, Prentice-Hall, New Jersey, 1977.

[5] H. Anton, *Calculus*, John Wiley & Sons, Inc., New York, 4^{th} edition, 1992.

[6] G.R. Ayers and J.C. Dainty, "Iterative blind deconvolution method and its applications", *Opt. Lett.*, Vol. 13, No. 7, pp 547–549, July 1988.

[7] G.R. Ayers, M.J. Northcott and J.C. Dainty, "Knox-Thompson and triple correlation imaging through atmospheric turbulence", *J. Opt. Soc. Am. A*, Vol. 5, No. 7, pp 963–985, 1988.

[8] H. Baher, *Analog and Digital Signal Processing*, John Wiley and Sons, Chichester, 1990.

[9] R.H.T. Bates and M.J. McDonnell, *Image Restoration and Reconstruction*, Oxford University Press, New York, 1986.

[10] D.S.C. Biggs and M. Andrews, "Conjugate gradient acceleration of maximum-likelihood image restoration", *Elec. Letters*, Vol. 31, No. 23, pp 1985–1986, November 1995.

[11] D.S.C. Biggs and M. Andrews, "Acceleration of iterative image restoration algorithms", *Appl. Opt.*, Vol. 36, No. 8, pp 1766–1775, March 1997.

[12] D.S.C. Biggs and M. Andrews, "Asymmetric iterative blind deconvolution of multiframe images", *Proc. of SPIE*, Vol. 3641, pp 328–338, 1998.

[13] M. Born and E. Wolf, *Principles of Optics*, University Press, Cambridge, Cambridge, U.K., 7^{th} edition, 1999.

[14] R.N. Bracewell, *The Fourier transforms and its applications*, Mc-Graw Hill Book Company, New York, 1965.

[15] T. Bretschneider, P.J. Bones and S. McNeill, "Resolution enhancement using multispectral remotely sensed imagery", D. Pairman and H. North, (editor), *Proc. of IVCNZ'99*, pp 109–114, 1999.

[16] S.C. Cain and M.M. Hayat, "Exploiting the temporal statistics of atmospheric tilt for improved short exposure imaging", *Signal Recovery and Synthesis, OSA Technical Digest*, Vol. 1, pp 65–67, 2001.

[17] R.C. Cannon, "Global wave-front reconstruction using Shack-Hartmann sensors", *J. Opt. Soc. Am. A*, Vol. 12, No. 9, pp 2031–2039, September 1995.

[18] R.C. Cannon, "Optimal bases for wave-front simulation and reconstruction on annular apertures", *J. Opt. Soc. Am. A*, Vol. 13, No. 4, pp 862–867, April 1996.

[19] K.R. Castleman, *Digital Image Processing*, Prentice-Hall, New Jersey, 1996.

[20] Centre National d'Etudes Spatiales (CNES), France, SPOT satellite imagery.

[21] T.F. Chan and C.-K. Wong, "Total Variation Blind Deconvolution", *IEEE Trans. Im. Proc.*, Vol. 7, No. 3, pp 370–375, March 1998.

[22] Z. Chen, M.A. Karim and M.M. Hayat, "Elimination of higher order aliasings by multiple interlaced sampling", *Opt. Eng.*, Vol. 38, No. 5, pp 879–885, May 1999.

[23] G. Chinn and S.-C. Huang, "A general class of preconditioners for statistical iterative reconstruction of emission computed tomography", *IEEE Trans. Med. Im.*, Vol. 16, No. 1, pp 1–10, February 1997.

[24] J.C. Christou, D. Bonaccini, N. Ageorges and F. Marchis, "Myopic deconvolution of adaptive optics image", *Messenger*, Vol. 97, pp 14–22, 1999.

[25] I. Clark, *Practical Geostatistics*, Applied Science Publishers Ltd, London, 1979.

[26] N.H. Clinthorne, T.-S. Pan, P.-C. Chiao, W.L. Rogers and J.A. Stamos, "Preconditioning methods for improved convergence rates in iterative reconstructions", *IEEE Trans. Med. Im.*, Vol. 12, No. 1, pp 78–83, March 1993.

[27] J.-M. Conan, L.M. Mugnier, T. Fusco, V.Michau and G. Rousset, "Myopic deconvolution of adaptive optics images by use of object and point-spread function power spectra", *Appl. Opt.*, Vol. 37, No. 21, pp 4614–4622, July 1998.

