Dissertation
submitted to the
Combined Faculties of the Natural Sciences and for Mathematics
of the Ruperto-Carola University of Heidelberg, Germany
for the degree of
Doctor of Natural Sciences

presented by

David Baddeley, MSc (Physics)

born in: Auckland, New Zealand

Oral examination: 31 October 2007

Precision measurements with SMI and 4Pi Microscopy

Referees: Prof. Dr. Cristoph Cremer

                                Prof. Dr. Josef Bille

Precision measurements with SMI and 4Pi Microscopy

Newly developed far field light microscopy techniques such as Spatially Modulated Illumination and 4Pi microscopy allow structural measurements on biological objects with a precision very much better than the conventional imaging resolution. This thesis focuses on the improvement of this measurement accuracy, encompassing both changes to the experimental SMI setup and new data analysis procedures. An aberration resistant fit algorithm for SMI microscopy is presented along with similar model based procedures for the analysis of 4Pi and confocal data. These are combined in a framework which allows for robust automatic analysis. A method for deconvolving 4Pi data which shows a refractive index mismatch induced phase shift is also introduced.

The devices and the associated algorithms were applied to various biological problems most notably the size of replication foci and the structure of the inner kinetochore. Virtual microscopy simulations were performed to support the interpretation of this data and to obtain insight into the results of the SMI gene structure measurements performed by other group members.

Präzisionsmessungen mittels SMI und 4Pi Mikroskopie

Neu entwickelte Methoden der Fernfeld-Lichtmikroskopie wie „Spatially Modulated Illumination" und 4Pi -Mikroskopie erlauben die Bestimmung von Strukturparametern von biologischen Objekten mit einer Genauigkeit, die deutlich besser als die herkömmliche optische Auflösung ist. Diese Dissertation legt den Schwerpunkt auf die Verbesserung dieser Messgenauigkeit, wobei sowohl Änderungen des experimentellen Aufbaus des SMI-Mikroskops als auch neue Methoden der Datenanalyse zum Einsatz kommen. Detailliert beschrieben wird ein Fit-Algorithmus für die SMI-Mikroskopie, der gegen Vezerrungen der Punktbildfunktion stabil ist, sowie ähnliche modell-basierende Methoden zur Analyse von 4Pi und konfokalen Daten. Diese Techniken sind derart zusammengestellt, dass eine robuste automatische Auswertung möglich ist. Eine Methode für die Entfaltung von 4Pi Bildern, die durch eine Fehlanpassung des Brechungsindex eine Phasenverschiebung aufweisen, wird ebenfalls vorgestellt.

Die Mikroskope und die damit verbundenen Algorithmen werden auf verschiedene biologische Fragestellungen, insbesondere auf die Größen von Replikations-foci und die Struktur des Inner-Kinetochor-Komplexes, angewandt. Es wurden Simulationen, basierend auf virtueller Mikroskopie, durchgeführt, um die Interpretation der Daten zu unterstützen sowie die Ergebnisse der SMI-Messungen von Gen-Strukturen besser verstehen zu können.

Contents

Notes for Readers of This Thesis

References to Matlab scripts

As part of this thesis, several Matlab scripts have been written for the analysis of SMI data. These form a package which goes under the name of *sviewer*[1] and is currently intensively used within the group. The package has grown throughout the course of my thesis, and is unfortunately somewhat sparsely documented. As I anticipate that my thesis might have to serve as de-facto documentation for future students I have included boxed sections such as that found below within the thesis at various locations. These boxed sections show which commands I would use to perform the task being discussed.

> **Sviewer Scripts**
>
> `script`: This does *something* when called from the Matlab command window.
>
> `another script`: Does something else, possibly depending on the results obtained in `script`.

Modulation Depth/Modulation Contrast

The definition of the modulation parameter used for SMI size measurements is somewhat different from that used in some older theses. For more details see Section 3.4.1.

---

[1]short for *SMI Data Viewer*, one of the first scripts

Chapter 1

Introduction

The resolution obtainable with a far-field microscope is limited by the laws governing the diffraction of light. More specifically, it is determined by the wavelength of light and the angles at which it can be made to interfere[1] [1]. For conventional optical microscopy, with good optics and oil immersion, this is approximately 240nm laterally in the focal plane and 550nm along the optical axis. Lens manufacturing having already reached a level of sophistication capable of achieving diffraction limited performance, the middle of the 20th century saw efforts to improve the resolution focusing on progression to shorter wavelengths and the development of electron microscopy.

Whilst electron microscopy remains the gold standard in high resolution microscopy, it has several disadvantages over light microscopy for many kinds of biological problems. The most important of these are as follows: specimens for electron microscopy must be dead, they must be thin (either surfaces of cryosections through larger structures), the possibilities for achieving contrast and specific labelling are very much more limited than with optical microscopy, and the whole process is somewhat time consuming.

A large number of interesting biological structures and processes occur on a scale which is only a little under the conventional imaging resolution. There has thus been a resurgence of interest in optical techniques which has resulted in several incremental improvements, starting with confocal microscopy, which is based on the idea that one can influence the apparent structure of the object being imaged by introducing structure into the illumination light. Other methods have built on this principle of structured illumination, often combining it with techniques which improve the effective range of angles which are captured/illuminated giving us SWM, SMI, PEM, OMX, 4Pi, axial tomography[2] and others. More recently, it has become possible to exploit non-linearities in the fluorescence process to circumvent the resolution limit completely and attain an effective resolution limited only by the signal level. Such nonlinear techniques include STED, PALM, and S-PEM.

In addition to the possibilities of improving the imaging resolution it is possible to estimate certain parameters, e.g. position or size, of an object with an accuracy significantly better than the resolution. Such estimations invariably require some prior information or

---

[1]for a simple lens, these angles are dependant on the numerical aperture (NA) of the lens

[2]abbreviations and the associated techniques are described in detail in Chapter 2

assumption about the object; in the case of position measurements on an isolated object positivity is generally sufficient. Spatially Modulated Illumination (SMI) microscopy is a technique developed in the Cremer lab which uses structured illumination in order to obtain a higher precision in the measurement of both object sizes and their positions.

My thesis has focused on improvements to the SMI microscopy technique, and its application to biological problems in the field of genome structure. Particular emphasis has been placed on improving the measurement accuracy and usability. This involved changes in the device hardware, control software and, above all, the data analysis algorithms. As not all problems can be addressed with SMI measurements, I have also spent some time on 4Pi microscopy, and on the analysis of confocal images.

Improvements to SMI experimental setup

Comparatively minor improvements have been made to the existing SMI setup. These are detailed in Chapter 9 and Appendix A. The most important changes are the introduction of a new camera, a new laser line at 568nm, and a rewrite of the control software.

Data analysis for SMI, 4Pi, and confocal microscopy

Several improvements to the procedure for the analysis of SMI data have been made, including the introduction of a fit algorithm which copes with an asymmetry in the axial profile, object identification, lateral position determination, and a high degree of automation. With the help of a new axial model function the tools developed for the SMI have also been applied to 4Pi data. As several of the structures we wish to measure may be too large for the SMI, routines (a lateral Gaussian fit and the *multi-point model*) were developed for high precision size estimation from confocal images. All data analysis routines are combined in a toolbox for Matlab[3]. For details of the data analysis algorithms see Chapters 7, 10 and 14.

4Pi deconvolution

Refractive index mismatch or slight misalignment can result in substantial, position dependant, changes in the 4Pi imaging properties. Due to the subsequent lack of shift invariance this complicates the deconvolution process needed to produce sidelobe free data. An algorithm which allows the deconvolution of 4Pi images in which such a phase change is present was developed and is presented in Chapter 8.

---

[3]The Mathworks Inc, Natwick MA

SMI measurements of replication foci and inner kinetochore proteins

The SMI microscope was applied to two biological problems in the field of nuclear structure. The first of these addressed the size of foci in which DNA is replicated during the S-phase, and went further to investigate whether there is any change in size as one progresses through the S-phase. Both the PCNA protein associated with the replication machinery and the newly replicated DNA were labelled. The result of these measurements, discussed in more detail in Chapter 11 was that the size of the foci is approximately 125nm and that this size is maintained throughout the S-phase.

A second biological application was the proteins CENP-A and CENP-C which form part of the inner kinetochore, the innermost part of the centromere where the spindle fibres attach during mitosis. These were examined in interphase nuclei, using multi-colour XFP fusion protein labelling. The results of these experiments (Chapter 13) are somewhat less conclusive than those with the replication foci, indicating a size near or larger than the upper limit for SMI observations and a structure which is somewhat more compact than most gene regions.

Simulation of SMI measurements of gene structures

Despite not being directly responsible for the various attempts made at SMI measurements during my thesis, I have been extensively involved in a supervisory/advisory role in both the data acquisition and analysis in several of these projects. As such I have been privy to several notable experimental observations, and the general realisation that the interpretation of SMI gene-region data is not necessarily trivial. Some of these observations are:

- Typical, Bacterial Artificial Chromosome (BAC) clone labelled, gene regions of a size of around 150-250kbp, show significant lateral structure, implying a size larger than the lateral resolution ($\sim$250nm).

- These same BAC clones often show significant, if still somewhat shallow, modulation which would correspond to a size in the 150-220nm range.

- Regions of the same length which are indirectly labelled with antibodies show more modulation.

- Regions of the same length labelled using COMBO-FISH typically show lots of modulation and a very large spread of *'size'* estimates.

- Little or no correlation is observed between modulation depth and probe length.

In order to attempt to explain some of these observations and to establish guidelines for the interpretation of SMI data on gene-regions, I performed several simulations in

which simulated gene structures were used to generate virtual SMI images which were then analysed with the usual analysis algorithms. The results of these simulations are shown in Chapter 12 and indicate that although a correlation between modulation depth and size for gene regions cannot, in general, be assumed, there are certain situations (very small structures) where this might still be valid. They also suggest the possibility to infer other structural information not directly related to size from the modulation shown by larger structures.

Part I

Background

Chapter 2

Fluorescence Microscopy

The discussion of microscopy in this chapter is by no means complete. Its aim is to provide a bare minimum of information in order to support the treatment of a microscope system as a linear system, and to underline the fundamentals of resolution enhancement. As such a detailed explanation of optical theory has been omitted. Good resources are [15] and [36] as well as several previous theses.

2.1 Fluorescence

Fluorescence is observed in many atomic or molecular systems and is characterised by the absorption of a photon and the emission, after a short delay, of a photon at longer wavelength. Fluorescence is an important contrast medium for light microscopy, with the most commonly used fluorescent substances being synthetic organic dyes, fluorescent proteins, and recently semiconductor 'quantum dots'.

2.2 Image Formation and Resolution

2.2.1 Image Formation as a Linear System

A linear system can be fully described by its impulse response. With knowledge of the impulse response, h, it is possible to calculate the system output, f, for an arbitrary input, g through the convolution

$$f = g \otimes h.$$

Whilst the optics in the microscope detection path can be considered linear at the intensities relevant for most microscopy approaches, detection methods (e.g. a CCD camera) are typically only sensitive to intensity. As taking the absolute value of a complex quantity is not normally considered to be a linear operation, one might suspect that the concept of an impulse response is not generally applicable to microscopy. When the contrast method

is insensitive to phase (e.g. fluorescence, incoherent transmitted light[1] or when the object does not produce any phase variations) the object is a real quantity. In common contrast methods the object is physically constrained to be positive (fluorescence) or negative (absorption)[2]. For a real quantity which is constrained to be either positive or negative, taking the absolute value does not represent a non-linear operation, and the concept of an impulse response can be used, giving us the so-called 'Point Spread Function' (PSF).

A somewhat more intuitive description is as follows: The Point Spread Function is the image of an infinitely small point. In order to form an image we place a copy of the PSF at every point in the object, weighted with the amplitude of the object at that point. Expressing this mathematically, we can describe our object as a weighted sum of delta functions;

$$g = \sum_{i=0}^{\infty} A_i \delta(\mathbf{r}_i)$$

Plugging this into the convolution integral,

$$f = \iiint_{-\infty}^{\infty} g(\mathbf{r} - \mathbf{s}) \cdot h(\mathbf{r}) \, d\mathbf{r}_x \, d\mathbf{r}_y \, d\mathbf{r}_x$$

gives

$$f = \sum_{i=0}^{\infty} A_i \iiint_{-\infty}^{\infty} \delta(\mathbf{r}_i - \mathbf{s}) \cdot h(\mathbf{r}) \, d\mathbf{r}_x \, d\mathbf{r}_y \, d\mathbf{r}_x$$
$$= \sum_{i=0}^{\infty} A_i \cdot h(\mathbf{s}_i)$$

Note that the linearity required for a PSF based description to be valid applies to the operation which maps the object to the image. In *nonlinear* microscopy techniques such as multi-photon excitation or STED, the nonlinearity is present in the excitation and can be effectively removed by replacing the true excitation structure with the fluorophores (non-linear) response to it[3], and the imaging can again be described as a linear process.

As convolution corresponds to multiplication in the Fourier domain, one can alternatively describe the system in terms its frequency response - the 'Optical Transfer Function' (OTF). This is obtained by taking the Fourier transform of the PSF. The PSF and OTF of a conventional widefield fluorescence microscope are shown in Fig. 2.1. This represents the standard against which all improved microscopy methods are measured.

---

[1] these methods are not entirely phase insensitive - as seen in the problems introduced by refractive index mismatch

[2] Objects with both positive and negative values can in principle be made positive by adding an offset, for example in transmitted light.

[3] assuming negligible absorption/scattering
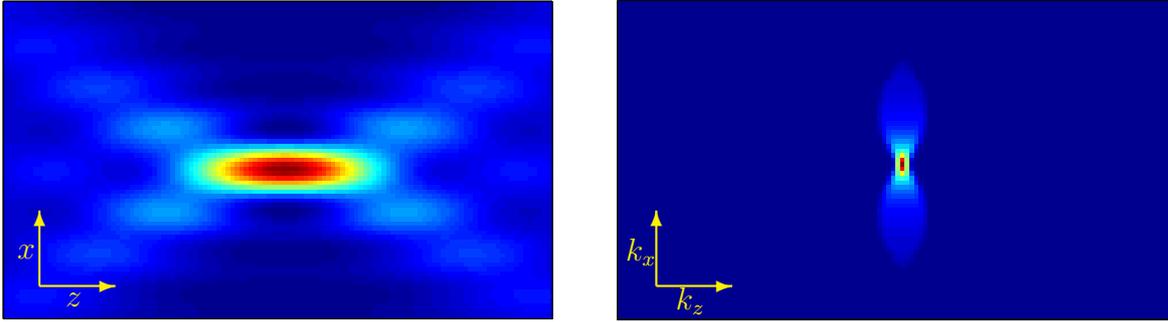
Figure 2.1: Widefield PSF (left) and OTF (right).

Note on PSF simulations

The PSF simulations used to produce the PSF and OTF images presented here and in the SMI and 4Pi sections are rather naive, being 2D simulations performed under the paraxial approximation. For confocal and 4Pi-C imaging modes, an infinitesimally small pinhole is assumed and excitation and emission wavelengths are assumed to be equal (this has a particularly large effect for the 4Pi-C measurements). The Numerical Aperture (NA) is also typically wrong. The results are thus best be interpreted as a schematic representation of the respective PSFs.

2.2.2 Resolution

There are various criteria for measuring the resolution of an optical system. The most common are detailed as follows:

Abbe Criterion

The Abbe criterion [1] stipulates the period of the smallest periodic structure which can be imaged through the lens, based on the refractive index of the immersion media n, the lens aperture angle θ, and the wavelength of the illuminating light λ. n and θ are often combined in a term $\mathrm{NA} = n\sin(\theta)$, called the numerical aperture. It is derived by considering the fact that the 0th and 1st orders of light from a diffraction grating need to be collected to construct an image of the grating, giving the following expression:

$$D = \lambda/2\mathrm{NA}$$

Rayleigh Criterion

According to the Rayleigh limit [78], two point like objects are said to be resolved when the distance between them is greater than the distance between the maximum and the first minimum of the diffraction pattern. For a 1-D geometry (e.g. a slit like aperture), this is

equivalent to the Abbe limit. For a spherical aperture, however, the following formula is obtained:

$$D = 0.61\lambda/\mathrm{NA}$$

Sparrow Criterion

The Sparrow limit [90] is the value of separation at which the dip in intensity between two points vanishes. In a mathematical sense this occurs at the distance where the first and second derivatives of the function obtained by the summation of the two of the PSFs becomes zero at the midpoint between the objects. For a circular aperture the following formula is obtained:

$$D = 0.47\lambda/\mathrm{NA}$$

FWHM Criterion

Another commonly used measure of resolution is the Full Width at Half Maximum (FWHM) of the PSF. In astronomy this is also known as the Houston Criterion [58]. It has the advantage of being applicable to arbitrary distributions[4]. For a circular aperture the FWHM has the following form:

$$FWHM = 0.51\lambda/\mathrm{NA}$$

Regardless of the exact formulation chosen, the resolution is a quantification of the smallest feature size which can be imaged, and hence the smaller the values, the better. A complementary, and in the context of the Abbe definition obvious, way to view resolution is to look at the region of OTF support. In contrast to the PSF based view of resolution, the larger the OTF the better. A good OTF is characterised by not only having some support over a large band, but by also having substantial support over a good portion of that band. Weak areas in the OTF will result in poor contrast for object frequencies in these areas. This is the reason why OMX microscopy is so much better than traditional confocal microscopy despite a similar region of OTF support.

2.3   Structured Illumination

One way of extending the resolution of a light microscope is structured illumination. An instructive analogy is to compare a microscope to a radio. Our ears are capable of detecting signals in a limited bandwidth from around 20Hz-20kHz. If we were to directly generate acoustic waves from an AM radio signal at ∼800kHz we would be unable to hear anything. In order to detect the information carried by the radio signal we multiply with the carrier frequency at 800kHz, generating the difference frequency which is now within our hearing range and can be converted to audio and understood.

---

[4]The Rayleigh limit fails spectacularly for a Gaussian distribution, for example.

In a microscope the objective lens corresponds to our ear - it can only pass a limited band of frequencies determined by its OTF. In order to access out of band information, the light collected by the objective must be the multiplication of the object with our *carrier frequency*. This is rather easy, as we know that the intensity of the emitted fluorescenceis proportional to both the illumination intensity and fluorophore density. If we introduce a pattern in the excitation intensity, the light collected by the lens will correspond to the mix frequencies from object and illumination.

The feature size in the illumination structure is, as per normal imaging, limited by the wavelength of the light. This poses a practical limit on the resolution improvement obtainable using linear structured illumination of approximately a factor of two, depending on the method used. The total response of a structured illumination system is obtained by multiplying the illumination and detection PSFs, or convolving the OTFs.

PEM, OMX etc . . .

Perhaps the most easily understandable variant of structured illumination microscopy is Patterned Excitation Microscopy (PEM) [46, 45], also known as OMX [37] and under various other acronyms. These forms of microscopy use structured illumination in the form of a lateral sinusoidal wave obtained by imaging a diffraction grating into the object space. The pattern frequency must pass through the objective, and hence has a frequency which is chosen to be slightly less than the critical frequency at the illumination wavelength. Due to the quadratic dependance of intensity on the E-field, the form of the pattern is $\sin^2$ or equivalently $1 + \sin$. The light collected by the objective is thus the sum of the normal response with mix products at $\pm$ the pattern frequency. These components can be separated by analysing images taken with different phases of the illumination pattern, and a high resolution image reconstructed by numerically moving the mix products back to their original locations in frequency space. The resolution improvement is a little less than a factor of two.

The Zeiss Apotome system is also a form of lateral structured illumination, but uses a much larger grating period, and thus obtains a much smaller resolution increase. It is unfortunately somewhat let down by its current reconstruction algorithm. Optical sectioning is achieved.

HELM, SW-TIRF

Harmonic Excitation Light Microscopy (HELM) [34] and Standing Wave TIRF [19] are similar to PEM, with the principle difference being that the illumination pattern is not imaged through the objective lens, but created by the interference of independently introduced laser beams. Because the illumination does not need to pass the objective, a resolution enhancement by a factor of more than two can be achieved, although the period of the pattern is still limited by the laser wavelength.
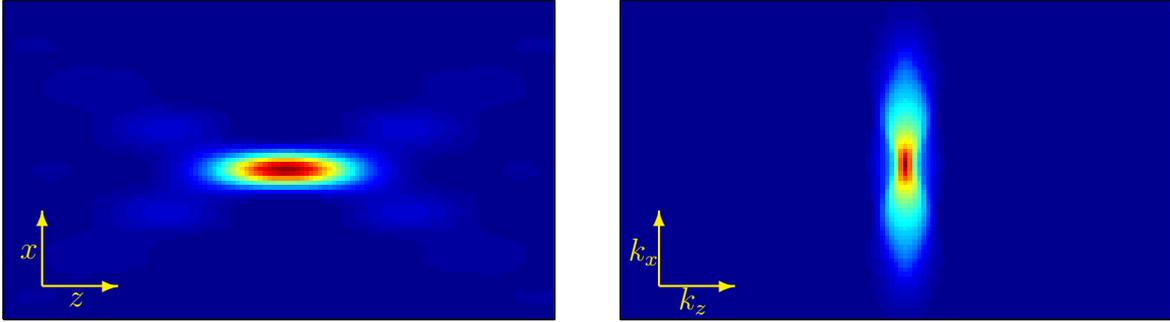
Figure 2.2: Confocal PSF (left) and OTF (right).