[28] S.R. Curtis, A.V. Oppenheim and J.S. Lim, "Signal reconstruction from Fourier transform signal information", *IEEE Trans. Acoust., Speech and Sig. Proc.*, Vol. 33, No. 3, pp 643–657, June 1985.

[29] J.C. Dainty and A.H. Greenaway, "Estimation of spatial power spectra in speckle imaging", *J. Opt. Soc. Am.*, Vol. 69, pp 786–790, 1979.

[30] G. Demoment, "Image reconstruction and restoration: overview of common estimation structures and problems", *IEEE Trans. Acoust., Speech and Sig. Proc.*, Vol. 37, No. 12, pp 2024–2036, December 1989.

[31] A.P. Dempster, N.M. Laird and D.B. Rubin, "Maximum likelihood from incomplete data", *J. of the Royal Stat. Soc.*, Vol. 39, No. 1, pp 1–38, 1977.

[32] R. Ditteon, *Modern Geometrical Optics*, John Wiley & Sons, Inc., New York, 1998.

[33] P. Duhamel and M. Vetterli, "Fast Fourier transforms: A tutorial review and a state of the art", *Sig. Proc.*, Vol. 19, No. 4, pp 259–299, April 1990.

[34] S. Esposito and A. Riccardi, "Pyramid wavefront sensor behaviour in partial correction adaptive optics systems", *Astron. Astrophy.*, Vol. 369, pp L9–L12, April 2001.

[35] O.J. Farrell and B. Ross, *Solved Problems: Gamma and Beta functions, Legendre polynomials, Bessel functions*, The MacMillan Company, New York, 1963.

[36] D.A. Fish, A.M. Brinicombe, E.R. Pike and J.G. Walker, "Blind deconvolution by means of the Richardson-Lucy algorithm", *J. Opt. Soc. Am. A*, Vol. 12, No. 1, pp 58–65, January 1995.

[37] R.A. Fisher, "On an absolute criterion for fitting frequency curves", *The Messenger of Mathematics*, Vol. 41, pp 155–160, 1912.

[38] R. Foy, *Instrumentation for ground-based optical astronomy, present and future*, "The photon counting camera CP40", pp 589–592, Springer-Verlag, New York, 1998.

[39] R. Foy and A. Labeyrie, "Feasibility of adaptive telescopes with laser probes", *Astron. Astrophy.*, Vol. 152, No. 2, pp L29–L31, November 1985.

[40] J.B. Fraleigh and R.A. Beauregard, *Linear Algebra*, Addison-Wesley Publishing Company, California, 2^{nd} edition, 1990.

[41] J.N. Franklin, "Well posed stochastic extensions of ill posed linear problems", *J. Math. Anal. Applic.*, Vol. 31, pp 682–716, 1970.

[42] D.L. Fried, "Statistics of a geometric representation of wavefront distortion", *J. Opt. Soc. Am. A*, Vol. 55, No. 11, pp 1427–1435, November 1965.

[43] D.L. Fried, "Optical resolution through a randomly inhomogeneous medium for very long and very short exposures", *J. Opt. Soc. Am. A*, Vol. 56, No. 10, pp 1372–1379, October 1966.

[44] D.L. Fried, "Analysis of the CLEAN algorithm and implications for superresolution", *J. Opt. Soc. Am. A*, Vol. 31, No. 5, pp 853–860, May 1970.

[45] D.L. Fried, "Anisoplanatism in adaptive optics", *J. Opt. Soc. Am. A*, Vol. 72, No. 1, pp 52–61, January 1982.

[46] D.L. Fried, "Post detection wavefront compensation", *SPIE Proceedings on Digital image recovery and synthesis*, pp 127–133, 1987.

[47] B.R. Frieden, "Restoring with maximum likelihood and maximum entropy", *J. Opt. Soc. Am. A*, Vol. 62, No. 4, pp 511–518, April 1972.

[48] M. Frigo and S.G. Johnson, "FFTW: An adaptive software architecture for the FFT", *Int. Conf. Acoust., Speech and Sig. Proc.*, Vol. 3, pp 1381–1384, 1998.

[49] M. Frigo and S.G. Johnson, FFTW, http://www.fftw.org, Feb 2000.

[50] J. Fryer and K. McIntosh, "Enhancement of image resolution in digital photogrammetry", *Photogram. Eng. and Remote. Sens.*, Vol. 67, No. 6, pp 741–749, June 2001.

[51] T. Fusco, J.P. Véran, J.-M. Conan and L.M. Mugnier, "Myopic deconvolution method for adaptive optics images of stellar fields", *Astron. Astrophy. Suppl. Series*, Vol. 134, No. 1, pp 193–200, January 1999.

[52] J.P. Garcia and E.L. Dereniak, "Mixed-expectation image-reconstruction technique", *Appl. Opt.*, Vol. 38, No. 17, pp 3745–3748, June 1999.

[53] S. Geman and D. Geman, "Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images", *IEEE Trans. Pat. Anal. and Mach. Intel.*, Vol. 6, No. 6, pp 721–741, November 1984.

[54] R.W. Gerchberg, "Super-resolution through error energy reduction", *Opt. Acta*, Vol. 21, No. 9, pp 709–720, September 1974.

[55] D.R. Gerwe, D.J. Lee and J.D. Barchers, "Supersampling multiframe blind deconvolution resolution enhancement of adaptive optics compensated imagery of LEO satellites", *Proc. SPIE Imaging technology and telescopes conference*, pp 187–205, 2000.

[56] D.R. Gerwe and M.A. Plonus, "Superresolved image reconstruction of images taken through the turbulent atmosphere", *J. Opt. Soc. Am. A*, Vol. 15, No. 10, pp 2620–2628, October 1998.

[57] J.C. Gillette, T.M. Stadmiller and R.C. Hardie, "Reduction of aliasing in staring infrared imagers utilizing subpixel techniques", *Opt. Eng.*, Vol. 34, No. 11, pp 3130–3137, November 1995.

[58] A. Glindemann and J.C. Dainty, "Object fitting to the bispectral phase by using least squares", *J. Opt. Soc. Am. A*, Vol. 10, No. 5, pp 1056–1063, May 1993.

[59] A. Glindemann, R.G. Lane and J.C. Dainty, "Estimation of binary star parameters by model fitting the bispectrum phase", *J. Opt. Soc. Am. A*, Vol. 9, No. 4, pp 543–548, April 1992.

[60] G.H. Golub and C.F. van Loan, *Matrix computations*, The Johns Hopkins University Press, Baltimore, 2^{nd} edition, 1993.

[61] J.W. Goodman, *Introduction to Fourier optics*, McGraw-Hill, Inc., San Francisco, 1968.

[62] J.W. Goodman, *Statistical optics*, John Wiley & Sons, Inc., New York, 1985.

[63] P. Goovaerts, *Geostatistics for natural resources evaluation*, Oxford University Press, New York, 1997.

[64] D. Granrath and J. Lersch, "Fusion of images on affine sampling grids", *J. Opt. Soc. Am. A*, Vol. 15, No. 4, pp 791–801, April 1998.

[65] J.J. Green and B.R. Hunt, "Improved restoration of space object imagery", *J. Opt. Soc. Am. A*, Vol. 16, No. 12, pp 2859–2865, December 1999.

[66] S.F. Gull and G.J. Daniell, "Image reconstruction from incomplete and noisy data", *Nature*, Vol. 272, pp 686, April 1978.

[67] S.F. Gull and G.J. Daniell, *Image formation from coherence functions in astronomy*, p 219, D. Reidel, Dordrecht, 1979.

[68] J. Hadamard, "Sur les problèmes aux dérivées partialles et leur signification physique", *Bull. Princeton Univ.*, Vol. 13, 1902.

[69] J. Hadamard, *Le problème de Cauchy et les équations aux dérivées partielles linéaires hyperboliques*, Hermann, Paris, 1932.

[70] A. Hald, "On the history of maximum likelihood in relation to inverse probability and least squares", *Statistical Science*, Vol. 14, No. 2, pp 214–222, 1999.

[71] E.L. Hall, *Computer image processing and recognition*, Academic Press, Inc., New York, 1979.

[72] C.A. Haniff, "Least-squares Fourier phase estimation from the modulo 2π bispectrum phase", *J. Opt. Soc. Am. A*, Vol. 8, No. 1, pp 134–140, January 1991.