Confocal Microscopy

Confocal microscopy is, strictly speaking, a form of structured illumination microscopy. Due to the large proportion of low frequency content in the illumination OTF, the effective resolution improvement is rather mediocre despite a total area of OTF support comparable to other structured illumination techniques. Optical sectioning is obtained. The confocal PSF and OTF are shown in Fig. 2.2. A more detailed description of confocal microscopy can be found in, for example, [107].

SMI, 4PI, I5M, STED, etc

All use structured illumination in combination with other techniques, and will be described in more detail later.

2.4   I5M

I5M combines widefield interferometric structured illumination with interference of the detected light to achieve a region of OTF support similar to that obtained with one-photon 4Pi-C. The OTF magnitude at certain frequencies is smaller than that of 4Pi microscopy making reconstruction potentially more difficult (although the superior detection efficiency offered by a widefield detection geometry may offset this to some extent) [40, 38, 12].

2.5   Non-linear Processes

If the response of the fluorophores to the excitation light can be made non-linear, higher order terms are present in the mix products obtained when using structured illumination. A quadratic nonlinearity will thus result in mix terms at 0, $\pm f_p$, and $\pm 2f_p$ where f_p is the frequency of the illumination pattern. A cubic nonlinearity gives $[0, \pm f_p, \pm 2f_p, \pm 3f_p]$ and so on. Nonlinearities can thus be exploited to access object information at even higher object frequencies. In order to do this however, the sample must be exposed to a significantly

higher light dose as either a larger number of phase sections (saturated PEM) [47, 39], or a smaller voxel size (multi-photon processes, STED, etc...) must be used.

2.5.1 Two Photon Excitation

Two photon excitation is a process whereby a fluorophore can be excited by the simultaneous absorption of two photons at approximately half the one-photon absorption wavelength. It is a complex quantum mechanical process but luckily a full understanding is unnecessary for its usage in microscopy. The probability of a two photon process is proportional to the intensity squared, making two photon excitation a 2nd order nonlinearity. The 2-photon cross-section of typical dyes is rather small, requiring a high instantaneous intensity to obtain a significant number of 2-photon events. Excitation is thus performed using a pulsed IR laser with ps or fs pulses (typically a Ti-Sapphire or the combination of Ti-Sapphire with an OPO). Peculiarly the 2-photon spectra of fluorescent dyes does not directly correspond to the 1-photon spectra, being broader, and with a peak at a wavelength somewhat shorter than half the 1-photon excitation peak.

Two photon excitation confocal microscopy [25] is well established and finds its application in deep tissue imaging due to reduced absorbtion and scattering of the near IR light. A significant resolution improvement is not obtained as, whilst the 2nd order nonlinearity should give a factor of two in resolution, operating at half the wavelength counteracts this. Optical sectioning is however obtained without the need for a pinhole.

Three photon excitation is also possible, with a probability even lower than that for the 2-photon process. It is a 3rd order nonlinearity.

2.5.2 Saturation

Another nonlinearity is the saturation of fluorescence excitation. As saturation is an exponential nonlinearity, an infinite number of higher order terms are produced, and an infinite resolution is theoretically possible. The amplitude of the higher order terms, however, decreases rapidly placing practical limits on the resolution achievable. Classical excitation saturation requires high illumination power, and faces significant problems due to photobleaching and specimen damage. This has not stopped an optical resolution of $\sim$50nm [39] from being achieved with this method.

STED

STimulated Emission Depletion (STED) [59, 26] microscopy uses another saturable process - stimulated emission. Whilst high intensities are required in order to saturate the stimulated emission, these remove the fluorophores from the excited state and the extent of photodamage is thus much lower. Due to the experimental geometry, a high resolution image can also be obtained directly without significant postprocessing. The non-linearity is also exponential. The current record for STED lies at around 16nm optical resolution in one direction [106].

Photoswitching

A third saturable process is the photoswitching seen in GFP variants such as DRONPA. This can be used in a (mathematically) similar way to STED, with the advantage that much lower laser powers are required. A microscope built to exploit this non-linearity is likely to have a widefield geometry and use structured illumination as in PEM[5]. A similar concept, under the acronym RESOLFT [54], has been proposed by S. Hell and can be seen as generalisation of STED to arbitrary multi-state systems.

2.6 Resolution Through Localisation ('Pointilism')

Another school of thought places the emphasis on localisation precision rather than absolute imaging resolution. If two objects can be separately identified, then their individual and hence mutual positions can be determined with a much higher precision than the imaging resolution. This is the basis of several techniques including SPDM. The major disadvantage of most techniques is the limited number of signals which can be localised within one microscope detection volume. This limitation precludes a true resolution improvement, and has been overcome in PALM.

2.6.1 SPDM

Spectral Precision Distance Microscopy (SPDM) [16, 22, 29, 27] is a technique whereby objects which would normally not be resolved are separated by labelling with different spectral signatures. Major limitations of the technique are the limited number of spectral signatures available (approximately 4 in a modern confocal microscope[6]) and the necessity of correcting the chromatic shift. Whilst the chromatic shift induced by the objective lens alone can be accurately measured and corrected for, refractive index mismatch within biological specimens produce spatially dependant variations which are extremely difficult to measure and correct. An additional concern with respect to several of the SPDM measurements performed in our group is the labelling of the individual object regions. The regions are often relatively large, and the labelling inefficient. This can result in changes in apparent position due to the statistical nature of the labelling. All in all, our experiments (S. Fenz [32], H. Mathee [71], and my own measurements on the kinetochore structure (Section 13.1)) suggest that an effective distance resolution better than $\sim$50nm is likely to be very difficult to obtain in normal biological specimens.

Extensions to the SPDM technique include the use of fluorescence lifetime, which has the potential to either increase the number of spectral signatures, or to reduce the problems associated with chromatic shift, and two photon excitation. When performing confocal

---

[5]Priv. Comm. R. Heintzmann

[6]Whilst several confocal manufacturers claim more channels, this usually involves linear unmixing and other tricks and still doesn't result in complete separation of the channels. As such it might be possible to perform imaging applications, but any quantitative analysis of the data when using more than around 4 channels is more than questionable.

SPDM, the majority of the chromatic shift comes from the chromatic shift present in the excitation[7]. Using 2-photon excitation it is possible to excite different spectral signatures with the same excitation wavelength, thus eliminating the excitation shift. If the detection pinhole is then relatively wide open, the contribution of the detection chromatic shift is minimal. This has important consequences for the comparison of SMI and 4Pi distance measurements.

2.6.2 PALM

Photo Activation Localisation Microscopy (PALM) is a recently developed technique [11, 53] which uses the photoconversion processes exhibited by certain fluorescent proteins in order to overcome the limit on the number of points which can lie within the observation volume of the PSF. In short, short wavelength light[8] is used to convert the protein from a non-fluorescent state to a fluorescent one. The trick is that the photoactivation is performed with a very low power such that the probability of activation is small. Only a small fraction of the available pool of protein will thus be activated. If the activation power is chosen appropriately, the chance that two activated proteins are present within the same detection volume will be negligible. The positions of all activated proteins can then be measured. If the active proteins are then bleached, and the activation process repeated it is possible to iteratively measure the locations of a large fraction of all protein molecules and thus create an image with a resolution equal to the localisation accuracy.

PALM has been demonstrated with the PA-GFP, Kaede and Eos-FP GFP variants as well as caged Rhodamine. A similar technique (named STORM) has been shown to work with a switchable construct formed with Cy3 and Cy5 [82], and it is likely that other switchable fluorophores such as DRONPA could be used. PALM, in its current state of development, does not suffer from problems associated with chromatic shift. It is however time consuming and susceptible to mechanical drift. To date, only one colour measurements on two-dimensional specimens have been demonstrated. The extension to further colours and/or 3-d measurements presents a significant technical and biochemical challenge.

2.6.3 Other Notable 'Localisation Microscopy' Examples

Other interesting experiments with relevance to localisation microscopy include the separation of individual quantum dots based on the statistical nature of their blinking [67], and the use of single molecule spectroscopy. When the sample is cooled below ∼3K the higher vibrational and translational levels are frozen out and, it is possible to distinguish individual fluorophores by the slight differences in the energy of the 0-0 transition, dictated by their surroundings. It is then possible to localise each molecule independently and create a high resolution image as for PALM [103, 14].

---

[7]This is an empirical observation, the cause of which is hard to pinpoint, although slight differences in angle or collimation of the various laser sources could provide an explanation.

[8]typically the 405nm diode laser line

Chapter 3

SMI Microscopy

The original idea on which SMI microscopy is based is Standing Wave Microscopy (SWM)[5, 65]. The original articles focused on an improvement in axial imaging resolution through the use of standing wave microscopy. Due to the large missing portion in the OTF, imaging applications are fraught with ambiguity, and more or less limited to thin objects such as the cytoskeleton in the leading edge of cells. A technique for the improvement of the OTF support through *excitation field synthesis* was published [63]. It appears that, on realising that even with excitation field synthesis SWM falls short of a true improvement in imaging resolution, the original authors abandoned further work on the subject. Interestingly Zeiss also experimented with SWM shortly after its initial publication [33]. The research into standing wave microscopy, under the acronym Spatially Modulated Illumination (SMI) microscopy, performed in the Cremer group [3, 31, 84, 91][1] has pursued a subtly different approach - the improvement of measurement accuracy. Unlike a true resolution improvement, this is indeed mathematically possible with a standing wave microscope, albeit with the requirement for some rather severe assumptions about object form. The theoretical background to SMI microscopy has been treated in detail in [91],[3] and [104], the treatment here will thus be restricted to an summary of the most important results, with the body of the proofs and derivations omitted.

3.1   Experimental Setup

The Spatially Modulated Illumination microscope is a widefield fluorescence microscope with structured illumination in the axial direction. This illumination structure is created through the interference of two counter-propagating laser beams, and thus has a $\cos^2$ form. The setup has evolved with time, but is in principle the same as detailed in [91]. Several components have been changed in the course of my thesis, so a brief description of the current stand follows.

The experimental setup is shown in Fig. 3.1. An Argon ion laser (currently Lexel 95-4, although a Spectra-Physics 2060 was used for some time whilst waiting for the opportunity

---

[1]NB: not, by any means, a complete list
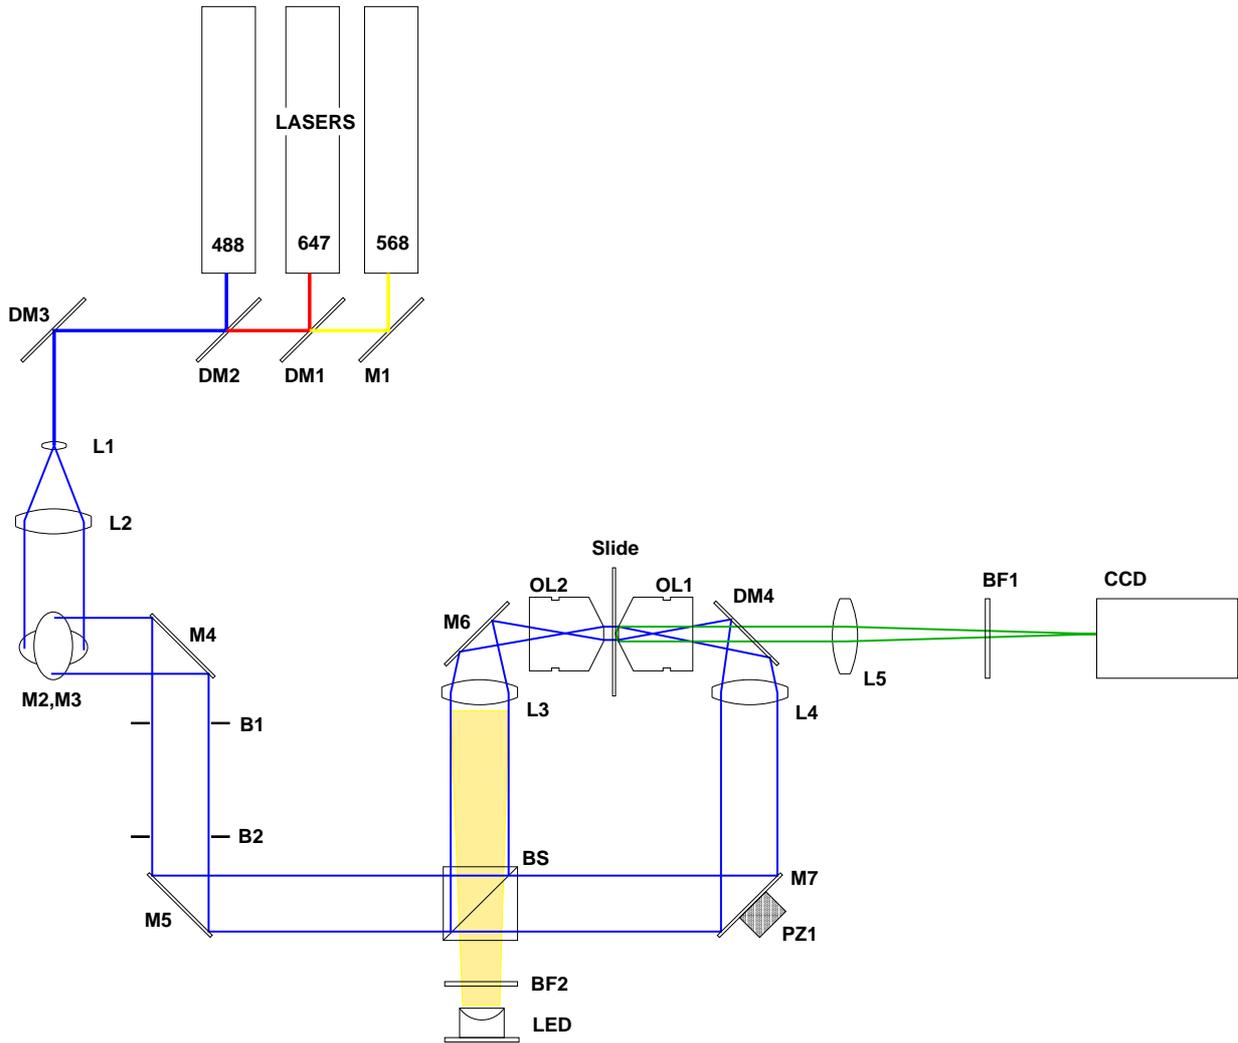
Figure 3.1: Experimental setup of the SMI microscope. For a complete desciption of the individual components see text.

to obtain a replacement tube) provides laser light at 488nm (approximately 750mW max power with current laser and tube). A Lexel 95-K Krypton laser provides ∼200mW at 647nm, and a new Lexel 95L-K Krypton laser was installed as part of this thesis to provide 568nm excitation (also ∼200mW). All lasers operate with a prism and in-cavity etalon to ensure single frequency operation. The three laser beams can be independently switched with shutters before being combined with the dichroic mirrors **DM1** and **DM2** (AHF part numbers F33-488). The dichroic **DM3** is identical to the excitation dichroic **DM4** and takes the place of a cleanup filter, reflecting the three laser lines whilst transmitting remaining plasma light (this proved necessary to remove a disturbing background signal seen with both 647 and 568nm excitations). The combined laser beam is then expanded to a beam diameter of approximately 2cm. The beam expander consists of two achromats

(**L1**& **L2** - focal lengths 10 and 100mm resp.) and is assembled in a microbank cage from Linos.

The expanded laser beam is then conveyed to the interferometer, passing a periscope[2] (**M2**, **M3**) which raises the beam to the height of the breadboard and though two iris diaphragms **B1** and **B2** which facilitate independent adjustment of the laser combination, beam expansion, etc, and the interferometer itself. They also act as an illumination field stop. On entering the interferometer, the beam is split using a cube beam splitter (**BS**, Edmund Optics - part number NT47-009), and directed to each of the two opposing objectives (**OL1**, **OL2**). An achromatic lens (**L3**, **L4**, focal length 60mm, Linos cat num. 322266) before each objective focuses into the back focal plane of the objective, resulting in a collimated beam in the object space. The objectives currently in use are oil immersion lenses from Leica with a *magnification*[3] of 100x and a NA of 1.4. The mirror **M7** is mounted on a piezo translation stage (**PZ1**), allowing the relative phase in the two interferometer arms to be varied. The sample is positioned between the two objectives using a 3-axis (x,y,z) stepper motor stage for coarse movements and a 2-axis (y,z) piezo stage for fine control and object scanning. For full details of the object holder, stepper motor and piezo stages see [91] or [3].

The fluorescence from the sample is collected by the right hand objective, passes through a dichroic mirror (**DM4**, AHF - F63-488) which reflects the laser lines and is imaged onto the CCD camera (PCO SensicamQE[4]) through a Leica 1.25x tube lens. A blocking filter (**BF1**), in a filter wheel, blocks any remaining laser light, and, depending on filter selection, out of band fluorescence[5]. The following blocking filters are currently available: a triple pass filter which only blocks the laser lines (AHF - F62-568), a dual pass blocking the 488 and 647nm lines[6] (AHF - F52-488), a bandpass for GFP/FITC (525±25nm, AHF - F42-017), a bandpass for TRITC/mRFP (610±30nm, AHF - F42-614), and a bandpass for YFP (535±25nm, AHF - F42-028)[7]. Control of object motion and image acquisition is performed by a PC computer running Windows and specially designed control software.

---

[2]A side effect of the periscope is to rotate the polarisation of the beam. The beam is vertically polarised on leaving the lasers and horizontally polarised after the periscope. Previous setups had used vertically polarised light in the interferometer - no significant difference in performance is observed as polarisation is preserved in both cases.

[3]Strictly speaking, infinity corrected objective lenses on their own do not have a magnification, the magnification figure is the magnification which would be obtained with a hypothetical 200mm focal length tube lens.

[4]for more details concerning the camera see section 9.1.

[5]The blocking filter is positioned after the tube lens for historical reasons; whilst a position before the tube lens would be theoretically preferable I could not detect a significant difference between the two geometries in my experiments and thus reverted to the position after the lens in order to avoid construction of a new filter wheel and because the current geometry effectively protects the camera from dust ingress.

[6]568nm excitation is not blocked.

[7]488nm excitation must be used, which is not ideal - it is also likely that the current dichroic will remove a good chunk of the YFP spectrum in blocking the 568nm line. The use of YFP with the current SMI setup is thus not advisable.
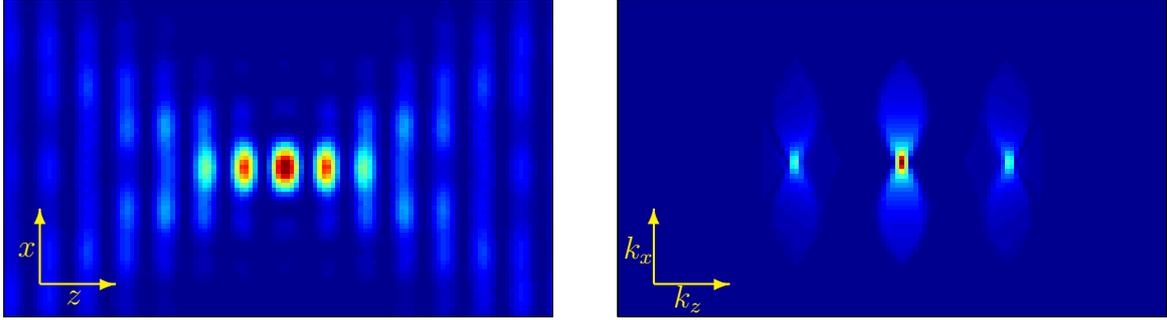
Figure 3.2: SMI PSF (left) and OTF (right).

3.2   Operating Principle

The interference of the laser beams from each of the objectives produces a sinusoidal standing wave pattern along the optical (z) axis which has the form:

$$I(z) = I_0 \cos^2\left(\frac{2\pi n \cos\left(\alpha\right) z}{\lambda_{\text{exc}}} + \phi\right) \tag{3.1}$$

Where λ_{exc} is the excitation wavelength, n the refractive index in the sample, z the distance along the optical axis, ϕ an arbitrary phase offset and α the angle between the two laser beams which, for normal alignment, can be considered to be small. In further consideration, α is assumed to be 0 (any errors due to non-zero α will manifest themselves as a reduction in the apparent refractive index).

   The object is moved axially through the illumination pattern and the microscope focus, with images acquired at each step with the CCD camera. This results in 3D image stacks with a PSF corresponding to the multiplication of a normal widefield PSF and the SMI standing wave pattern. A simulated SMI PSF is shown in Fig. 3.2. In the frequency domain our illumination pattern consists of 3 delta functions at 0 and at $\pm$ the frequency of the the interference pattern. The OTF is thus composed of three copies of the widefield OTF at the locations of the delta functions.

   Looking at the SMI OTF, it is clear why a true improvement in imaging resolution is not achieved - there is a large gap between the widefield portion of the OTF and the additional support provided by the illumination structure. No information is obtained about the object spatial frequencies within this gap. The illumination structure created by the interference between two counter-propagating beams is however the shortest possible structure which can be created at a particular wavelength. The end result is that SMI microscopy, whilst not an imaging technique, quite likely[8] provides the most robust access of any linear structured illumination technique to high axial frequency components.