[73] R. C. Hardie, K. J. Barnard and E. Armstrong, "Joint MAP registration and high-resolution image estimation using a squence of undersampled images", *IEEE Trans. Im. Proc.*, Vol. 6, No. 12, pp 1621–1633, December 1997.

[74] R.C. Hardie, K.J. Barnard, J.G. Bognar and E.A. Watson, "High-resolution image reconstruction from a sequence of rotated and translated frames and its application to an infrared imaging system", *Opt. Eng.*, Vol. 37, No. 1, pp 247–260, Jan 1998.

[75] R.C. Hardie, M.M. Hayat, E. Armstrong and B. Yasuda, "Scene-based nonuniformity correction with video sequences and registration", *Appl. Opt.*, Vol. 39, No. 8, pp 1241–1250, March 2000.

[76] C.M. Harding, R.A. Johnston and R.G. Lane, "Fast simulation of a Kolmogorov phase screen", *Appl. Opt.*, Vol. 38, No. 11, pp 2161–2170, April 1999.

[77] J.W. Hardy, *Adaptive Optics for Astronomical Telescopes*, Oxford University Press, Inc., New York, 1998.

[78] S. Haykin, *Communication Systems*, John Wiley & Sons, Inc., New York, 3^{rd} edition, 1994.

[79] J.A. Högbom, "Aperture synthesis with a non-regular distribution of interferometer baselines", *Astron. Astrophy. Suppl. Series*, Vol. 15, No. 3, pp 417–426, June 1974.

[80] T.J. Holmes and Y.-H. Liu, "Acceleration of maximum-likelihood image restoration for fluorescence microscopy and other noncoherent imagery", *J. Opt. Soc. Am. A*, Vol. 8, No. 6, pp 893–907, June 1991.

[81] H.S. Hou and H.C. Andrews, "Cubic splines for image interpolation and digital filtering", *IEEE Trans. Acoust., Speech and Sig. Proc.*, Vol. 26, No. 6, pp 508–517, December 1978.

[82] T.S. Huang and R.Y. Tsay, *Advances in computing vision and image processing*, "Multiple frame image restoration and registration", pp 317–339, CT:JAI, Greenwich, 1984.

[83] F.O. Huck, N. Nhalyo and S.K. Park, "Aliasing and blurring in 2-D sampled imagery", *Appl. Opt.*, Vol. 19, No. 13, pp 2174–2181, 1980.

[84] B.R. Hunt, "Deconvolution of linear systems by constrained regression and its relationship to the Wiener theory", *IEEE Trans. Autom. Control*, Vol. 17, No. 5, pp 703–705, October 1972.

[85] B.R. Hunt, "Super-Resolution of Images: Algorithms, Principles, Performance", *Int. J. of Imag. Sys. and Tech.*, Vol. 6, pp 297–304, 1995.

[86] B.R. Hunt and P. Sementilli, *Description of a Poisson imagery super-resolution algorithm*, Vol. 25, Astron. Soc. of the Pac., San Francisco, 1992.

[87] E.C. Ifeachor and B.W. Jervis, *Digital Signal Processing: A Practical Approach*, Addison-Wesley, Essex, England, 1993.

[88] E.H. Isaaks and R.M. Srivastava, *An introduction to applied geostatistics*, Oxford University Press, New York, 1989.

[89] E.T. Jaynes, "Prior probabilities", *IEEE Trans. Sys. Sci. and Cyber.*, Vol. 4, No. 3, pp 227–241, September 1968.

[90] S. Jefferies and J.C. Christou, "Restoration of astronomical images by iterative blind deconvolution", *The Astrophy. Journal*, Vol. 415, No. 2, pp 862–874, October 1993.

[91] F.A. Jenkins and H.E. White, *Fundamentals of Optics*, McGraw-Hill Kogakusha, Ltd., Tokyo, 3 edition, 1957.

[92] R.W. Johnson and G.J.M. Aitken, "Image truncation in stellar speckle imaging", *J. Opt. Soc. Am. A*, Vol. 9, No. 11, pp 1955–1963, November 1992.

[93] R.A. Johnston, *Inverse problems in astronomical imaging*, PhD thesis, University of Canterbury, Christchurch, New Zealand, November 2000.