---

[8]The considerations are such that the results can easily be skewed, depending on assumptions made about s/n and other factors. I5M and 1-photon 4Pi-C both theoretically allow access to higher frequency components, albeit at a low amplitude.

3.2.1 Excitation Field Synthesis/ More Angles

When contemplating equation 3.1, a logical conclusion is that one could fill in a portion of the gap by varying the relative angle of the laser beams. This has been extensively discussed in [3]and [63]. The end result of such considerations is that the maximum region of support which could be obtained in this fashion is equivalent to that of 1-photon 4Pi-A. A discussion of preliminary experiments in this direction can be found in Chapter 15.

3.3 Technical Details

3.3.1 Image Acquisition

Object Scan

In performing an object scan the sample is moved along the optical axis through the focus of the objective (and simultaneously through the standing wave pattern) with an image acquired at each step with the CCD camera. This results in a 3D image stack from which both position and size information can be extracted. A step size of 20nm between slices had previously been used, resulting in a considerable oversampling of the signal. Analysis of the noise behaviour, particularly with respect to the Gaussian readout noise, led to the conclusion that it would be better to take a larger step size and proportionally longer integration time. The axial step size is thus now 40nm. This is still oversampled, given a Nyquist step size of around 70nm, but not as drastically.

Phase Scan

In addition to an object scan, it is also possible to perform a phase scan in which the position of the object remains constant and the position of the standing wave is moved. This has the disadvantage of only being able to measure objects within one plane, and that it is not possible to directly measure the period of the interference pattern (which depends on both the angle between the two beams and the specimen refractive index). Position information is also only obtained modulo the standing wave period.

If, however, the objects in which one is interested are all in the one plane, a phase scan will be faster and will require a lower photon-dose for the same detected intensity. This makes it an attractive possibility for in-vivo measurements. It is also possible to conceive a combined object and phase scan in which the object was moved in relatively large steps (ie $\sim$200-250nm) and a phase scan made at each z position. The position information and ability to analyse 3D objects would then be retrieved. In theory it might be possible for a combined scan to consist of fewer individual images and yet have the same information content as a normal object scan, in practice this is unlikely.

3.3.2 Sample Preparation

Current SMI sample preparation uses standard slides, coverslips, embedding and immersion. This has the disadvantage that light from the 2nd objective must pass through the object slide. For a collimated beam on the optical axis, this is not too much of a problem, although it does become difficult to position the beam waist in the focal plane. Use with converging illumination is however impossible due to the short working distance of the lens. This effectively rules out many possibilities for introducing further structure into the illumination, a problem which can be solved by moving to a preparation technique similar to that used on the 4Pi.

3.4 Size Measurements

As the images are formed through the convolution of the object with the PSF, the fringe structure, or *modulation* seen in the SMI PSF will be blurred out by the object, with little or no modulation visible for large objects. SMI size measurements are based on the quantification of the degree of this blurring. Given an object form, the amount of residual modulation visible in the image can be translated into an estimate of the object size. In order to make such measurements, it is sufficient to contemplate an intensity profile taken through the object in the direction of the optical axis, the so called AID (see Fig. 3.3).
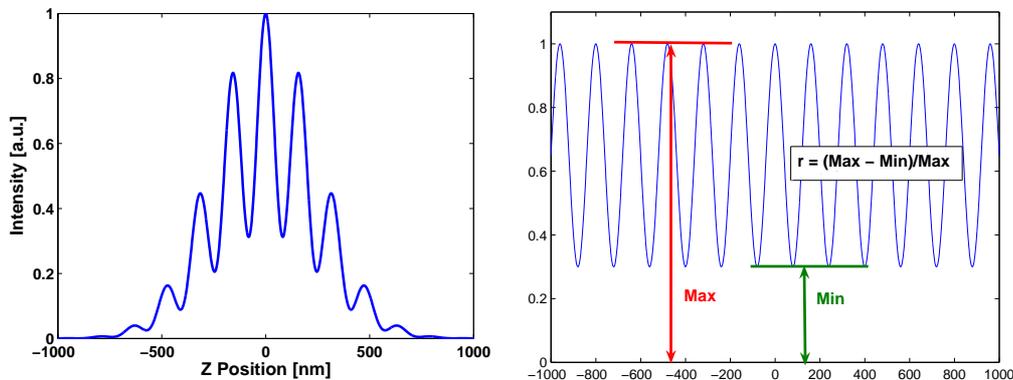
Figure 3.3: Schematic profile along the z-axis through an object (the *AID*, **left**) and definition of the modulation depth parameter used for size estimation (**right**).

3.4.1 Modulation Contrast/ Modulation Depth

In earlier SMI theses and publications this residual modulation was quantified with the use of a so called *modulation contrast*, defined as the ratio between the inner and outer envelopes of the AID. A *modulation 'contrast'* of 0 thus corresponds to the case of an infinitely small object where the complete PSF modulation is retained, and 1 to the case when no modulation is visible in the image. This definitions has the disadvantage that it is

contradictory to the conventional definitions of contrast used in optics and interferometry. One would expect a higher contrast to indicate that more structure can be seen and that the resulting image is less smooth. Indeed interferometric fringe contrast is mathematically defined [43, 60] as the amplitude of the modulating part of the signal over the total amplitude.

To avoid the problems associated with the previous definition of *modulation contrast*, the term *modulation depth* is used throughout this thesis[9]. Modulation depth is defined, as per interferometric fringe contrast, as the fraction of the signal which shows a modulation. For those wishing to compare with previous works using the questionable definition of contrast, the conversion is easy, with the modulation depth, r, being equal to one minus the modulation 'contrast', R.

$$r \equiv \frac{I_{\max} - I_{\min}}{I_{\max}} = 1 - R \tag{3.2}$$

3.4.2 Modulation Depth and Size for Spherical Objects

Given an object form, it is possible to calculate the relationship between modulation depth and size by calculating the convolution between the equation for the axial projection of the object and that for an axial profile through the PSF. As the PSF is relatively broad in the axial direction, one can approximate the form with that of the illumination alone, leading to the convolution of the object and a sin wave. This convolution can often be solved analytically. These calculations are described in detail in previous theses, particularly that from Christian Wagner [104]. The result of such a calculation for a spherical object is shown in Fig. 3.4. One can see that the modulation depth decreases monotonically with increasing object size up to a size of around 240nm, after which a low amplitude ringing is seen. In the region from 0 to around 200nm this relationship is one to one and can be inverted[10] to obtain a mapping from measured object modulation depth to object size. The low gradient at the small size end of the relationship, along with the poor signal to noise ratio typically obtained from small objects results in a practical lower limit at around 40nm.

3.5 Distance Measurements

Due to their smaller width, the axial position of the fringes present in a SMI image can be determined with a much greater accuracy than that of a normal widefield PSF. This is the basis of SMI distance measurements. The challenges presented by SMI distance measurements include the ambiguity as to which fringe corresponds to the object, in the calibration of wavefront tilt and curvature, focal plane tilt, and the relationship between

---

[9]This may seem overly pedantic, and possibly is. The previous definition of modulation contrast was however the cause of considerable confusion at the start of my thesis and a certain pig-headedness prevents me from accepting the somewhat misleading definition

[10]numerically or through use of a lookup table

Figure 3.4: Relationship between diameter and modulation depth for a spherical object, assuming a refractive index of 1.47, complete interference and an angle of 0° between the two beams. Calculations are shown for the 3 wavelengths available in the current SMI setup.
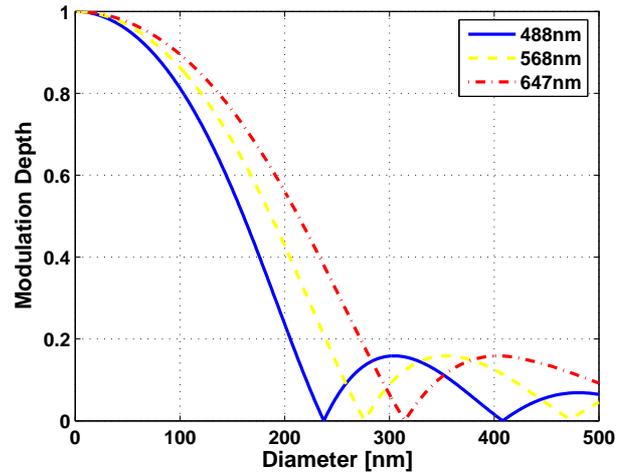
the wavefronts of the different colour channels. This is discussed in further detail in the algorithms section.

3.6   Influence of Object Form

3.6.1   On Size Measurements

As mentioned earlier, size measurements require an assumption regarding the object form. This is extensively discussed in [104]. A short summary follows. Fig. 3.5 shows the curves relating the modulation depth to object size for various object forms. One can see that different object forms lead to drastically different responses. Particularly disturbing is the example consisting of two delta functions which leads to a periodic solution which is not invertible over any part of the curve. The profiles corresponding to the axial projection of a sphere and a cube (green and red respectively) also show ringing, which can be attributed to the high frequency components associated with the object edges.

   Part of the reason that the curves in Fig. 3.5 are so different is the choice of size parameter. In each case this is relatively intuitive (FWHM or diameter), but diameter in particular is often either poorly defined or not a very good indicator of object extension. There are fundamental issues with the use of diameters (or even FWHMs) to compare objects of different shape[11]. Luckily statistics has a well defined and robust definition of distribution width at hand - the standard deviation. Fig. 3.6 shows the modulation curves when plotted against the standard deviation of the dye distribution. The agreement is significantly better than in Fig. 3.5, a strong dependance on object form and the problem of multiple solutions however remain.

---

[11]This also applies to confocal measurements, hence the use of a standard deviation based size parameter in the confocal evaluations (Chapter 14).
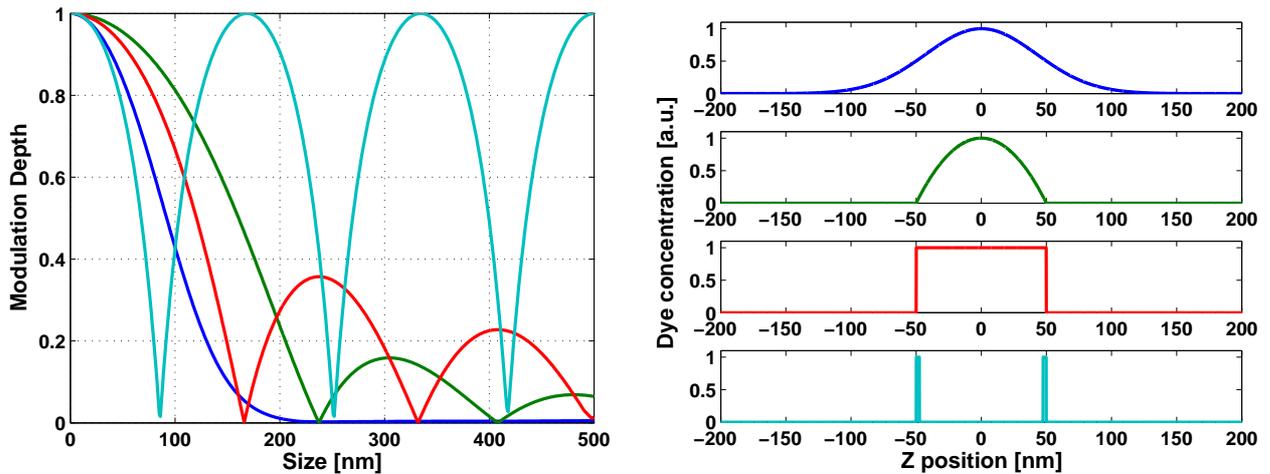
Figure 3.5: Left: Size-modulation depth curves for various object forms. **Right:** The corresponding forms, each with the same *size*. For the Gaussian curve, the FWHM was taken as the size parameter, for all others the diameter.

A size definition for SMI microscopy

In the light of the above considerations, it might seem logical to express SMI size estimates in terms of the standard deviation, indeed Christian Wagner strongly advocated this. This approach, however, has the disadvantage that it is very difficult to visualise 'an object having a standard deviation of σ', particularly when not familiar with the technique[a]. The sizes are in any case still dependant on the assumed model. A compromise is to continue using diameter or FWHM and to proceed this with a somewhat awkward phrase such as *'if the object were a sphere/Gaussian it would have a diameter/FWHM of ...'*. As the reference objects are spherical, a spherical object model is typically used and all size estimates in this thesis are *sphere equivalent sizes*.

---

[a]This could, somewhat perversely, be seen as an advantage as it would force users to devote a little more time to understanding the technique before trying to interpret the data.

In order to understand the shape dependant effects it is instructive to consider the Fourier transform of the objects, as shown in Fig. 3.7. On closer inspection one sees that the calibration curves have the same form as the absolute value of the Fourier transform. When we consider that the process of SMI measurement is essentially equivalent to measuring the ratio of the spectral content at the frequency of the modulation to that at zero frequency, this is no great surprise. It does however allow us to apply our familiar wisdom regarding the properties of Fourier transforms to the prediction of SMI behaviour on unknown objects.
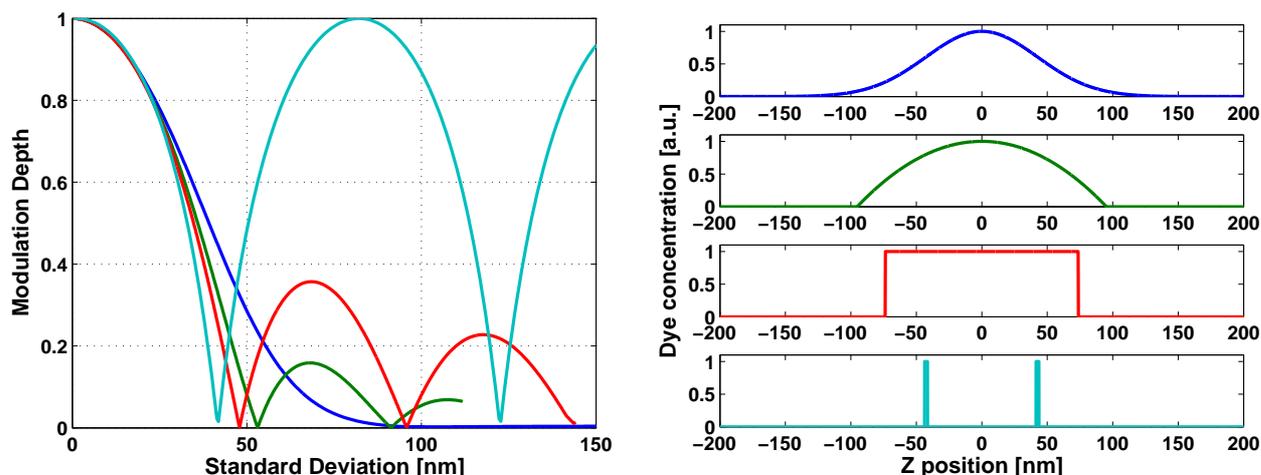
Figure 3.6: Left: Modulation depth curves for the same objects as in Fig. 3.5, with standard deviation on the x-axis instead of an arbitrary size parameter. **Right:** Object forms scaled to have the same standard deviation.

Figure 3.7: Fourier transform of the objects in Fig. 3.5. Note that the modulation depth curves have the same form as the absolute value of the Fourier transform of the objects.
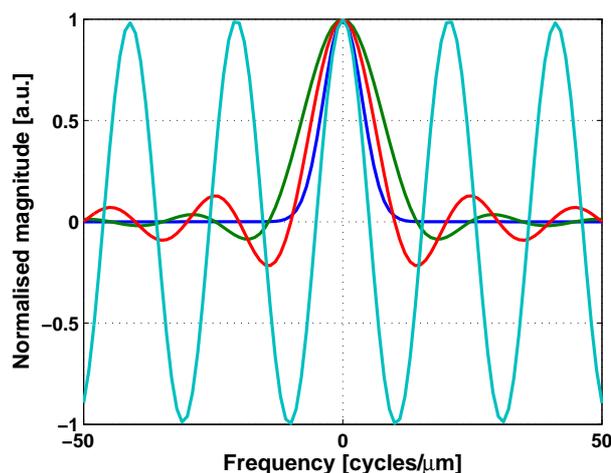
When can we use SMI nanosizing?

It should be clear by now that an absolute size estimate is impossible without detailed knowledge of the object form. A scaling relationship is however maintained for several object forms. Looking at the modulation depth curves, we can say that a usable scaling relationship is maintained over that part of the curve which is monotonically decreasing and yet still higher than any subsequent peaks. In general the object should be compact, isolated, and relatively smooth. One can see that for periodic, quasi-periodic, and sparse objects there is virtually no chance that a scaling relationship will be maintained. Unfortunately polymer random walks (e.g. DNA) are both sparse and quasi-periodic. For more details see Chapter 12.
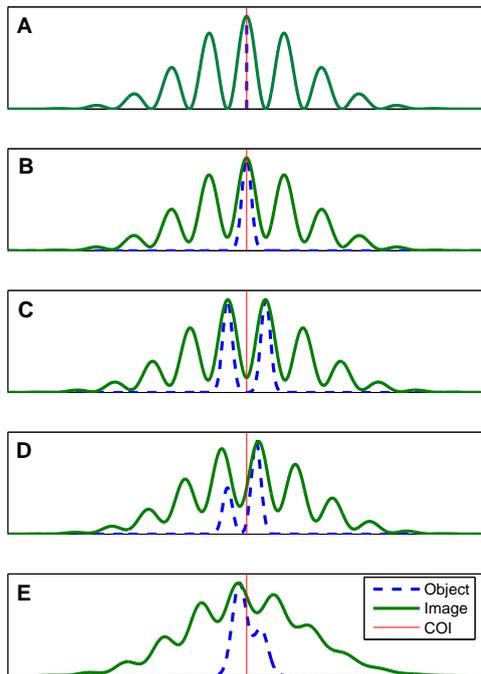
Figure 3.8: Schematic axial profiles showing some of the possible effects of object form on z position measurements. When imaging a delta function (**A**), the peak of the corresponding profile corresponds with the position of the object. The same is true for the Gaussian in **B**. For the pathological case of two Gaussians separated by half the wavelength (**C**) the peak positions have very little to do with the object centre. Other, more probable distributions (**D** and **E**) show that in the general case the peak positions do not correspond to the object position. On closer inspection one can also see that they do not exactly match the position of the peak of the object either.

3.6.2 On Distance Measurements

Whilst significant effort has been expended on the influence of object form on size measurements, analysis of its influence on distance measurements is restricted to a few brief (and rather theoretical) comments as part of the analysis of asymmetrical dye distributions ([104], pp 28-30). From this it is clear that the measured position of the fringe maxima does not necessarily coincide with the object centre of intensity. This topic obviously deserves a little more consideration if we are to determine what kinds of structures can benefit from the potential improvement in distance measurement accuracy.

To better understand the problem, it is useful to consider the examples depicted in Fig. 3.8. The case where the position of the fringes is furthest away from the true position (**C**) is similar to the type of structures which create problems for size measurements. Once again it is helpful to consider the Fourier transform of the objects. If the position of the fringes is to represent that of the object, the phase of the Fourier transform at the modulation frequency must be equal to that at the origin. For a real, symmetric object the Fourier transform is also real. This means that the phase is either 0 or pi and that sensible distance measurements can be made for object sizes up to the first zero crossing of the Fourier transform, more or less the same region in which size measurements can be made. For asymmetric objects the case is a little more difficult and depends on both the available photon statistics and the precise object form - one must determine the point where

the bias in fringe position measurements is larger than the accuracy achievable through a conventional centre of intensity measurement. It seems likely that, assuming one can rule out a periodic or quasi-periodic structure, distance measurements on arbitrary objects will be sensible if these objects show enough modulation to give a size estimate less than around 100nm.

3.7   Calibration

3.7.1   Object Size

Christian Wagner, in his doctoral thesis [104], introduced the imbalance of intensity between the two interferometer arms as a systematic error requiring calibration. An intensity difference will lead to incomplete interference and thus a decrease in the measured modulation depth. Measurements on beads of known size allow this change in modulation depth to be calculated and corrected for. All quantitative size measurements are thus accompanied with bead measurements from which the unmodulated fraction of the illumination light is determined (`umod` in the analysis software). In light of recent considerations it seems unlikely that the `umod` values previously obtained could be derived from an intensity difference alone - for more information see section 9.5.5.

3.7.2   Distances /Positions

In order to make accurate position measurements it is necessary to determine the tilt of both wavefronts and focal plane to the object slide (or another suitable reference plane). For multi-colour measurements the tilt of the wavefronts is likely to be different for the different excitation wavelengths and one must additionally measure the chromatic shift and phase relationships. The calibration of distance measurements requires the preparation of a relatively dense layer of multi-coloured beads or similar objects either on top of or below the object to be measured. For more details see Chapter 10. Due to the many variables and the non-trivial preparation requirements, fully corrected multi-colour distance measurements in biological objects have yet to be performed. The closest approach was by Susanne Fenz in her diplom thesis [32].

Chapter 4

4Pi Microscopy

4Pi microscopy is a form of far field confocal fluorescence microscopy which uses interference of the excitation and/or detection light to result in an increase in the effective acceptance angle and hence numeric aperture of the system. The technique is well described in [50, 51] and several subsequent works and the discussion here will focus on the commercial variant available in our lab. Despite the substantial increase in the effective detection aperture it is nonetheless not possible to cover the whole 4Pi steradians of solid angle, resulting in the sidelobes (or ghost images) typical of 4Pi images[1].

4.1 Optical Setup and Operating Principle

A schematic layout of the interferometer section of the 4Pi microscope is shown in Fig. 4.1. In addition to the components shown in the schematic, the second objective lens is mounted on a 3-axis piezo in order to allow the two foci to be accurately overlayed. **PZ1** is also in practice a 3-axis (piston and $2\times$ tilt) piezo allowing the relative tilt of the wavefronts to be adjusted as well as the phase. A somewhat coarser control of path length and tilt can be performed with thumbscrews behind **M2**. Should one wish to bring the detected light to interference (4Pi-B & C), components for dispersion compensation are also necessary.

The coherent overlay of the two foci leads to the illumination pattern shown in Fig. 4.2-A. Even with confocal detection, the resulting OTF does not deliver enough support to avoid ambiguities such as are found with SWM. Bringing both the detected and emitted light to interference (4Pi-C, Fig. 4.2-B) improves the situation, but needs a large stokes shift to allow unambiguous deconvolution. A robust OTF support is only possible by shifting to 2-photon excitation (Fig. 4.2-C) - the quadratic dependence on excitation intensity results in lower sidelobes and the longer wavelength moves them further away from the centre of the detection PSF, further decreasing their effective height. This is equivalent to shifting the regions of additional support in the OTF closer to the centre, such that there is now an

---

[1]Note that true 4Pi excitation or detection would, as the system is still bandlimited, show a small amount of ripple as observed laterally in the Airy disk. Localised axial ghost images, or sidelobes, would however not be present.
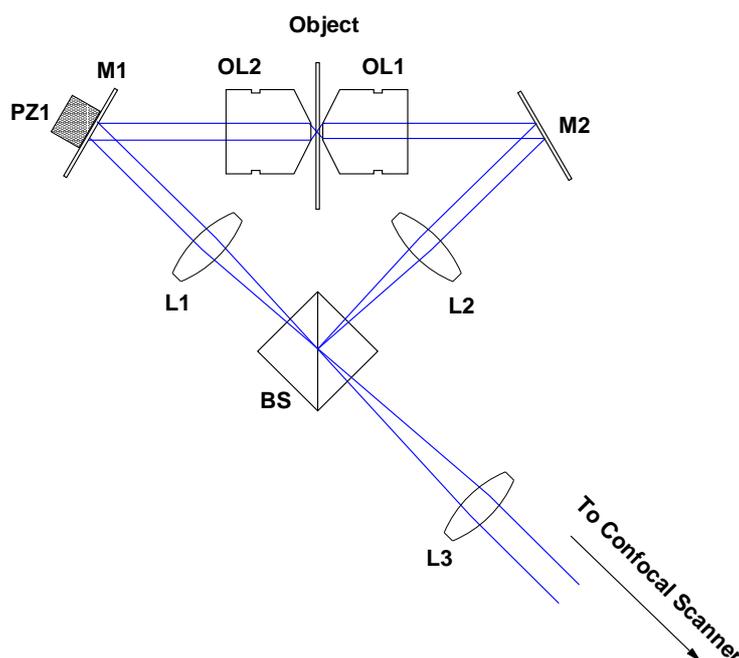
Figure 4.1: Optical layout of the 4Pi microscope. Illumination light comes from a conventional confocal scanner, and is split with a 50/50 beamsplitter **BS** and focussed into the sample by the two opposing objectives **OL1** and **OL2** where it interferes, creating the characteristic 4Pi illumination PSF. The emitted fluorescence follows the same path back to the scanner and is detected as per confocal microscopy. The piezo **PZ1** allows adjustment of the relative path length in the two arms and hence the phase of the interference. The relay lenses, **L1**, **L2** and **L3** compensate for the extra distance between the objective and the rest of the imaging optics.

	Excitation	Detection
4Pi-A	interfering	confocal
4Pi-B	confocal	interfering
4Pi-C	interfering	interfering

Table 4.1: 4Pi Variants, note that all variants are possible with either 1 or 2 photon excitation and that type B is practically never used.

overlap. This can be further improved by letting the detection interfere, giving 2-photon 4Pi-C. For a summary of 4Pi variants see Table 4.1.

4.2   Image Processing

4.2.1   Preprocessing

To increase sensitivity the commercial 4Pi uses avalanche photodiodes (APDs) for detection. This has the disadvantage that the APDs have a very low ($\sim$3bit) dynamic range.
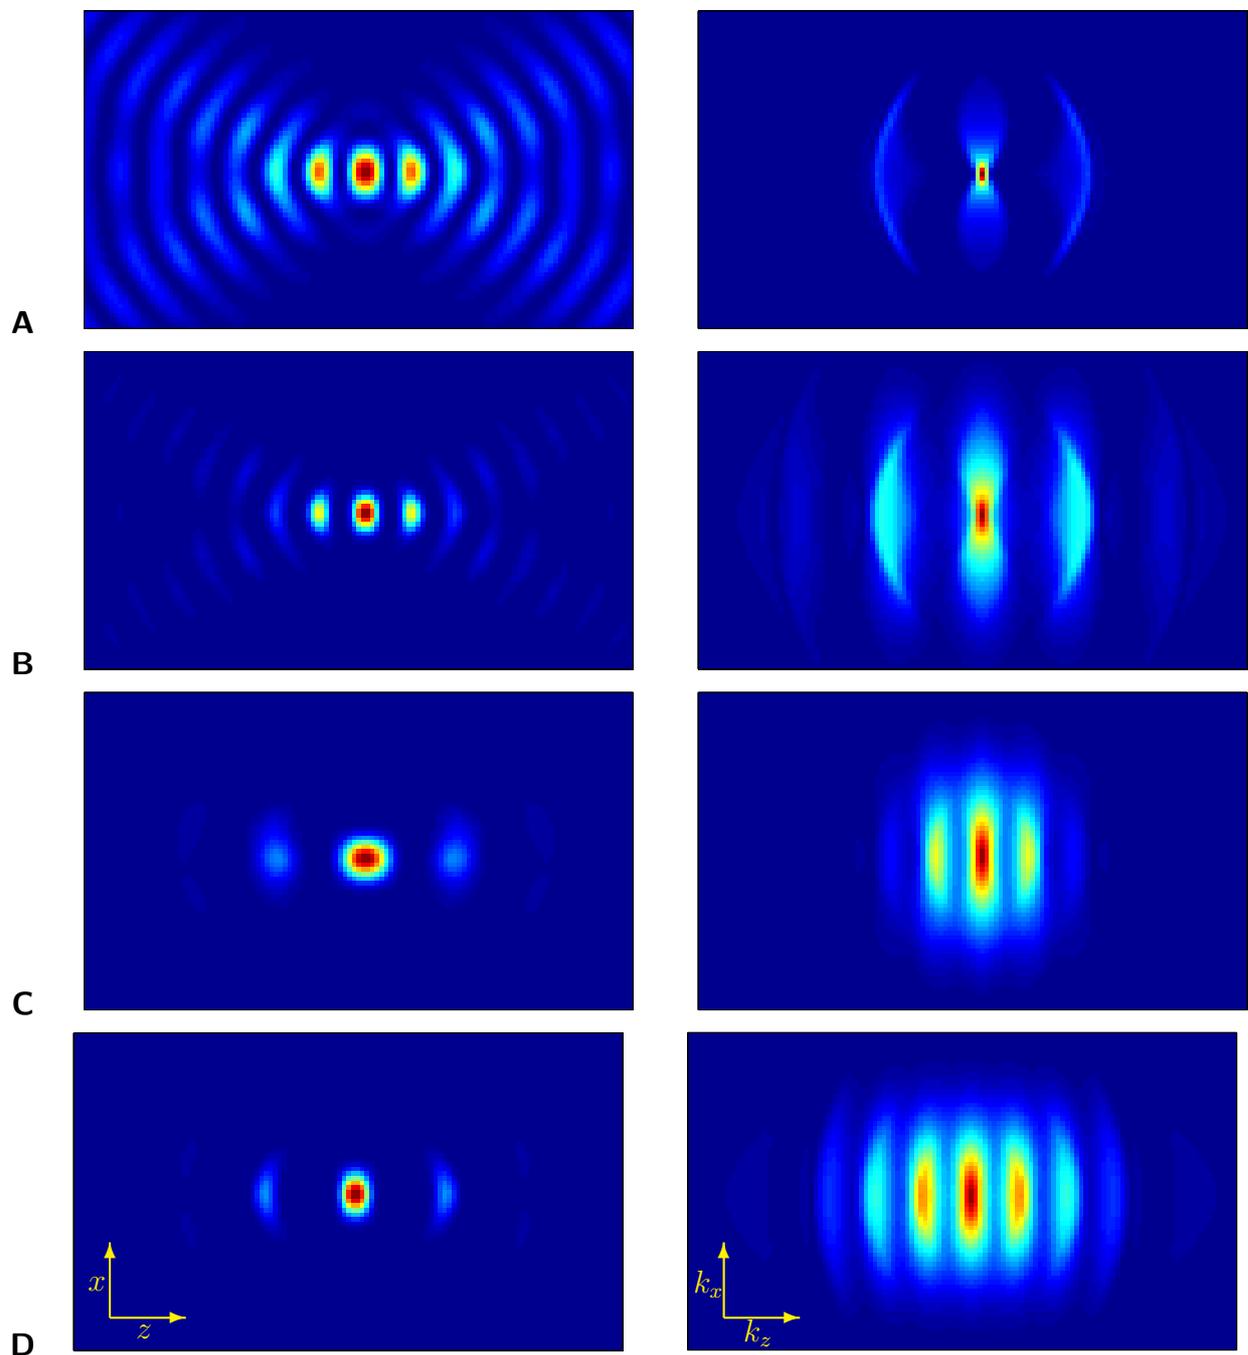
Figure 4.2: 4Pi PSFs (left) and OTFs (right). **A:** 1-photon 4Pi excitation (NB: no detection component), **B:** 1-photon 4Pi-C, **C:** 2-photon 4Pi-A, **D:** 2-photon 4Pi-C. Note PSFs are only approximate (see note in section 2.2.1).

Figure 4.3: 5 point inverse filter. The 4Pi PSF is approximated as 5 points along the optical axis. The sidelobes of first order are 40% of the central peak, the 2nd order sidelobes 5%. Inverting a 15×15 matrix with the coefficients on the major diagonals results in the inverse filter shown. Note that the magnitude of the terms in the inverse filter falls quickly with increasing distance from the central peak.
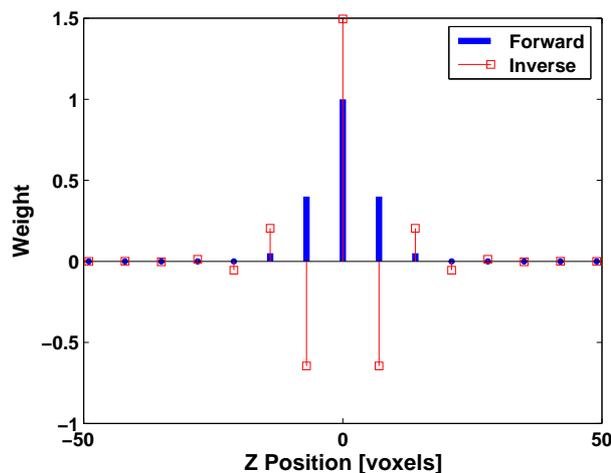
It is thus normal to significantly oversample the images during acquisition and to blur prior to data evaluation to recover an image with usable dynamic range and signal to noise characteristics. This blurring is typically performed with a Gaussian filter kernel. Other choices of filter (eg. Butterworth) are possible and have the potential to allow a sharper cutoff and hence less unwanted bluring[2]. Preliminary experiments with a Butterworth filter give good results, albeit with some ringing not present when using a Gaussian.

4.2.2   Sidelobe Removal

As mentioned earlier, 4Pi images are characterised by sidelobes, or ghost images on either side of the true structure. In order to obtain an unambiguous image, these sidelobes must be mathematically removed [87].

3 or 5 Point Inverse Filter

The three or five point inverse filters rely on the assumption that the PSF is separable in axial and lateral responses, and that the axial response consists of scaled copies of the central response. This allows one to express sidelobe formation as the weighted sum of either 3 or 5 shifted copies of the image, depending on whether one considers the 2nd order sidelobes or not. Coefficients for an inverse filter can then be determined by inverting a matrix containing these coefficients as diagonals.

Example 5-point inverse filters are shown in Figures 4.3, 4.4 and 4.5. In order to accomplish successful sidelobe removal without excessive artifact generation, the length of the filter must be sufficient such that the terms which get cut off are insignificantly small. Unfortunately the Leica 4Pi software does not allow the length of the inverse filter to be specified and deconvolution of simulated data produces significant artifacts[3], suggesting

---

[2]A Gaussian filter has a very slow transition from passband to stopband, making it hard to chose a kernel size which delivers good attenuation of out of band information without effecting the desired signal.

[3]experiments performed at the 4Pi workshop in Göttingen, 18-22 Sept. 2006.
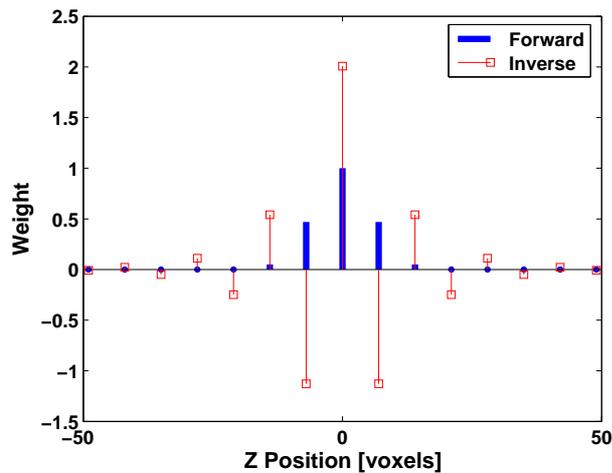
Figure 4.4: 5 point inverse filter. The 4Pi PSF is approximated as 5 points along the optical axis. The sidelobes of first order are 47% of the central peak, the 2nd order sidelobes 5%. Inverting a 15×15 matrix with the coefficients on the major diagonals results in the inverse filter shown. Note that the increase in the sidelobe height results in a significant increase in the weights at high shift in the inverse.
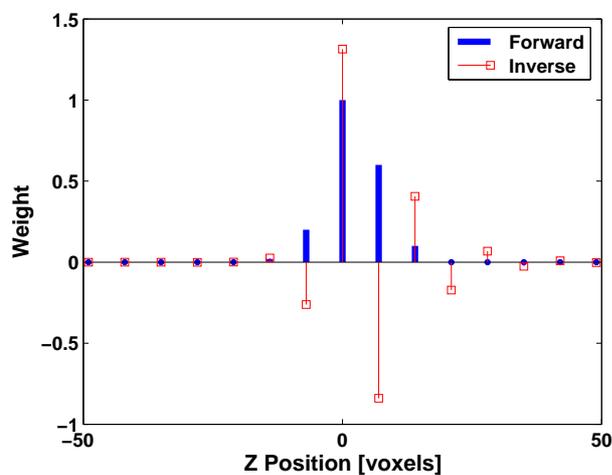
Figure 4.5: 5 point inverse filter of an asymmetric PSF. The 4Pi PSF is approximated as 5 points along the optical axis. The sidelobe weights are [0.02, 0.2, 1, 0.6, 0.1]. Inverting a 15×15 matrix with the coefficients on the major diagonals results in the inverse filter shown. Note that the magnitude of the terms in the inverse filter at high shifts is significant.

a short (probably length 5) filter. It might thus be advisable to refrain from using the internal Leica 5-point deconvolution unless the PSF is symmetric and sidelobe height very low. Figure 4.6 shows the impulse and step responses of the filter in Fig. 4.4 with an inverse filter of length 15 and 5. The issue is also discussed in [94].

Wiener Filter

One can also use a Wiener filter, the main disadvantage of which being that it probably won't completely remove the sidelobes[4], and that it is almost as much effort as a real deconvolution. It does have the advantage of being mathematically (much) more sound than the 5 point inverse filter and very much more predictable than most ML deconvolutions. For more information on Wiener filters see Chapter 6.

---

[4]Due to the fact that the Wiener filter effectively represents a weighted sum of the inverse and the original image. This could be deemed an advantage as a good tradeoff between sidelobe suppression and noise amplification will be obtained.
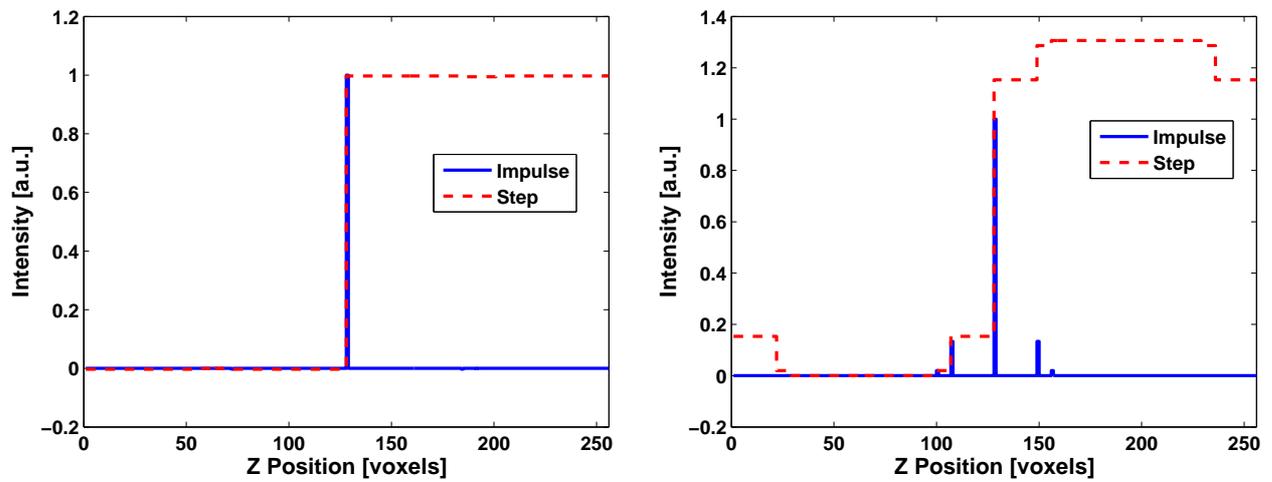
Figure 4.6: Left: Impulse and step responses for the scenario in Fig. 4.4 when using the 15 point inverse filter. **Right:** The same when the inverse filter is truncated to 5 elements.

Maximum Likelihood (ML) Deconvolution

The technique which is likely to produce the best images is a full blown maximum likelihood deconvolution. This is however much slower than the 5 point inverse filter and somewhat more susceptible to a poor PSF measurement so in practice the 5 point method is the most used technique. A more detailed description of Maximum Likelihood deconvolution is presented in Chapters 6 and 8.

Chapter 5

Biological Structures in the Cell Nucleus and their Labelling

The basic principles of biological systems investigated in this thesis are generally well described in elementary biology textbooks, and it's most unlikely that I'd be able to offer a great deal of additional insight. Indeed when it comes to microscopic imaging it's arguably sufficient to treat biological systems as abstract mathematical objects. There are however a few pitfalls, particularly when it comes to labelling, which must be considered in the design of experiments and in the analysis and interpretation of the resulting data. If the tone of this chapter seems overly negative then only because I've attempted to highlight the potential problems. When not otherwise stated the information presented here has been paraphrased from [2].

5.1   The Nucleus

The nucleus is a compartment in the cell which contains the cellular DNA along with the machinery required for transcription, replication and repair. The nucleus also contains lots of other things, not all of which are directly associated with the processing of the genetic information (eg. proteolysis foci). It is however generally accepted that the principle function of the nucleus is the coordination of the genetic information.

Various models have been proposed for the internal structure of the nucleus(e.g. [23, 17, 74]), with earlier models proposing a more or less random distribution of DNA within the nucleus and more recent models indicating a much higher level of organisation. At present there is still no real consensus on a model for chromatin organisation. Rather than attempting to introduce one of the models I thus prefer to highlight a few facts that are known and well accepted (NB this is by no means a complete list) [73].

- The DNA is packed by looping around complexes of histones, the nucleosomes, giving the 10nm fibre. A higher level of compaction is achieved in the 30nm fibre by stacking the nucleosomes, and it is largely accepted that even higher order compaction mechanisms exist. It is however anything but clear exactly what form the chromatin

takes within the intact nucleus - recent studies point to a heterogeneous state where both 10nm and 30nm fibres are present [24].

- There are sub-compartments within the nucleus (eg. the nucleolus, PML bodies, splicing speckles etc...).

- The distribution of DNA within the nucleus is not uniform, with more (heterochromatic) and less dense (euchromatic) regions. The heterochromatic regions tend to be found towards the nuclear periphery and tend to be transcriptionally inactive.

- The individual chromosomes remain separated, forming chromosome territories. The relative location of these territories seems to depend on cell type and can change during differentiation. The interior structure of the chromosome territories and the degree of overlap (if any) between neighbouring CTs remains a topic of speculation [17].