[94] R.A. Johnston and R.G. Lane, "An improved method for deconvolving a positive image", *Opt. Comm.*, Vol. 181, pp 267–278, July 2000.

[95] A.K. Katsaggelos, *Digital image restoration*, pp 1–20, Springer-Verlag, Berlin, 1991.

[96] L. Kaufman, "Implementing and accelerating the EM algorithm for positron emission tomography", *IEEE Trans. Med. Im.*, Vol. 6, No. 1, pp 37–51, March 1987.

[97] S.M. Kay, *Fundamentals of statistical signal processing: estimation theory*, Prentice-Hall, New Jersey, 1993.

[98] R.G. Keys, "Cubic convolution interpolation for digital image processing", *IEEE Trans. Acoust., Speech and Sig. Proc.*, Vol. 29, No. 6, pp 1153–1160, December 1981.

[99] K.T. Knox and B.J. Thompson, "Recovery of images from atmospherically degraded-short exposure photographs", *Astrophy. J.*, Vol. 193, pp L45–L48, 1974.

[100] D. Kouznetsov, V.V. Voitsekhovich and R.Ortega-Martinez, "Simulations of turbulence-induced phase and log-amplitude distortions", *Appl. Opt.*, Vol. 36, pp 464–469, 1997.

[101] E. Kreyszig, *Advanced engineering mathematics*, John Wiley & Sons, Inc., New York, 7^{th} edition, 1993.

[102] D. Kundur and D. Hatzinakos, "A novel blind deconvolution scheme for image restoration using recursive filtering", *IEEE Trans. Sig. Proc.*, Vol. 46, No. 2, pp 375–390, February 1998.

[103] A. Labeyrie, "Attainment of diffraction limited resolution in large telescopes by Fourier analysing speckle patterns in star images", *Astron. Astrophy.*, Vol. 6, pp 85–87, 1970.

[104] R.L. Lagendijk and J. Biemond, *Iterative identification and restoration of images*, Kluwer Academic Publishers, Boston, 1991.

[105] R.G. Lane, "Blind deconvolution of speckle images", *J. Opt. Soc. Am. A*, Vol. 9, No. 9, pp 1508–1514, September 1992.

[106] R.G. Lane, "Methods for maximum-likelihood deconvolution", *J. Opt. Soc. Am. A*, Vol. 13, No. 10, pp 1992–1998, October 1996.

[107] R.G. Lane and R.H.T. Bates, "Automatic multidimensional deconvolution", *J. Opt. Soc. Am. A*, Vol. 4, No. 1, pp 180–188, January 1987.

[108] R.G. Lane and M. Tallon, "Wave-front reconstruction using a Shack-Hartmann sensor", *J. Opt. Soc. Am. A*, Vol. 31, No. 32, pp 6902–6908, November 1992.

[109] T.R. Lauer, "Combining undersampled dithered images", *Publ. Astron. Soc. Pac.*, Vol. 111, pp 227–237, February 1999.

[110] N.F. Law and R.G. Lane, "Blind deconvolution using least squares minimisation", *Opt. Comm.*, Vol. 128, pp 341–352, July 1996.

[111] W-Y.V. Leung, R.M. Clare and R.G. Lane, "Blind deconvolution of speckle images constrained by wavefront sensing data", *To be published in Proc. SPIE Image Reconstruction from Incomplete Data II* , July 2002.

[112] J. Liang, D.R. Williams and D.T. Miller, "Supernormal vision and high-resolution retinal imaging through adaptive optics", *J. Opt. Soc. Am. A*, Vol. 14, pp 2884–2892, 1997.

[113] A.W. Lohmann, G. Weigelt and B. Wirnitzer, "Speckle masking in astronomy: triple correlation theory and applications", *Appl. Opt.*, Vol. 22, No. 24, pp 4028–4037, December 1983.

[114] L.B. Lucy, "An iterative technique for the rectification of observed distribution", *The Astrophy. Journal*, Vol. 79, No. 6, pp 745–754, June 1974.

[115] D.G. Luenberger, *Optimization by Vector Space Methods*, Wiley, New York, 1969.

[116] D.G. Luenberger, *Linear and Nonlinear Programming*, Addison-Wesley Publishing Company, Inc., Massachusetts, 1984.

[117] Matlab Version 6.0.0.88 Release 12, The MathWorks, Inc., September 2000.

[118] C.L. Matson, "Fourier spectrum extrapolation and enhancement using support constraints", *IEEE Trans. Sig. Proc.*, Vol. 42, No. 1, pp 156–163, January 1994.

[119] C.L. Matson, "Resolution, linear filtering, and positivity", *J. Opt. Soc. Am. A*, Vol. 15, No. 1, pp 33–41, January 1998.

[120] B.C. McCallum, "Blind deconvolution by simulated annealing", *Opt. Comm.*, Vol. 75, No. 2, pp 101–105, February 1990.

[121] N. Metropolis, A. Rosenbluth, M. Rosenbluth, A. Teller and E. Teller, "Equation of state calculations by fast computing machines", *J. Chem. Phys.*, Vol. 21, No. 6, pp 1087, June 1953.

[122] D.T. Miller, "Retinal imaging and vision at the frontiers of adaptive optics", *Phys. Today*, Vol. , pp 31–36, January 2000.

[123] N. Miura, M. Inoue and N. Baba, "Blind deconvolution constrained by Knox-Thompson spectral phases for photon-limited stellar images", *Opt. Comm.*, Vol. 178, pp 297–308, May 2000.

[124] N. Miura, N. Sawai and N. Baba, "Stellar image reconstruction from phase difference", *Opt. Eng.*, Vol. 118, pp 466–472, August 1995.

[125] N. Miura, N. Sawai and N. Baba, "Stellar image reconstruction by minimizing errors for bispectral phases", *Opt. Comm.*, Vol. 138, pp 275–278, June 1997.

[126] L.M. Mugnier, C. Robert, J.-M. Conan, V. Michau and S. Salem, "Myopic deconvolution from wave-front sensing", *J. Opt. Soc. Am. A*, Vol. 18, No. 4, pp 862–872, April 2001.

[127] I.J. Nagrath and M. Gopal, *Control Systems Engineering*, New Age International (P) Limited, New Delhi, 2^{nd} edition, 1981.

[128] M.K. Ng, R.J Plemmons and S. Qiao, "Regularization of RIF blind image deconvolution", *IEEE Trans. Im. Proc.*, Vol. 9, No. 6, pp 1130–1134, June 2000.

[129] R. Nityananda and R. Narayan, "Maximum entropy image reconstruction- a practical non-information theoretic approach", *J. Astrophy. and Astron.*, Vol. 3, pp 419–450, 1982.

[130] J. Núñez and J. Llacer, "A general Bayesian image reconstruction algorithm with entropy prior. Preliminary application to HST data", *Publ. Astron. Soc. Pac.*, Vol. 105, pp 1192–1208, October 1993.

[131] R.J. Noll, "Zernike polynomials and atmospheric turbulence", *J. Opt. Soc. Am. A*, Vol. 66, No. 3, pp 207–211, March 1976.

[132] A. Papoulis, *The Fourier integral and its applications*, Mc-Graw Hill Book Company, New York, 1962.

[133] W.K. Pratt, "Generalised Wiener filter computation techniques", *IEEE Trans. Compt.*, Vol. C-21, pp 636–641, 1972.

[134] W.H. Press, *Numerical recipes in C: the art of scientific computing*, Cambridge University Press, Cambridge, 2^{nd} edition, 1992.

[135] J. Primot, G. Rousset and J.C. 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, September 1990.

[136] R. Ragazzoni, "Pupil plane wavefront sensing with oscillating prism", *J. Mod. Opt.*, Vol. 43, No. 2, pp 289–293, 1996.

[137] S. Ramo, J.R. Whinnery and T. van Duzer, *Field and waves in communication electronics*, John Wiley & Sons, Inc., New York, 3^{rd} edition, 1994.

[138] C.H. Reinsch, "Smoothing by spline functions", *Num. Math.*, Vol. 10, pp 177–183, 1967.

[139] W.H. Richardson, "Bayesian-Based Method of Image Restoration", *J. Opt. Soc. Am. A*, Vol. 62, No. 1, pp 55–59, January 1972.

[140] F. Roddier, *Progress in Optics XIX*, "The effects of atmospheric turbulence in optical astronomy", pp 281–376, North-Holland Publishing Company, Amsterdam, 1981.