- Certain chromatin regions show an affinity to other structures within the nucleus (e.g. telomeres often reside at the nuclear periphery; proteins forming a link between the telomeres and Lamin have been identified).

- The structure of the nucleus is not static - significant dynamic interactions have been observed [64].

- Proteins which 'do stuff' with DNA are often found in clumps, or 'foci'. This is commonly linked to a 'functional relationship', a conclusion which is not without its hazards (there are physical mechanisms, eg. entropic crowding, which could explain foci formation without the need for any functional relationships [69, 68].

5.1.1 Gene Loci

A gene is the segment of DNA which codes for a specific protein. When referring to a gene locus, one usually means the 3D object comprising the DNA and all associated proteins. It is thought that changes in this structure are likely to be required in order to allow access by the transcription machinery and that one might be able to detect such changes through the use of microscopy and correlate the structure of a gene locus to its regulatory state. In vivo experiments with a large artificial plasmid injected into cells and labelled using the Lac repressor system have shown considerable decondensations under the influence of a bacterial promoter [100]. To the best of my knowledge, firm evidence for decondensation of either individual gene loci or, in fact, any endogenous chromatin have yet to be shown. This is probably, at least in part, due to the difficulty of finding appropriate in vivo labelling strategies for endogenous DNA.

5.2 Labelling Techniques

Before one can measure a biological object it must be labelled. Several labelling strategies are possible, some of the more common methods are outlined below. They deserve mention here as with the current labelling methods it is dangerous to assume that the distribution fluorescent labelling is a good representation of that of the underlying object, and a consideration of the possible effects of labelling is thus essential for meaningful interpretation of microscopic data. Most labelling techniques involve adding the dye to the sample, letting it bind and then washing the unbound excess away. Some labelling techniques are particularly sensitive to variations in the washing procedures resulting in protocols which are, to the un-schooled physicist, the source of much scepticism. A good reference for labelling techniques is the Molecular Probes handbook [41].

5.2.1 DAPI, Phalloidin, Lipophile Dyes etc...

One of the most direct labelling techniques is to use a fluorescent chemical compound with an affinity to a particular structure. Examples are DAPI where the dye molecule binds preferentially to AT rich sequences in the minor groove of the DNA double helix [98]. Other DNA dyes such as Hoechst work on a similar principle. Another example of direct affinity labelling is lipophile dyes (e.g. DiI), whose chemical properties make them more soluble in lipids than in aqueous solution leading to a selective labelling of membrane structures. Another affinity label is Phalloidin [21]. Phalloidin labelling is slightly different from the above methods in which the dye molecule itself possesses a specificity in that a fluorophore is conjugated to the previously non-fluorescent phalloidin toxin which has a strong affinity for F-actin. Unfortunately specifically interacting chemical compounds have only been identified for a few of the many structures which one would like to visualise, and more flexible techniques are needed.

5.2.2 Immunolabelling

Another approach to labelling is to (ab)use the mechanism used by our immune system to recognise foreign proteins. Antibodies against the protein of interest can be created, and these can in turn be conjugated with one or more fluorophores. In order to increase both the signal strength and the flexibility of the system, a two-layer system is typically used in which unlabelled primary antibodies from one species (eg. mouse-anti-*target*) are visualised with a more generic, fluorescently labelled, secondary antibody from a different species (eg. goat-anti-mouse-FITC). Three layer protocols are also possible. The increase in signal obtained when using multiple layer antibody protocols comes from the fact that there are often several *epitopes* on the primary antibody, allowing several secondary antibodies to bind.

There are several properties of antibody labelling which must be considered when using them for high resolution microscopy. The most obvious of these is size. As shown in Fig. 5.1, the size of an antibody, at around 12nm when fully extended, is considerable. When
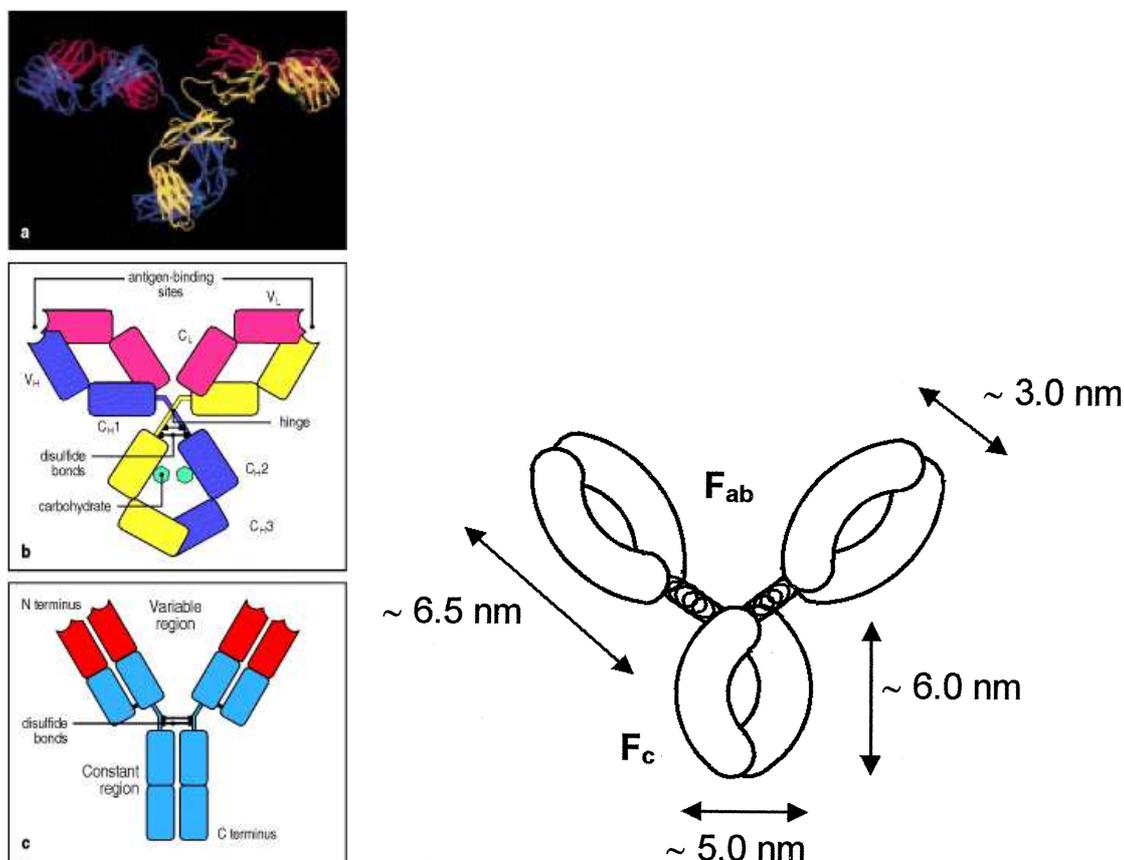
Figure 5.1: Crystal structure of human antibody IgG, along with schematic diagrams illustrating functional structure and dimensions. Images: [57, 83]

using two-layer protocols, the fluorophore can thus be up to 24nm away from the target. Another is the fact that each antibody has at least two[1] antigen binding sites, which can lead to cross-linking and alterations in the structure being measured. A partial solution is to use so called *Fab* fragments [] which are directly labelled with fluorophore(s). Fab fragments are obtained by cleaving the antibody at the hinge and disposing of the the Fc constant region. They are thus much smaller and have only one antigen binding site.

Another problem is that several antibodies also bind to structures other than those of interest. This unspecific binding is typically addressed by adding milk powder[2] to the cells before labelling, with results which then depend on the quantity of milk powder and the duration over which it is added.
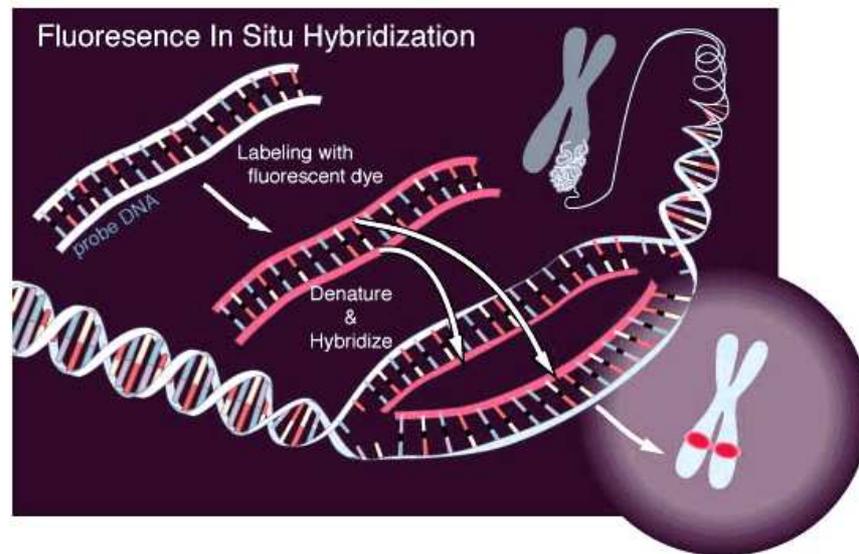
Figure 5.2: Schematic illustration of FISH. (*Source: http://en.wikipedia.org/wiki/Fluorescent_in_situ_hybridization*)

5.2.3 Conventional FISH

Fluorescence In Situ Hybridisation or FISH [18] is a technique for the labelling of specific DNA sequences. DNA matching the sequence of interest is obtained[3], labelled, cut into shorter segments, and hybridised with the target sequence. When hybridising to double stranded DNA, this must first be denatured using either heat or chaotropic agents, which is possibly the most important disadvantage of this method. It is not certain what the effect of this denaturation is going to have on higher order chromatin structure and whether it is necessary to (at least temporarily) remove the histones in order for the probe DNA to hybridise. Many people studying FISH labelled structures maintain that the hybridised DNA reassembles around the histones and reconstitutes the structure, a view which I find difficult to accept. There are two ways in which FISH probes can be labelled:

Direct

When using direct labelling, fluorophores are directly attached to the DNA probe, either through the inclusion of a modified base during PCR or through techniques such as nick translation.

---

[1]IgG has two

[2]or fetal calf serum, both are concentrated sources of protein, and the idea is that these proteins will compete with the antibodies for binding to the unspecific epitopes.

[3]normally as a Bacterial Artificial Cromosome (BAC) clone amplified using PCR, but potentially also using chemical synthesis.

Indirect

It is also possible to introduce a substance which can be recognised in a follow up reaction, either by antibodies or through other affinity relationships such as that between biotin and (strept)avidin. This has the effect of combining the disadvantages of both FISH and immunofluorescence and typically results in a much higher level of background and unspecific binding. It is however usually possible to obtain a somewhat higher signal level, and indirect labelling is currently the easiest way to perform FISH with quantum dots or for electron microscopy.

5.2.4   COMBO-FISH

COMBO-FISH [42] is a novel variant of oligo-nucleotide FISH in which short sequences which cluster at the locus of interest are selected from the genome databank and chemically synthesised. It has the advantage of being able to label essentially arbitrary locations[4] and of allowing probe sizes which are significantly smaller than with conventional FISH[5]. The COMBO-FISH technique, as practiced in the Cremer lab during my thesis, still shows several teething problems such as poor specificity and a rather miserable binding efficiency (see 5.3.4) which have been somewhat of a limiting factor in its application.

Once the above problems have been overcome COMBO-FISH has considerable promise as a diagnostic tool, and for approximate gene position measurements. Due to the relatively low labelling density intrinsic to the method (around 30 oligos, each of ∼15bp randomly distributed over a much larger - hundreds of Kbp - stretch of DNA) it is more than questionable as to whether, even with 100% binding, the labelling density would be sufficient to allow any structural (eg size) measurements.

5.2.5   Fluorescent Proteins

Several naturally occurring proteins are fluorescent, one of the first and the most famous of which is GFP, derived from the jellyfish *Aequorea victoria* [99], although fluorescent proteins covering the whole visible spectrum are now available [88]. One of the most elegant ways of labelling proteins of interest is to create a gene in which the sequence for the target protein has been fused to a fluorescent protein. This *fusion protein* can then be expressed in the target cells, resulting in a specific labelling of said protein. As the labelling is produced by the cell itself, the technique is applicable both in-vivo and after fixation. There is also no need for washing, removing the temptation to wash until the expected structure is obtained.

---

[4]Technically BAC probes can do this too, as the requirements for being able to make a BAC probe are the same as for having been able to sequence the DNA in the first place.

[5]A much touted advantage, which with several of the probes used within the group during my thesis is of questionable relevance due to the fact that although only some hundreds of base pairs are labelled, these are randomly spread over a much larger region (eg. 180kbp).

There are however a few questions which must be considered when using XFP labelling - the most important of these are whether the fusion proteins retain their original function and localisation, and the level of the protein present in the cell. These problems are, to some extent, a result of the most common and easiest way of getting the protein expressed and visible (using a viral transfection with a copy[6] of the fusion protein gene under the control of a strong promotor, often of bacterial origin). With such a system it is quite possible that one produces a significant excess of the protein which can be toxic and/or lead to signal in locations where the protein would not normally be present. The problem of protein over-expression can be addressed by placing the fusion protein under the control of its endogenous promotor, potentially combined with stable integration and a knockout of the endogenous gene. The maintenance of function is generally tested by seeing whether the fusion protein can restore viability in a cell with an endogenous knockout. From a more physical perspective, XFP proteins are considerably larger than organic fluorophores (although not quite as large as antibodies) and some fluorescent proteins have a tendency to form dimers (or higher n-mers, commonly tetramers).

5.2.6 Other

There are several other labelling methods in common use, a couple deserving mention follow:

Lac and Tet: are two closely related techniques for the in-vivo labelling of DNA. Lac and Tet operons are bacterial DNA sequences to which the bacterial repressor proteins Lac or Tet bind. A labelling is accomplished by inserting these sequences near the gene to be observed and by expressing XFP-Lac or XFP-Tet fusion proteins. These proteins bind to their respective operons, making them visible [10].

PNAs: As an alternative to DNA based oligonucleotides it is possible to synthesise oligos with a peptide instead of a sugar-phosphate backbone. These Peptide Nucleic Acids have the advantage of having a higher affinity to DNA/RNA than DNA or RNA based probes, and a slower degradation when used in living cells.

Biotin and (Strept)avidin: The proteins avidin and streptavidin show a high affinity for the vitamin biotin. This affinity reaction can be used in an indirect labelling strategy where one of the pair is attached to the target and the other fluorescently labelled.

5.3 Labelling and Sampling

Whilst most people consider sampling theory when acquiring images, and choose voxel sizes small enough to meet the Nyquist criterion for the optical system being used, it is

---

[6]possibly several copies

Figure 5.3: Gaussian and reconstruction from samples after sampling at $f_s = FWHM/2.35$. Note that whilst a Gaussian technically has infinite support, the effect of aliasing is almost negligible. A Gaussian is admittedly the best-case scenario, but it seems likely that many biological objects will also show limited aliasing related effects.
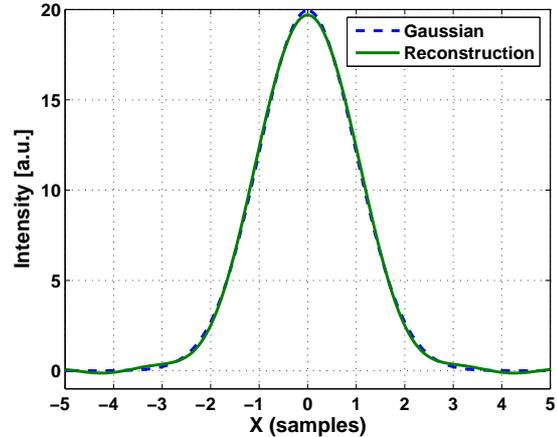
largely disregarded when it comes to designing labelling strategies, or when considering whether it is possible to extract structural information from microscopic images. It is thus worthwhile considering the effect of the labelling on our measurements of structure, and to see if we can apply any of the concepts of sampling theory to labelling.

The sampling theorem [89], commonly called the Nyquist or Nyquist-Shannon theorem, states that a band limited signal can be accurately reconstructed if and only if it is sampled with a frequency (f_s) greater than twice the highest frequency component in the signal (f_c):

$$f_s > 2f_c \tag{5.1}$$

5.3.1 Sampling theory applies to band-limited signals

Arbitrary biological objects are generally not band limited with the consequence that even for an arbitrarily high labelling density it is unlikely that the sampling theorem can be satisfied in a rigorous sense. Luckily most biological objects are isolated, positive and non-periodic implying a decline in the Power Spectral Density (PSD) towards higher spatial frequencies. This should limit the effect of aliasing, should the object be (necessarily) undersampled. A reasonable interpretation of sampling theory could thus be to assume that aliasing effects are small and require that the sampling was sufficient for the measurement accuracy[7] which one wishes to achieve. Care should however be taken when the object is quasi-periodic as is the case with membrane bilayers and similar structures.

5.3.2 Labelling takes discrete values and is not uniformly spaced

Sampling theory assumes that the samples are evenly spaced and can take arbitrary values. In labelling, the labels are not equally spaced and can only take two values (there or not there), for a comparison see Fig 5.4. The accuracy of the representation with such sampling

---

[7]Note that in cases where sub-resolution measurements are being made, this is NOT the imaging resolution.

depends on the density of the sampling (Fig 5.5). One way to get around this problem is to say that the number of labels within a volume element, the size of which corresponds to Nyquist sampling of the structure, should be high enough that two such elements having the same underlying density should contain approximately the same number of fluorophores.
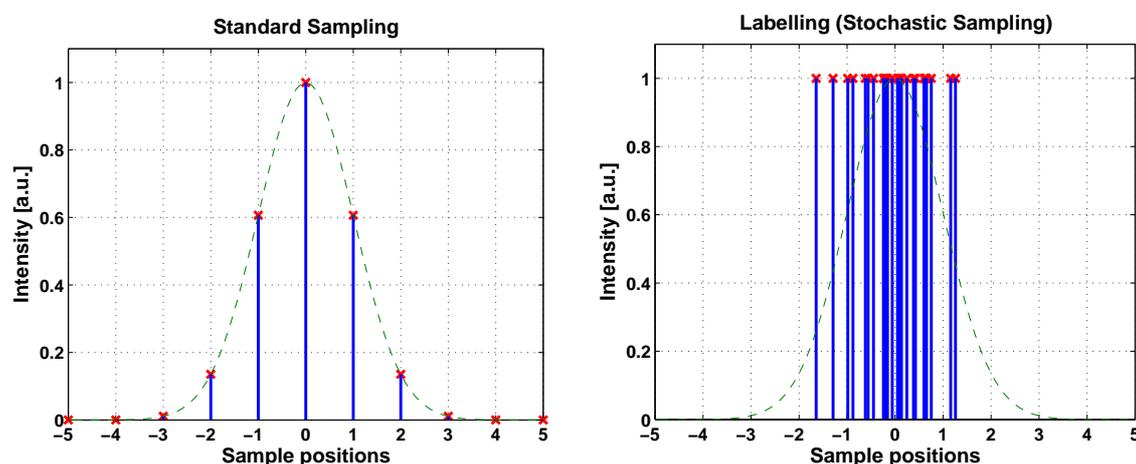
Figure 5.4: Comparison of sampling as described by sampling theory and what one obtains with labelling.

Note on sampling simulations

For *standard* sampling the samples were the value of the function at the evenly spaced sampling points, with the reconstruction being the sum of sinc functions at each of the sample points, weighted with the value at that point (as required by sampling theory). For the stochastic labelling, label positions were randomly chosen using a Monte-Carlo process in which the object function determined the likelihood of placing a label at a given position. In this case, the reconstructions were obtained by taking the sum of sinc functions of constant amplitude, located at each of the label positions.

Using a sinc^2 which, whilst mathematically not correct in light of the requirements of sampling theory, is somewhat closer to the imaging process gives qualitatively similar results.

5.3.3   Rule of thumb

A rigourous criterion based on the above considerations would more or less require an infinite labelling density. From the simulation results shown in Fig 5.5 it appears that useful, albeit somewhat error prone, information can be obtained once the labelling density exceeds approximately 10 times the sampling rate that would be required for Nyquist sampling of the object concerned. For example, a measurement accuracy of 50nm would require a labelling density of at least 10 fluorophores per 25nm^3 volume, or 0.00064 per nm^3.
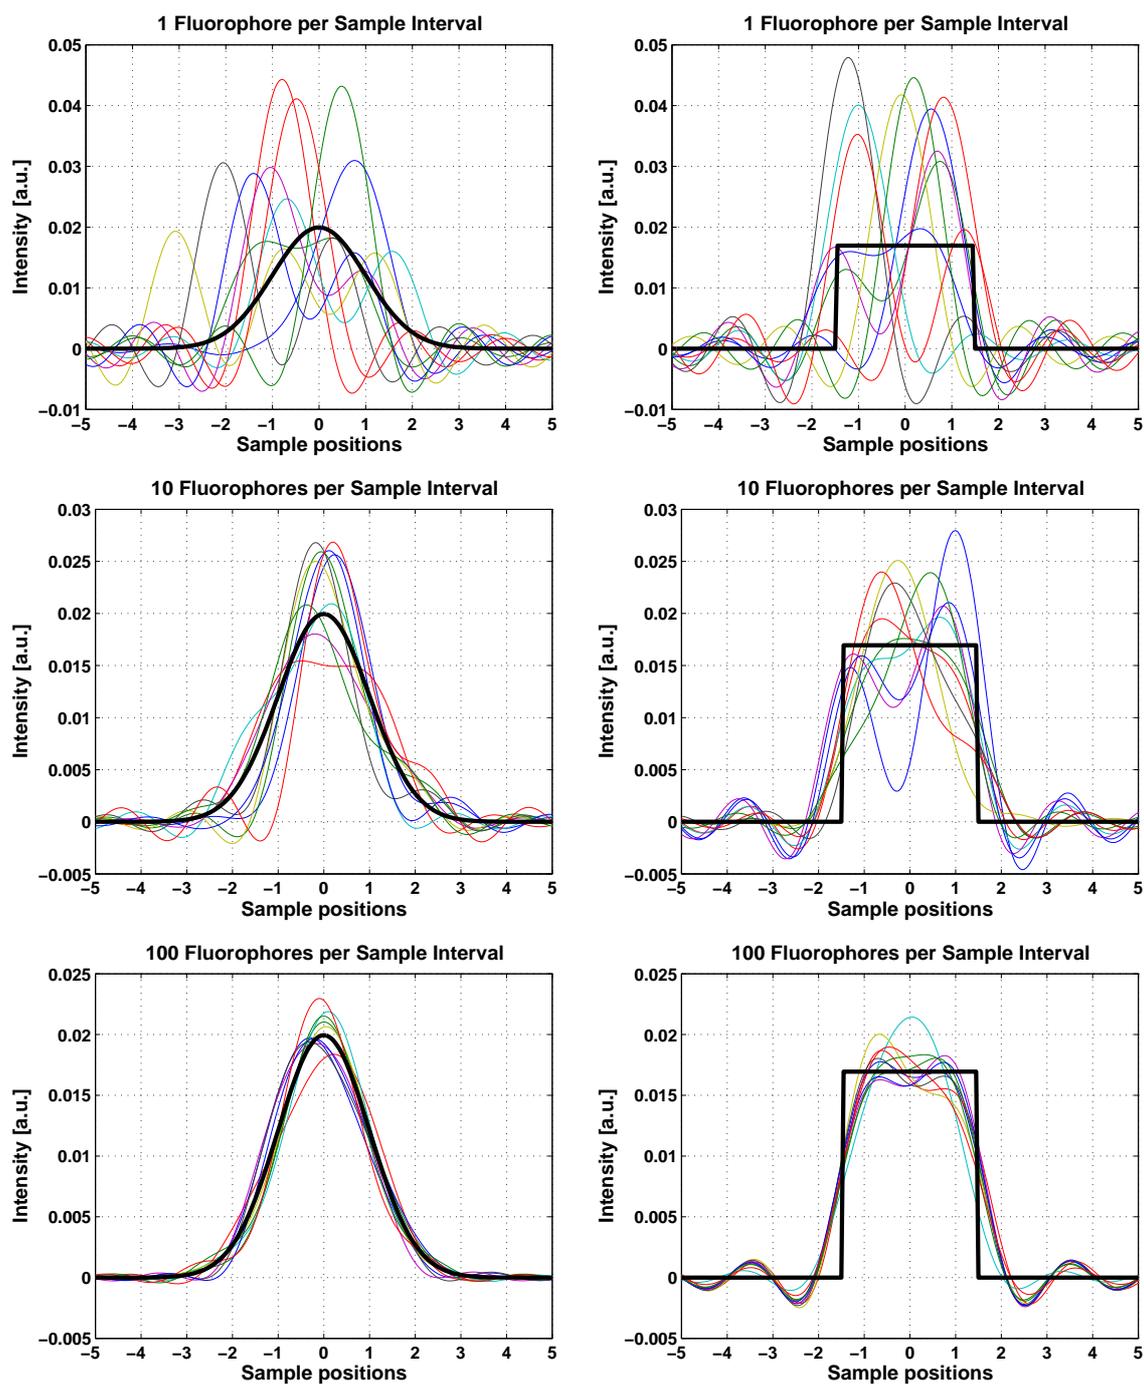
Figure 5.5: Effect of labelling density on object reconstruction. For the Gaussians, the given label density is for one of the central samples.

This has rather disturbing consequences when one considers combining labelling techniques such as COMBO-FISH with high precision structural measurements such as might

be obtained using SMI microscopy. In order to be able to distinguish two populations from SMI measurements, a measurement accuracy of better than around 30-50nm is essential. Assuming a 30nm fibre, the volume of a 180kbp gene segment[8] is $\pi \times 15^2 \times 10 \times 180 \approx 1.3 \times 10^6 \text{nm}^3$. Sampling for 50nm accuracy would thus require at least ~ 1000 fluorophores along the length of the segment. Current COMBO-FISH probes have a theoretical maximum of around 30 labelling sites if all oligos were to bind, with a poor binding efficiency reducing this considerably.

5.3.4 Labelling Efficiency

Most labelling reactions are a chemical equilibrium, in which not all sites will be labelled. When the equilibrium lies such that only a few of the sites are labelled, a corresponding decrease in the effective sampling rate will be obtained. In all cases, a poor labelling efficiency will be detrimental to the reliability of the method, particularly when a significant amount of unspecific labelling is also present. It is thus advisable to analyse the labelling efficiency before any further attempt at data analysis. For FISH studies, some information about the labelling efficiency can be obtained by counting the number of spots observed within each cell.

5.4 Fixation and Embedding

In order to preserve the structure of dead cells through the labelling and observation steps, it is necessary to use a chemical fixative. These usually work by either chemically cross-linking or precipitating proteins/DNA etc. in the sample. The most commonly used fixatives for light microscopy are methanol and acetic acid (MAA) or paraformaldehyde (PFA). There is much debate as to the relative merits of the two procedures, with various groups uncompromisingly advocating one method and discrediting the other. What is clear is that both methods involve compromises regarding structural preservation and also with respect to other aspects of the observation procedure such as labelling efficiency, and that the relative merits depend strongly on the structure being observed. In the end one must try and choose the fixation method most appropriate to the question being asked, potentially trying and comparing both methods. For more information see [8].

Another aspect of biological sample preparation which is possibly not given the amount of attention that it deserves is the embedding of the samples. Refractive index mismatch can result in a significant worsening in the optical characteristics of a microscope system [48] and the sample preparation techniques used in routine work rarely provide more than an approximate optical match (eg. glycerol embedding for use with oil objectives). This can be improved by a careful choice of embedding medium and the use of several steps to ensure the medium diffuses completely into the sample [93]. Such careful control of the sample embedding is essential for 4Pi microscopy and highly advantageous for other high resolution techniques, even confocal microscopy.

---

[8]as used by Connie [35]

Chapter 6

Parameter Estimation and Deconvolution

The source for much of the information in this chapter is the lecture notes from an excellent course on *Inverse Problems* given by Sze Tan and Colin Fox. The course notes are available online at *http://www.math.auckland.ac.nz/~phy707/* and I would strongly recommend them to anyone looking for a deeper mathematical understanding of deconvolution and related processes.

Both deconvolution and the fitting of models to experimental data fall into the field of *inverse problems*, and there is thus a high degree of similarity in the underlying mathematics. In general a functional, A where $d = A(f)$, is assumed which maps the parameters, f, of an underlying system/object onto the data d. In the case of model fitting A is the model function and, in the case of deconvolution A is a description of the imaging process and f the underlying object. If the mapping is linear, the vector notation

$$\mathbf{d} = A\mathbf{f}$$

can be used.

An inverse problem, ie. estimating $\mathbf{f}$ from $\mathbf{d}$, is said to be *well posed* if the following three conditions are true [97]:

1. A solution exists for every $\mathbf{d}$ in the data space.

2. This solution is unique in the image space.

3. The inverse mapping $\mathbf{d} \rightarrowtail \mathbf{f}$ is continuous.

Note that for a band limited mapping condition 2 is in general not fulfilled - the deconvolution of microscopic images is thus ill-posed.

Another important property of inverse problems is their conditioning, which for a discrete inverse problem, more or less translates to the number of parameters which have to estimated from a given number of data points. The concept of *well-conditioned* and *poorly-conditioned* is somewhat fuzzier than that of *well* and *poorly-posed*, but in general

one can say that the larger the number of data points in comparison to the number of parameters, the better conditioned a problem is. Poor conditioning generally results in an amplification of the influence of noise in the result.

The experimentally measured data values $\mathbf{d}$ are usually contaminated with noise, giving us $\mathbf{d} = A\mathbf{f} + \mathbf{n}$, where $\mathbf{n}$ is the contribution of additive noise. As $\mathbf{n}$ is unknown, it is not possible to retrieve $\mathbf{f}$ by simply applying the inverse mapping A^{-1} to the data. We thus try and obtain a *'mathematicians' best guess'* estimation of the solution $\hat{\mathbf{f}}$. For Gaussian distributed noise of constant amplitude it can be shown that this is when the square of the distance between the data and the forward mapping of the approximate solution is minimised (the 'least squares' solution):

$$\hat{\mathbf{f}} = \arg\min \left\| \mathbf{d} - A\hat{\mathbf{f}} \right\|^2 \tag{6.1}$$

As Poisson noise is multiplicative, not additive, the best estimate in the case of Poisson noise is obtained by comparing the ratio of the data to the forward transformed estimate:

$$\hat{\mathbf{f}} = \arg\min \sum_i \left(1 - \frac{d_i}{A_i\hat{\mathbf{f}}} \right) \frac{\partial A_i\hat{\mathbf{f}}}{\partial \hat{f}_k} \tag{6.2}$$

For cases where one has Gaussian noise of different (but known) amplitudes it is possible to modify eqn. 6.1 to include a weighting factor W:

$$\hat{\mathbf{f}} = \arg\min \left\| W \left(\mathbf{d} - A\hat{\mathbf{f}} \right) \right\|^2 \tag{6.3}$$

In this case the matrix W is a diagonal matrix where the entries on the diagonal are given by the error in each data point by $W_{ii} = 1/\sigma_i$ with σ_i being the standard deviation of data point d_i. It can be shown that, if the properties of the noise are unknown, the standard least squares method gives the best results.

6.1   Fitting

In most cases where a function is to be fitted to experimental data, the number of data points significantly exceeds the number of parameters to be estimated making the problem *well conditioned*. Fitting problems are also usually *well posed*, although it is sometimes necessary to apply constraints to ensure this[1]. The fitting problem is thus as simple as solving one of the equations 6.1, 6.2, or 6.3. For a linear operator, A, this is typically done using linear algebra, for nonlinear mappings with some form of gradient-descent, iterative solver.

Most of the fits used in this thesis are non-linear and thus use a gradient descent solver, specifically the Levenberg-Marquardt [66, 70] algorithm. One feature of an iterative algorithm is that one must specify a starting point in parameter space, and that only the

---

[1]particularly for periodic functions

local minimum is found. Whilst potentially a problem should the global solution be desired, it also allows us to solve, with the choice of appropriate starting parameters, some problems (eg: fitting periodic functions) which would otherwise be poorly posed.

6.2 Regularisation

Regularisation is a technique which allows sensible estimates of parameters to be obtained for problems which are either *ill posed*, *poorly conditioned*, or both. To do this one must include additional, prior, information about sensible values for the parameters. Deconvolution problems are typically both *ill posed* and *poorly conditioned* as one is trying to estimate the same number of parameters as one has data points, and as the mapping (ie: the PSF) is usually band-limited, implying multiple objects could be responsible for any one image. This means that some form of regularisation is virtually essential for any useful deconvolution.

Constraints which are often used to regularise deconvolution are positivity, favouring low pixel values, and smoothness. This often takes the form of minimising both data misfit (as above) and some mathematical *'likelihood'* operator applied to the image data. For example minimising the norm of the parameter set prevents the noise leading to extreme values, and minimising derivative norms allows one to select for smoothness.

6.3 Common Deconvolution Techniques

6.3.1 Wiener Filter

Probably the simplest deconvolution method is the Wiener filter. This exploits the fact that if the forward problem is strictly a convolution (ie: linear and shift invariant) with some kernel h, then this can be expressed in the Fourier domain as a multiplication:

$$d(\mathbf{z}) = f(\mathbf{z}) * h(\mathbf{z}) \iff D(\mathbf{k}) = F(\mathbf{k}) \times H(\mathbf{k}) \tag{6.4}$$

The obvious way to recover f would thus be to divide D by H, which however doesn't work due to the noise present in d. The Wiener filter thus introduces regularisation[2] resulting in the following expression:

$$\hat{F} = \frac{D\overline{H}}{H\overline{H} + \lambda^2} \tag{6.5}$$

where $\overline{}$ denotes complex conjugation and λ is a regularisation parameter. The mathematically optimum λ can be derived from the signal to noise ratio.

The Wiener filter has the advantage of being simple, robust and fast (non-iterative). It is however optimised for Gaussian noise, only works for shift invariant problems, and does not allow positivity or other more complicated constraints to be included. Nonetheless for

---

[2]equivalent to minimising the norm in the spatial domain

general purpose deconvolution it produces results which are often very nearly as good as those offered by more advanced deconvolution methods.

6.3.2 Tikhonov Regularised Deconvolution

Tikhonov (also known as Tikhonov-Miller) deconvolution introduces regularisation to the least squares problem by minimising both data misfit and a likelihood function:

$$\hat{\mathbf{f}} = \arg\min\left[\left\|(\mathbf{d} - A\hat{\mathbf{f}})\right\|^2 + \lambda^2 \left\|L(\hat{\mathbf{f}} - \mathbf{f}_{\mathrm{def}})\right\|^2\right] \tag{6.6}$$

The regularisation parameter λ allows one to weight the solution towards a good fit to the data or a high *likelihood* (good agreement with the prior). λ can be chosen rigourously by looking at the so-called 'L-curve' which plots the log likelihood against the log of data misfit. Tikhonov deconvolution is technically also a Gaussian noise solution, although by introducing weights into the data misfit term it can be adapted for more flexible noise situations. An iterative solver is typically used and Tikhonov regularisation lends itself well to the efficient conjugate gradient solver [6]. When using a conjugate gradient solver, Tikhonov regularised deconvolution typically shows fast and stable convergence. Whilst not intrinsically present, a positivity constraint can be introduced by clipping the estimate to zero at each iteration.

6.3.3 Richardson Lucy

R-L deconvolution [79] is a commonly used iterative deconvolution procedure which is mathematically correct for Poisson noise. It has the disadvantages of converging somewhat slowly and, more importantly, being unregularised, although it has an intrinsic positivity constraint. Due to the lack of regularisation a vanilla R-L implementation will typically diverge after a few[3] iterations when applied to microscopic images. Useful R-L implementations thus kludge some form of regularisation by either limiting the number of iterations or by low-pass filtering between iterations. It is also possible to accelerate the process through the use of so-called 'over-relaxation factors' to almost the speed of a C-G based solver [44], although the mathematical basis of this 'optimisation' is more than a little sketchy.

6.3.4 Others

There are several other popular deconvolution methods such as Maximum-A-Posteori-Inference (MAPI - based on Bayesian statistics), van Cittert (iterative inverse filtering), Expectation Maximisation (EM - a somewhat ambiguous term used by Huygens to describe

---

[3]usually somewhere between 20 and 100, although this depends on the data and PSF in a non-trivial fashion.

their algorithm) which I have omitted here along with some less serious methods such as subtracting the images in adjacent z planes from another.

Maximum Likelihood Deconvolution

One often hears the term *Maximum Likelihood Deconvolution*, and it is not always clear what is meant. In general ML-deconvolution is a fairly broad term which encompasses any deconvolution which attempts to find the mathematically best solution to the inverse imaging problem under a certain noise assumption (and potentially a certain prior). If one knows that the noise is Gaussian, then Tikhonov deconvolution and Wiener filtering are ML-deconvolutions whereas R-L is technically not. If one knows the noise is strictly Poisson, then R-L, EM and friends are ML methods whilst Tikhonov and Wiener are technically not. If one doesn't know exactly what the noise is doing (or doesn't care) then Tikhonov etc. are back in business[a]. It can get quite tricky to decide if and when to call a certain algorithm ML or not, and the easiest way is probably to do it by exclusion - any algorithm which doesn't make a mathematically rigourous attempt to take account for both noise and the complete 3D structure of the PSF is definitely not ML.

---

[a]as the ML solution when one doesn't know anything about the noise is least squares.

Part II

Algorithms

Chapter 7

Model Based Evaluation of SMI and 4Pi Data

7.1 Introduction

As mentioned in Chapter 3 the SMI data evaluation method attempts to extract information about the object size and position from an axial profile taken through the object. In order to extract the size and position in a quantitative fashion, it is desirable to fit a mathematical model function to the data. Previous algorithms [91] have relied on manual profile extraction and often some extent of manual start parameter selection. They have also exhibited either poor convergence (especially in biological specimens with a poor signal to noise or in the presence of an aberrated Point Spread Function (PSF)) or the loss of position information due to symmetrification [104], and have required extensive manual postprocessing.

The new algorithm developed in this thesis builds on the previous efforts whilst overcoming several of their shortcomings. It features automatic start parameter selection and is robust against PSF aberration whilst retaining position information. It requires very little postprocessing and, when coupled to a new object finding algorithm, allows the whole process to be automated.

7.2 Fit Algorithm

7.2.1 Fit to the Axial Profile

Previous fit algorithms had attempted to fit both the $\cos^2$ corresponding to the illumination standing wave and the widefield PSF envelope. For the envelope they either used a sinc^2, the mathematical form expected from theory, or approximated this with a Gaussian. This resulted in a model function similar in form to Eqn 7.1.

$$I\left(z\right) = A \cdot \text{sinc}^2\left(\frac{z - z_0}{\alpha}\right)\left(1 - \beta + \beta \cos^2\left(k\left(z - z_0\right) + \phi\right)\right) \tag{7.1}$$

Where A is a scaling factor, α is related to the width of the envelope, z_0 is the position of the object, k is proportional to the frequency of the standing wave pattern, ϕ the phase offset, and β a measure of fringe contrast. Such a model function works well in the case of an ideal PSF. When, however, the PSF is aberrated due to, for instance, refractive index mismatch these models no longer correctly describe the data resulting in a poor fit. Even the small amount of aberration present in a beads sample prepared under optimal conditions has a noticeable effect on fit efficiency. Some improvement is obtained by symmetrifying the PSF by Fourier transforming the data, taking the absolute value, and then the inverse transform [104]. This however loses position information and in the case of more significant aberration can still be insufficient.

The new algorithm avoids this problem by estimating the envelope from the data itself, in effect fitting only the component due to the illumination structure, resulting in the following model function.

$$I\left(z\right) = A \times \mathrm{Env}\left(z\right)\left(1 - \beta + \beta\cos^2\left(k\left(z - z_0\right) + \phi\right)\right) \tag{7.2}$$

The empirical envelope, $\mathrm{Env}(z)$, is extracted by low pass filtering our data to remove the extra frequencies introduced through the illumination structure and obtain the curve we would from a normal widefield microscope. One side effect of removing the envelope from the fit is that although we still have position information from the position of the interference fringes, the fit no longer tells us anything about the location of the PSF envelope. This information is necessary in order to remove the ambiguity introduced by the periodic nature of the illumination structure. The algorithm thus extracts the position of the envelope separately using a modified Centre of Intensity (COI) based approach and uses this to resolve the ambiguity in the fringe based position.

The start parameters are automatically extracted from the profile by finding the locations of the central maximum, the next fringe on one side, and the local minimum between the two. This allows rough estimates of all relevant parameters to be made. To increase the accuracy of these estimates (particularly that of standing wave period), the data is interpolated using lowpass interpolation prior to estimation.