[141] F. Roddier, "Triple correlation as a phase closure technique", *Opt. Eng.*, Vol. 60, pp 145–148, 1986.

[142] F. Roddier, "Curvature sensing and compensation: a new concept in adaptive optics", *Appl. Opt.*, Vol. 27, pp 1223–1225, 1988.

[143] F. Roddier, *Adaptive Optics in Astronomy*, "The design of an adaptive optics system: Theoretical aspects", pp 25–56, Cambridge University Press, United Kingdom, 1999.

[144] M.C. Roggemann, "Optical performance of fully and partially compensated adaptive optics systems using least-squares and minimum variance phase reconstruction", *Comp. and Elec. Eng.*, Vol. 18, No. 6, pp 451–466, 1992.

[145] M.C. Roggemann and T.J. Schulz, "Algorithm to increase the largest aberration that can be reconstructed from Hartmann sensor measurements", *Appl. Opt.*, Vol. 37, No. 20, pp 4321–4329, July 1998.

[146] M.C. Roggemann, T.J. Schulz, C.W. Ngai and J.T. Kraft, "Joint processing of Hartmann sensor and conventional image measurements to estimate large aberrations: theory and experimental results", *J. Opt. Soc. Am. A*, Vol. 38, No. 11, pp 2249–2255, April 1999.

[147] M.C. Roggemann and B. Welsh, *Imaging through turbulence*, CRC Press, Florida, 1996.

[148] M.C. Roggemann and B.M. Welsh, "Signal-to-noise ratio for astronomical imaging by deconvolution from wave-front sensing", *Appl. Opt.*, Vol. 33, No. 23, pp 5400–5414, August 1994.

[149] M.C. Roggemann, B.M. Welsh and J. Devey, "Biased estimators and object-spectrum estimation in the method of deconvolution from wavefront sensing", *Appl. Opt.*, Vol. 33, No. 24, pp 5754–5763, August 1994.

[150] A. Rosenfeld and A.C. Kak, *Digital Picture Processing, Vol. 1*, Academic Press, Inc., New York, 2^{nd} edition, 1982.

[151] S.A. Sallberg, B.M. Welsh and M.C. Roggemann, "Maximum *a posteriori* estimation of wave-front slopes using a Shack-Hartmann wave-front sensor", *J. Opt. Soc. Am. A*, Vol. 14, No. 6, pp 1347–1354, June 1997.

[152] I.J. Schoenberg, "Spline functions and the problem of graduation", *Proceedings of the National Academy of Science*, Vol. 52, pp 947–950, 1964.

[153] J.M. Schuler and D.A. Scribner, "Increasing spatial resolution through temporal super-sampling of digital video", *Opt. Eng.*, Vol. 38, No. 5, pp 801–805, May 1999.

[154] T.J. Schulz, "Multiframe blind deconvolution of astronomical images", *J. Opt. Soc. Am. A*, Vol. 10, No. 5, pp 1064–1073, May 1993.

[155] P.J. Sementilli, B.R. Hunt and M.S. Nadar, "Analysis of the limit to superresolution in incoherent imaging", *J. Opt. Soc. Am. A*, Vol. 10, No. 11, pp 2265–2276, November 1993.

[156] L.A. Shepp and Y. Vardi, "Maximum likelihood reconstruction for emission tomography", *IEEE Trans. Med. Im.*, Vol. 1, No. 2, pp 113–122, October 1982.

[157] D.G. Sheppard, B.R. Hunt and M.W. Marcellin, "Iterative multiframe superresolution algorithms for atmospheric-turbulence-degraded imagery", *J. Opt. Soc. Am. A*, Vol. 15, No. 4, pp 978–992, April 1998.

[158] D.L. Snyder and M.I. Miller, *Random point processes in time and space*, Springer-Verlag New York, Inc., New York, 2^{nd} edition, 1991.

[159] A. Sommerfeld, *Optics, Lectures on Theoretical Physics, Vol.IV*, Academic Press, Inc., New York, 1954.

[160] H. Stark, (editor), *Image Recovery: Theory and Application*, Academic Press, Inc., Orlando, Florida, 1987.

[161] G.C. Steward, *The Symmetrical Optical System*, Cambridge University Press, Cambridge, 1928.