The relationship between object size and fringe contrast (β) (Fig 3.4) has no analytic inverse, so object sizes were previously determined graphically from the measured contrast. This step has been replaced by computing a numerical solution after each fit. A correction for unequal beam intensities as discussed in [104] is also implemented. The fit process is schematically depicted in Fig. 7.1.
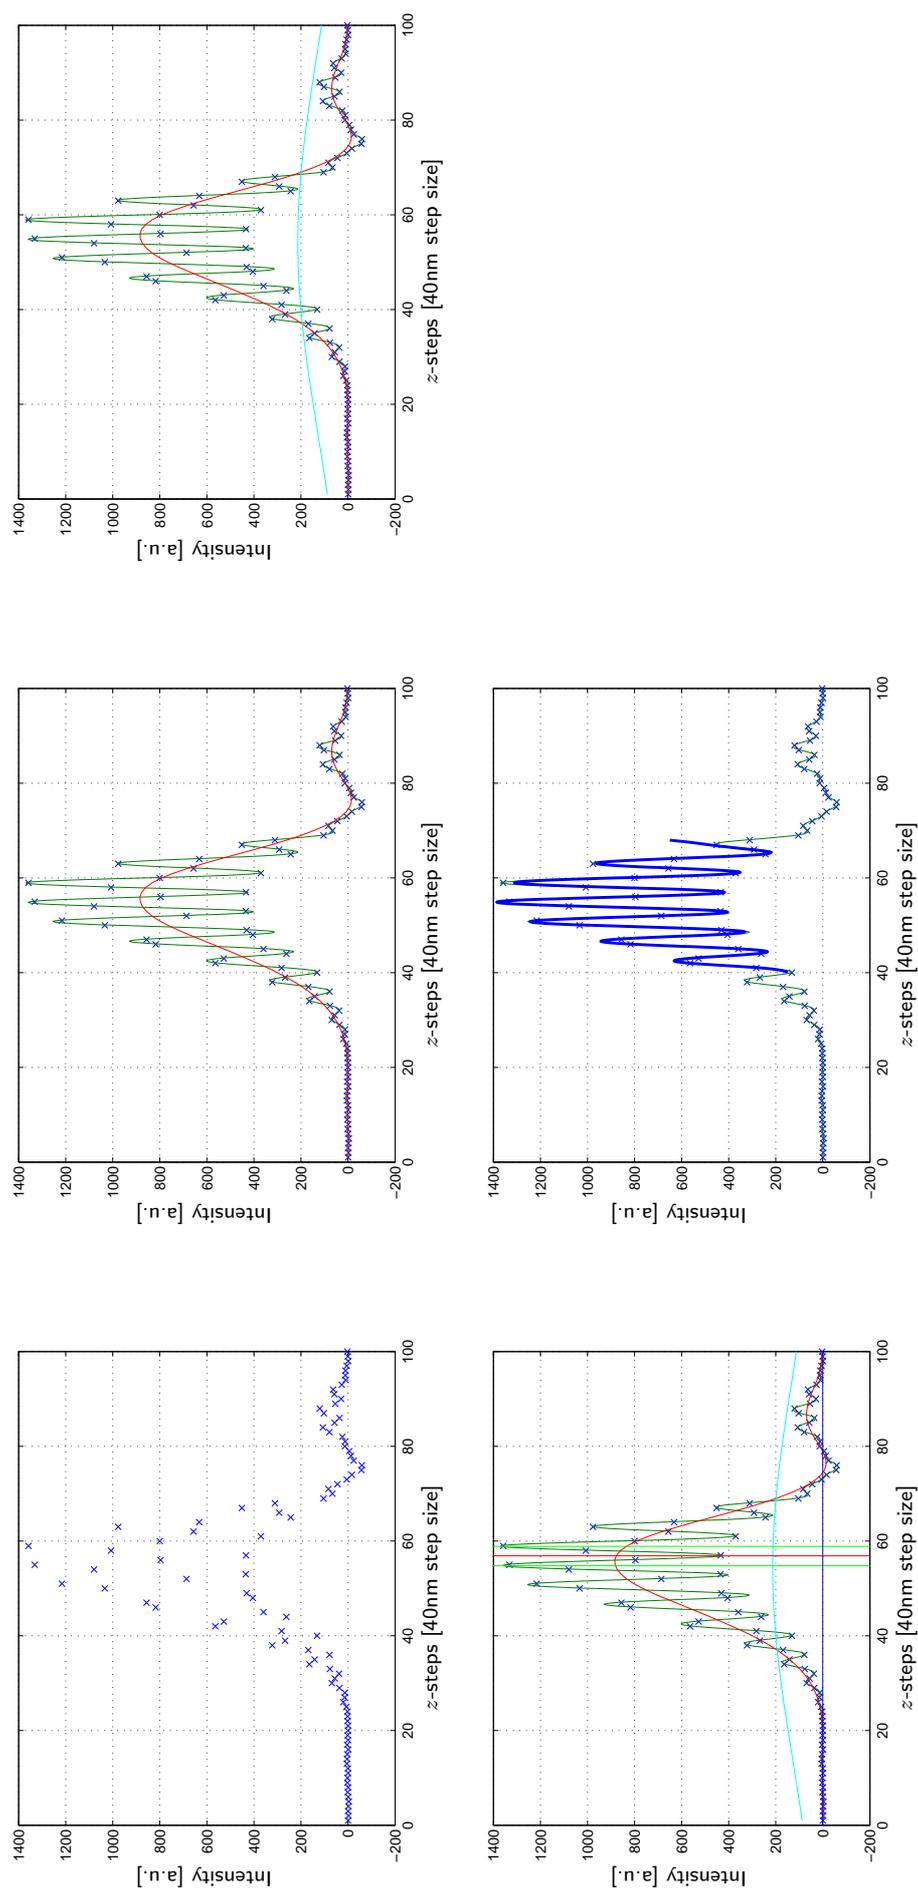
Figure 7.1: **Top left:** Raw intensity values from an axial profile through a bead image. **Top middle:** Raw data with interpolation (green) and estimate of the PSF envelope (red) derived through low pass filtering. **Top right:** Low pass filtering with a stronger filter gives the light blue line, the fit will be performed over the region where the envelope is higher than this line. **Bottom left:** The position and height of the maximum, adjacent peak, and the minimum between the two are used to estimate start parameters for the fit. **Bottom right:** The fit is performed, resulting in the thick blue line.

Figure 7.2: A 2D Gaussian is fitted to the image obtained by averaging over the in-focus slices. Sample: 110nm *transfluor* beads from Molecular Probes.
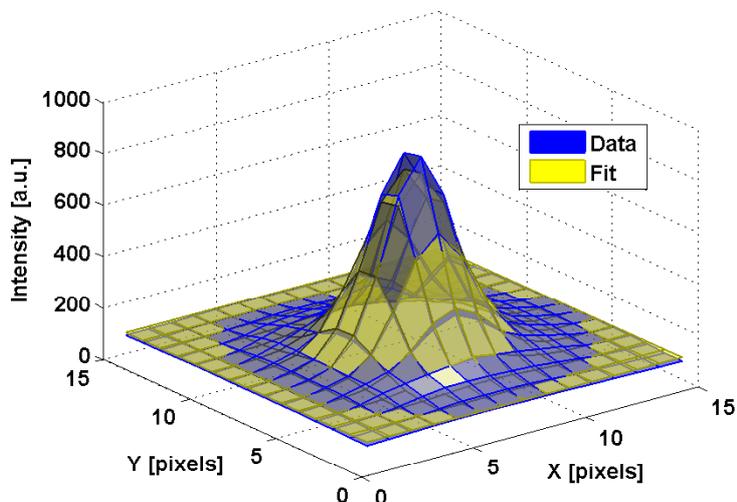
7.2.2   Lateral Gaussian Fit

In addition to the axial profile, a 2D Gaussian is fitted to an average projection of the slices in which the object is in focus ($\sim \pm$ 200nm above and below the axial maximum of the object). This allows a precise determination of the object position in the X-Y plane. It is also possible to estimate the lateral extent of the object based on the FWHM of the fitted Gaussian (see Chapter 14). The Gaussian fit may also be used to determine a background value for the axial fit, eliminating the need for the subtraction of a background ROI. The comparison of size estimates obtained using background subtraction, and those using a background estimate based on the lateral fit is a useful method of quality control for the fit.

7.3   Object Identification

SMI size estimation is useful for point like objects, and one would hence like to locate all point like objects within an image. Approaches based on simple thresholding are both subjective and rarely perform well given either densely clustered objects or a large range of object brightness. Many SMI images fall into both of these categories. A method combining filtering and detection at multiple thresholds was thus developed which allows the independent identification of objects over a wide range of intensities and with a moderate degree of clustering.

The first step is a reasonably narrow ($\sigma \approx 1$ pixel) lowpass Gaussian filter, followed by a slightly wider ($\sigma \approx 3$ pixels) Gaussian highpass. The resulting bandpass filter serves to suppress both noise, and any objects (including a non-uniform background) which are larger than a few pixels. The concept is similar to that of scale spaces [56], and has the effect of restricting the image to those features which are of a size interesting for SMI microscopy (larger objects - for example chromosome territories - which are inhomogeneous
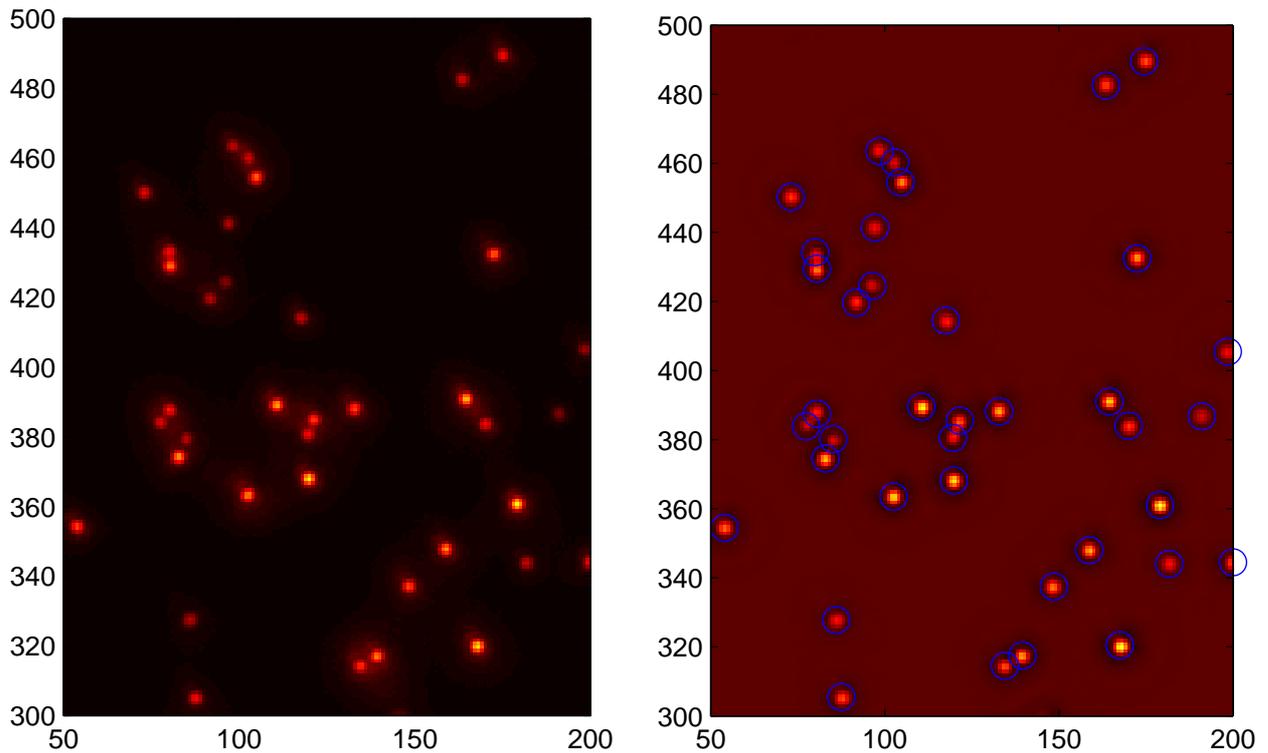
Figure 7.3: Left: Maximum projection from a beads sample. **Right:** Filtered image showing location of identified points. Note that with a simple threshold based method it would not have been possible to chose a threshold which separated the pair of objects in the centre of the image ($\sim$(125,385)) whilst still detecting the weak object at the right side of the image ($\sim$(190,385)). This effect is much more prominent in most biological images. The axes in both cases are in 108nm pixels.

will be reduced to a a set of foci corresponding to those features of the structure at the scale to which the filter is sensitive).

Following this the image is processed at progressively decreasing thresholds with the objects found at each threshold being removed from the image before proceeding to the next. In order to remove the objects found at each step, the thresholded image is blurred with a Gaussian approximation to the PSF and subtracted from the working image. This means that bright objects will disappear from the image upon detection, allowing weaker objects in their vicinity to be detected in subsequent steps. This is in principle similar to the CLEAN [55, 20] processing of radio astronomy data. If the PSF estimate is too small, the subtracted signal would be smaller than the real image of the object, leaving a doughnut shaped structure which could be detected again at lower thresholds. To avoid this we use a PSF estimate which is slightly too large.

7.4   Implementation

Due to the probabilistic nature of image formation, the noise in the images is expected to follow a Poisson distribution, although a small Gaussian component from the detector read out noise is also expected. A weighted least squares algorithm where the variances of each pixel are assumed to be proportional to the pixel intensity plus a small constant term corresponding to the camera read out noise of $\sigma = 4$ counts has thus been used. As well as modelling the system response, the small constant term serves to improve the stability of the fit in the presence of low pixel counts, an area where the use of intensity dependant weights is traditionally problematic. When compared in simulations to a maximum likelihood fit of the same data, no significant difference in the fit accuracy was observed, with the least squares method having the additional advantage of producing sensible error estimates for the fitted parameters. An unweighted least squares, and weighted least squares without the small constant term both produced results which were significantly worse.

The fits are performed using a Levenberg-Marquardt nonlinear weighted least squares fitting algorithm written in Matlab. The whole process is automated, with the intensity profiles required from the fit extracted at the location of each object found during the object identification step.

Sviewer Scripts

`connectToDB`: Establishes a connection to the SQL database on *verleihnix* where parameters for the data analysis are stored.

`svlauncher`: Displays a GUI allowing SMI, confocal, or 4Pi files to be loaded. The loaded data is typically available as the *global* variable `d` (to see this from the command window it might be necessary to declare it ie: `global d;`).

`sviewer`: Displays the GUI shown in Fig. 7.4. Also accessible through `svlauncher`.

`ofind3`: Performs the object identification procedure described in 7.3. The result is the matrix `pos` containing pairs of (x, y) coordinates identifying the found points.

`dofits`: Performs the fit procedure for each of the points in `pos`. Results vary depending on the type of data - for SMI `params` has the results of the axial fits, `pbs` those from the lateral Gaussian, `ss` the axial size estimates, `zpos2s` the COI based z-position estimates, and `I` the indices of those points for which the fit is deemed to have been successful (there are also several more SMI relevant variables eg: `zexacts`, `perror`, and potentially `params_5point` - for the full list see `dofits.m`).
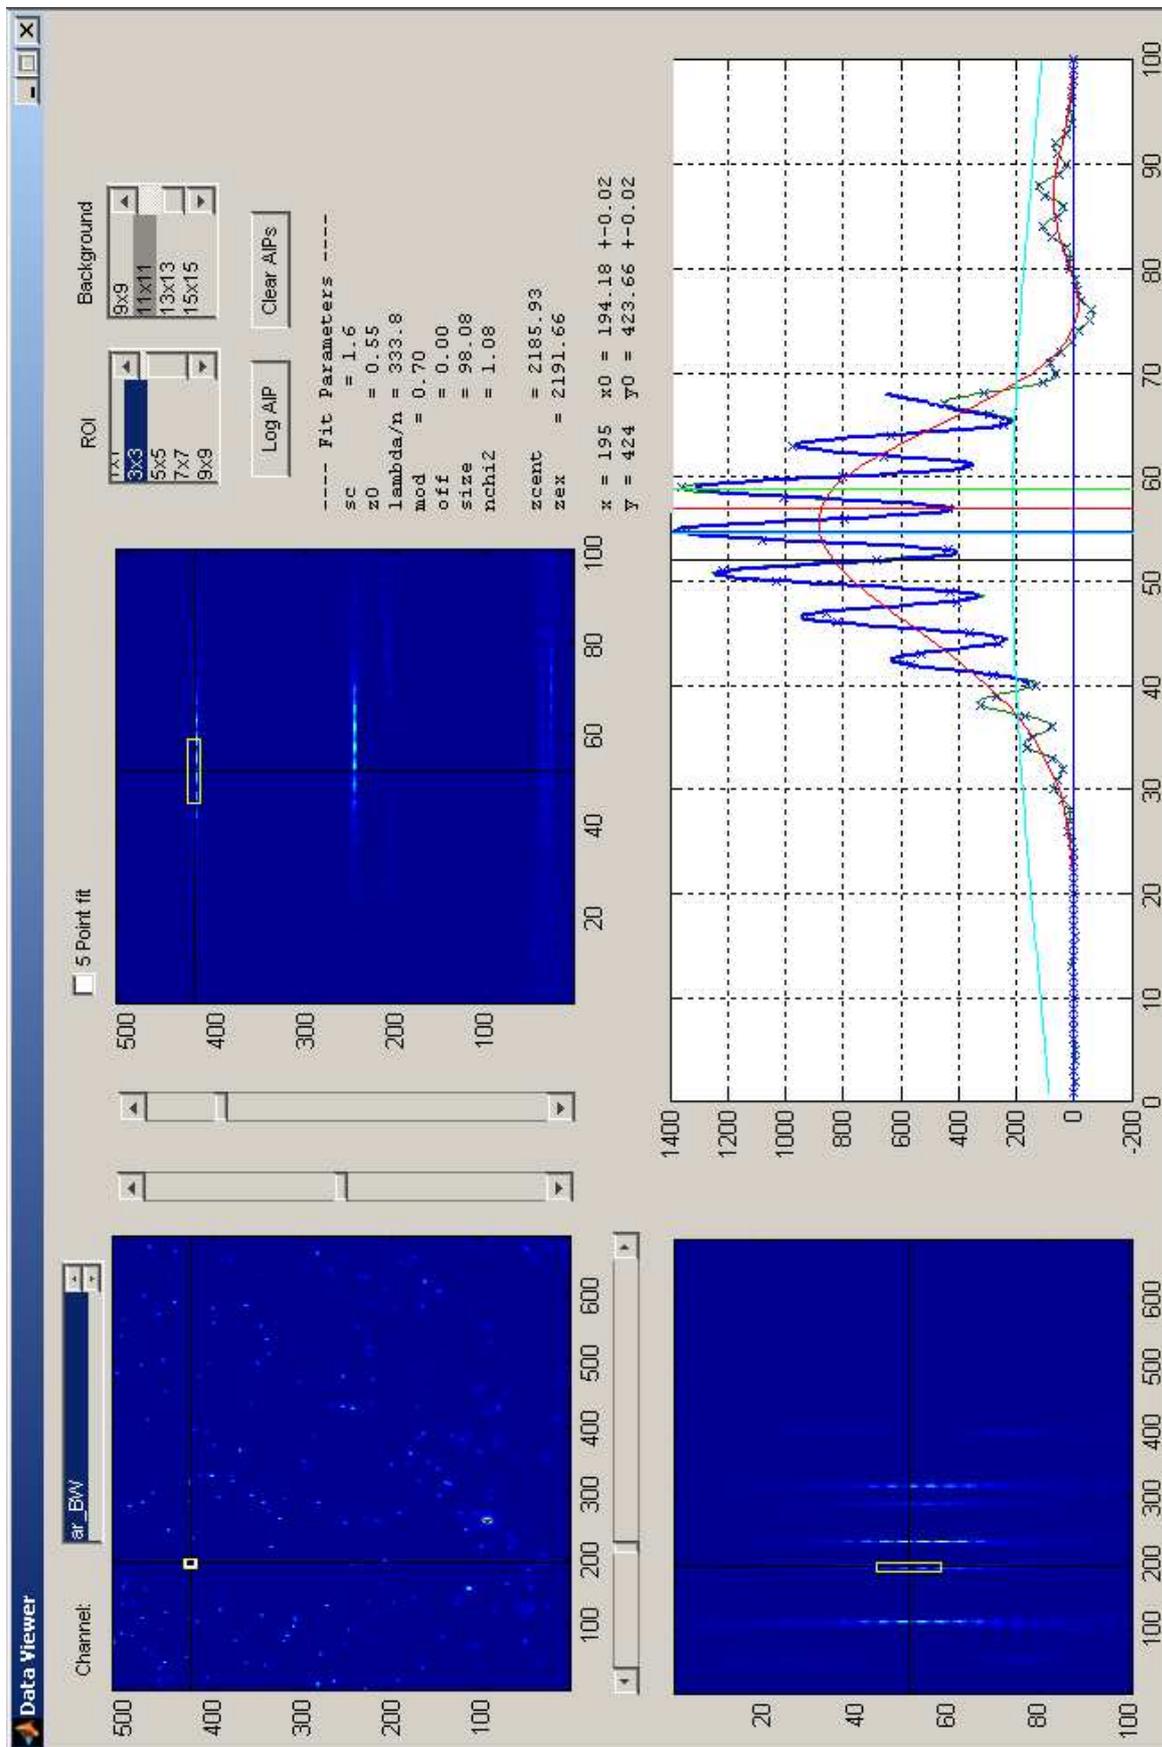
Figure 7.4: Screenshot of the 'sviewer' GUI when analysing a beads sample. This interface allows the manual inspection of individual loci, showing both the fit curve, and the fitted parameters. The result of the lateral Gaussian fit is shown in a separate window.

7.5   Results

Figure 7.4 shows a screenshot of the interactive data viewer where fitting can be performed by clicking on individual points. Using the object identification procedure the finding of objects and the subsequent fitting of profiles can be automated. Scripts for the generation of reports also exist. Figures 7.5 and 7.6 show the pdf report generated by applying the evaluation procedure to a cell in which replication foci were labelled. This report consists of a summary of all detected foci as well as data which shows how the detected sizes depend on, for example, intensity. A table of all measurements is also saved. Further results of the algorithm can be found in Chapters 10, 11 and 13 and in [71, 72, 32, 7].

Sviewer Scripts

Several commands are available for viewing and saving the results produced by the `ofind` and `dofits` combination.

`showres`: Displays a maximum intensity projection of the data, a graphical representation in which the objects are rendered as spheres at there locations, coloured according to estimated size, and a histogram of the estimated sizes.

`savefits_rep`: Saves the fit data in a *.txt* file readable by Excel, and as a ps/pdf report including the figures produced by `showres`.