[162] G. Strang, *Linear Algebra and its Applications*, Harcourt Brace Jovanovich College Publishers, Florida, 3^{rd} edition, 1986.

[163] W.A. Strauss, *Partial Differential Equations: An Introduction*, John Wiley & Sons, Inc., New York, 1992.

[164] C.R. Subrahmanya, *Image formation from coherence functions in astronomy*, p 261, D. Reidel, Dordrecht, 1979.

[165] C.R. Subrahmanya, *Bulletin of the astronomical society of India*, Vol. 8, pp 5, 1980.

[166] V.I. Tatarskii, *Wave Propagation in a Turbulent Medium*, McGraw-Hill, New York, 1961.

[167] E. Thiébaut, "Speckle interferometry with a photon-counting detector", *Astron. Astrophy.*, Vol. 284, pp 340–348, 1994.

[168] E. Thiébaut and J.-M. Conan, "Strict *a priori* constraints for maximum-likelihood blind deconvolution", *J. Opt. Soc. Am. A*, Vol. 12, No. 3, pp 485–492, March 1995.

[169] A.N. Tikhonov, "Resolution of ill-posed problems and the regularization method (in Russian)", *Dokl. Akad. Nauk SSSR*, Vol. 151, pp 501–504, 1963.

[170] A.N. Tikhonov and V.Y. Arsenin, *Solutions of ill-posed problems*, John Wiley & Sons, Washington D.C., 1977.

[171] H.J. Trussell, "The relationship between image restoration by the maximum *a posteriori* method and a maximum emtropy method", *IEEE Trans. Acoust., Speech and Sig. Proc.*, Vol. 28, No. 1, pp 114–117, February 1980.

[172] M. Unser, "Splines: a perfect fit for signal and image processing", *IEEE Signal Processing Magazine*, Vol. 16, No. 6, pp 22–38, November 1999.

[173] M. Unser, A. Aldroubi and M. Eden, "B-spline signal processing: part I", *IEEE Trans. Sig. Proc.*, Vol. 41, No. 2, pp 821–833, February 1993.

[174] M.A. van Dam and R.G. Lane, "Wave-front slope estimation", *J. Opt. Soc. Am. A*, Vol. 17, No. 7, pp 1319–1324, July 2000.

[175] M.A. van Dam and R.G. Lane, "Theoretical performance of phase retrieval on a subdivided aperture", *To be published in Opt. Eng.*, Vol. 41, No. 6, June 2002.

[176] H.L. van Trees, *Detection, estimation, amd modulation theory, vol. I*, John Wiley & Sons, New York, 1967.

[177] H. Wackernagel, *Multivariate Geostatistics*, Springer-Verlag Berlin Heidelberg, Germany, 1995.

[178] E. Wallner, "Optimal wave-front correction using slope measurements", *J. Opt. Soc. Am. A*, Vol. 73, No. 12, pp 1771–1776, December 1983.

[179] D.O. Walsh and P.A. Nielson-Delaney, "Direct method for superresolution", *J. Opt. Soc. Am. A*, Vol. 11, No. 2, pp 572–579, February 1994.

[180] G.P. Weigelt, "Modified astronomical speckle interferometry", *Opt. Comm.*, Vol. 21, pp 55–59, 1977.

[181] K.W. Wolf, *Integral Transforms in Science and Engineering*, Plenum Press, New York, 1979.

[182] P.W. Wong and C. Herley, "Area-based interpolation for scaling of images from a CCD", *IEEE ICIP*, pp 905–908, 1997.

[183] Y. Yang, N.P. Galatsanos and H. Stark, "Projection-based blind deconvolution", *J. Opt. Soc. Am. A*, Vol. 11, No. 9, pp 2401–2409, September 1994.

[184] Y.-L. You and M. Kaveh, "A regularisation approach to joint blur identification and image restoration", *IEEE Trans. Im. Proc.*, Vol. 5, No. 3, pp 416–428, March 1996.

[185] Y.-L. You and M. Kaveh, "Blind image restoration by anisotropic regularization", *IEEE Trans. Im. Proc.*, Vol. 8, No. 3, pp 396–407, March 1999.

[186] H.T. Yura and M.T. Tavis, "Centroid anisoplanatism", *J. Opt. Soc. Am. A*, Vol. 2, No. 5, pp 765–773, May 1985.