`ssbounds`: Calculates upper and lower bounds for the size estimates.

`show_lots_more_res`: Shows various interesting plots including dependance of size on threshold used to detect objects. Needs results from `ssbounds`

`ressurect`: Reloads data from a *.txt* file such as is produced by `savefits`.
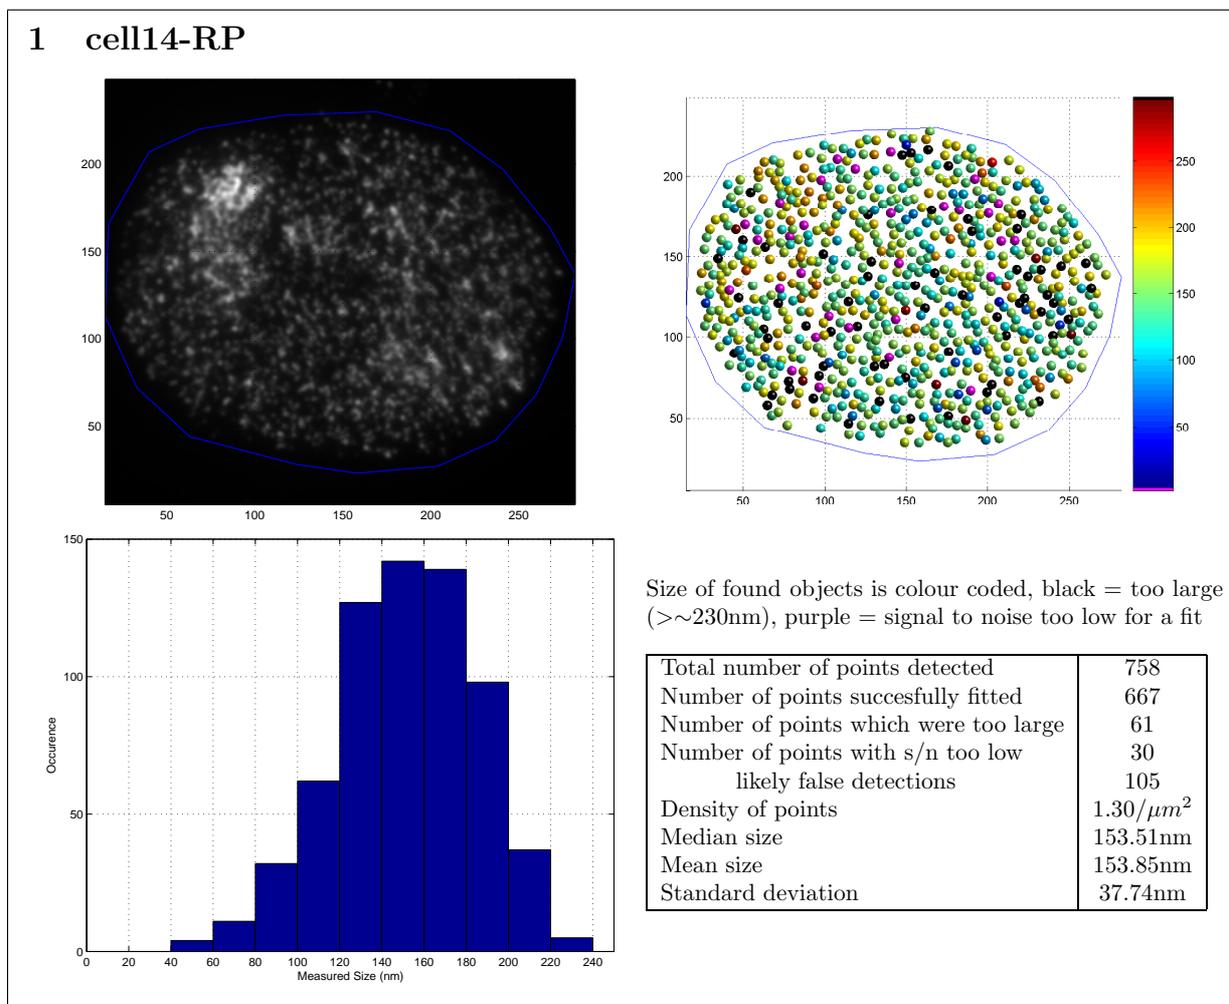
Figure 7.5: Page 1 of analysis output (of an example cell in which replication foci are labelled) obtained using the *sviewer* package. An overview image of the cell is given, along with a colour coded representation of the detected points and their estimated sizes. The distribution of sizes within the cell and a few basic statistics are also shown.
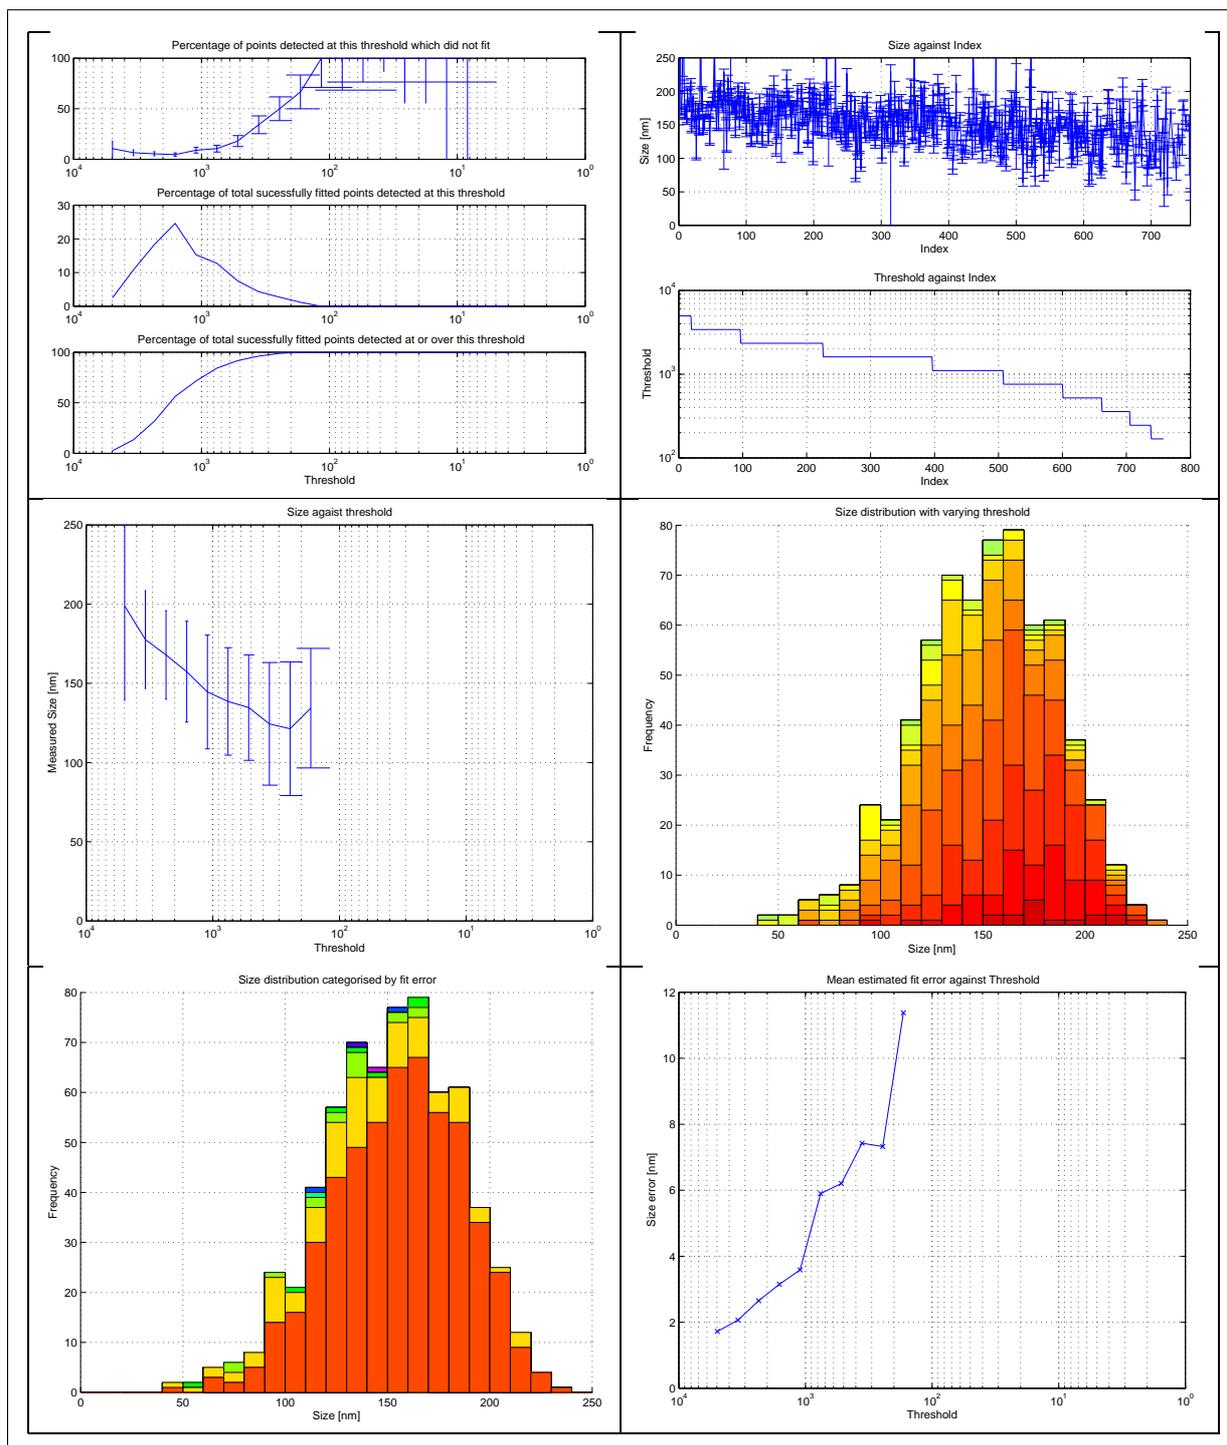
Figure 7.6: Page 2 of the analysis output (sample as in Fig. 7.5). The figures here relate the estimated sizes and the size errors to the threshold at which they were detected. Also shown is the number of loci detected, and the percentage of these successfully fitted at each threshold.

7.6 Automation

By writing a script to automatically load data files and execute the other scripts, it was possible to fully automate the process or data evaluation. Some manual intervention is required at the start to determine a suitable threshold range for the particular class of specimen, after which the rest of the analysis can be unattended.

Sviewer Scripts

`batch_proc`: Automatically processes all *.kdf* files in a directory.

7.7 Adaptation to 4Pi

The same data analysis scheme can easily be applied to 4Pi data, the only real change needed is a new model function which corresponds to the 4Pi images.

7.7.1 4Pi Model Function

The model function used for 4Pi data analysis is as follows:

$$I(z) = h\left(z|z_0, w, n/\lambda, \phi\right) * p\left(z|r\right) \tag{7.3}$$

where $h(z)$ is the axial component of the PSF, $p(z)$ the projection of the fluorophore distribution along the z-axis, and the $*$ operator represents convolution. The PSF is approximated as:

$$h(z) = \mathrm{sinc}^2\left(\frac{9}{5}\frac{z - z_0}{w}\right)\left[\mathrm{sinc}^2\left(\frac{z - z_0}{w}\right)\cos^2\left(\frac{2\pi n\left(z - z_0\right)}{\lambda} + \phi\right)\right] \tag{7.4}$$

where z_0 is the z-position, w is related to the width of the envelope, n/λ the wavelength in the medium, and ϕ an arbitrary phase offset. The object model used is a sphere, giving a parabolic axial projection:

$$p(z) = \begin{cases} I_0(r^2 - z^2) & z \in [-r, r] \\ 0 & \text{else} \end{cases} \tag{7.5}$$

where I_0 is the object intensity, and r the radius.

7.7.2 Results

The modified fit procedure can be successfully applied to 4Pi images, as shown in the example fit to a 100nm bead (Fig. 7.7). The behaviour of the fitting procedure when applied to beads is discussed in much more detail in Elvira Steinwands thesis [94]. One feature which invites comment, however, is the inaccuracy in the estimates of bead size (the example bead from Fig. 7.7 measures 250nm despite a real size of 100nm). The sizes

Figure 7.7: Fit of the 4Pi model function to an experimental profile from a bead measurement. The extracted parameters are: z-position $(z_0) = 60.85$ voxels, FWHM of envelope = 903nm, phase of interference $(\phi) = 21°$, wavelength in media (λ/n) = 560nm, size = 250nm, sidelobe height = 55%.
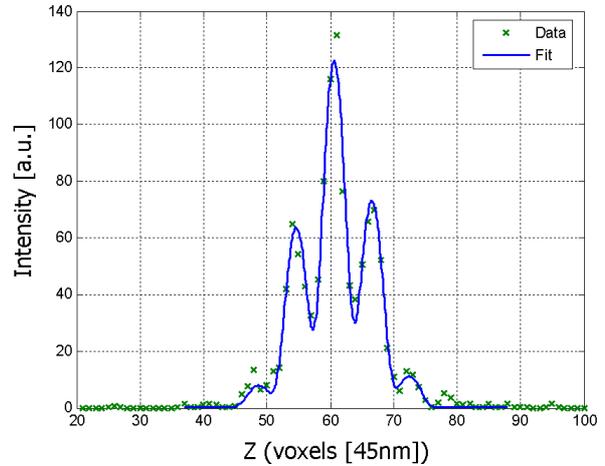
Figure 7.8: Distribution of refractive indices within a cell nucleus mounted in *VectaShield* obtained by fitting the model function to a 4Pi image of replication foci.
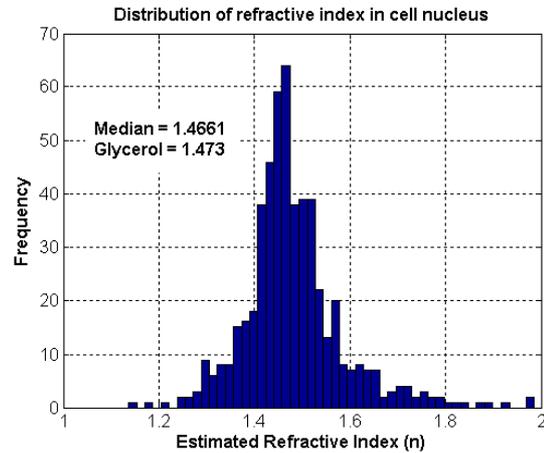
are typically significantly overestimated and there are considerable variations between the mean sizes of measurements taken at different times. Within one series, however, the distribution of size estimates is typically relatively narrow. There are several possible causes for this, two of which are instrumental vibrations, and the limited coherence of the 2-photon laser. Vibrations would tend to blur the signals, whilst a slight misalignment of the optical path lengths would effect the modulation depth. Both effects have been shown to occur when observing the mirror typically included as part of the 4Pi sample preparation.

An example of the application of the method to biological specimens is shown in Fig. 7.8 where the fit was applied to each of the foci in an image of replication foci. Through comparison of the wavelength in the medium (λ/n) as obtained from the fit with the actual laser wavelength it was possible to estimate the refractive index at each of the points within the nucleus.

Chapter 8

4Pi Deconvolution With a Variable Phase

8.1   Introduction

Note: this chapter is almost completely borrowed from *"4Pi microscopy deconvolution with a variable PSF phase"*, Baddeley et Al., 2006 [4].

As mentioned earlier, sidelobe removal through some form of digital postprocessing is an essential part of 4Pi imaging. Various techniques for sidelobe removal exist, including simple 3 or 5 point based filters and the application of traditional ML deconvolution processes.

One area where all current algorithms have difficulties is when the PSF depends on the position in the sample. The major PSF variation observed in 4Pi microscopy is a phase
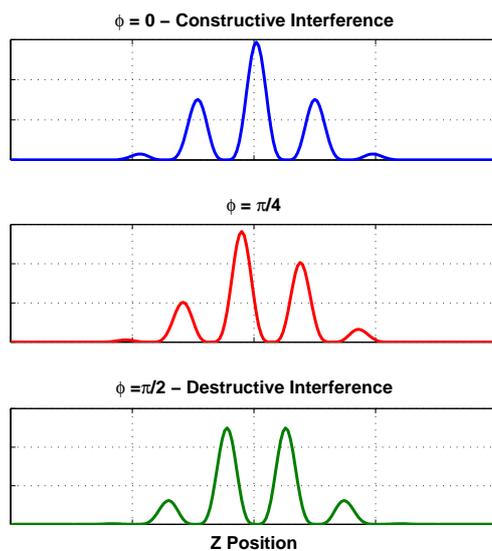
Figure 8.1: Schematic representation a z-axis profile through a 2-photon 4PiA PSF showing the effect of a varying phase. This is seen as a shift of the interference fringes with respect to the maximum of the envelope.

shift in the underlying interference pattern which is caused by a difference in the optical path length from each objective [28, 86]. This is manifested by a shift from constructive to destructive (or anywhere between) interference at the centre of the PSF, as shown in Fig. 8.1. When constant over the whole specimen this can easily be corrected with the phase adjustment piezo or by using a suitable PSF measurement with a conventional ML algorithm. When this varies throughout the object it is however impossible to correct for using the existing algorithms. Such a variation is commonly seen in biological specimens where precise control of the refractive index inside the cell is difficult. Various hardware solutions to this problem have been presented [86, 49, 13], involving continuous compensation of the phase during measurements and monitoring both outputs of the beamsplitter prism. None of these options are presently available of the commercial 4Pi, although it is planned to introduce a mechanical phase compensation. The methods presented here allow one to simulate images with such a spatially varying phase shift, to estimate the phase shift from acquired images, and to perform reconstruction of experimental images.

8.2   The Phase Shifted 4Pi PSF

Most deconvolution methods require a computationally inexpensive method of computing the forward transformation, $\mathbf{d} = H\mathbf{f}$ (the mapping from the vector of object voxels, $\mathbf{f}$, to the data voxels, $\mathbf{d}$), and also in several cases the conjugate transpose of this mapping. In the case of a spatially invariant PSF this is a convolution and can be inexpensively performed using Fourier domain multiplication. When the PSF is not spatially invariant, as is the case with a 4Pi PSF in non-ideal optical conditions, then this approach is not applicable. In the special case of a phase shift in the interference maxima of the 4Pi PSF, the forward transform, and its conjugate transpose, can be approximated as the linear combination of a small number of convolutions. In the following the focus will be on 2-Photon 4Pi type A, as this the type of 4Pi microscope that is currently commercially available, and as the inverse problem is not even moderately well posed for the 1 photon variants. The principles are however applicable to all 4Pi variants.

The 4Pi PSF can be described as an envelope corresponding to the PSF of a confocal microscope operating at the respective excitation and emission wavelengths $I_{\text{env}}(x, y, z) = I_{\text{det}}I_{\text{ex}}$ or in the case of 2-photon excitation $I_{\text{det}}I_{\text{ex}}^2$, multiplied by the interference pattern produced by the coherent addition of light from both objectives [1] $I_{\text{fringes}}(x, y, z)$ giving

$$I_{\text{PSF}} = I_{\text{env}}(x, y, z)I_{\text{fringes}}(x, y, z). \tag{8.1}$$

These interference fringes correspond to the wavefronts of the excitation beam, which are plane in the focus of a Gaussian beam and acquire curvature at increasing distance on either side. If the envelope, I_{env}, is sufficiently small, such that the detected intensity is very low before the curvature of the fringes becomes noticeable, it is possible to approximate the

---

[1]This is true for 4Pi A. In the cases of 4Pi B the interference pattern produced by the detected light back-projected into the object space must be taken, and in C the sum of both.

fringes as a plane wave interference pattern having only a z-dependence, $I'_{\text{fringes}}(z)$. This approximation forms the basis of the 3 and 5 point deconvolution schemes extensively used in 4Pi microscopy, and is a good approximation in the two-photon case. For plane wave interference and one photon excitation, the normalised fringe pattern is expected to be $I'_{\text{fringes}_{1\text{photon}}}(z) = \cos^2\left(\frac{2\pi nz}{\lambda_{\text{ex}}} + \phi\right)$ where n is the refractive index, λ_{ex} the excitation wavelength, and ϕ a phase term. In the case of two photon excitation, the detected intensity is proportional to the square of the excitation giving $I'_{\text{fringes}}(z) = \cos^4\left(\frac{2\pi nz}{\lambda_{\text{ex}}} + \phi\right)$. In the general case, and in samples without ideal optical conditions, ϕ cannot be thought of as a constant and must be regarded as a function, $\phi(\mathbf{r})$, of the position $\mathbf{r} = (x, y, z)$ in the object space. Substituting in Eqn. 8.1 gives us an approximate expression for the 2-photon 4Pi PSF in the presence of an arbitrary, spatially varying, phase shift:

$$I_{\text{PSF}} \approx I_{\text{env}}(\mathbf{r})\cos^4\left(\frac{2\pi nz}{\lambda_{\text{ex}}} + \phi(\mathbf{r})\right) \tag{8.2}$$

The mapping $(\mathbf{d} = H\mathbf{f})$ responsible for the imaging process can be expressed as

$$d(\mathbf{r}) = \int f(\mathbf{s}) I_{\text{PSF}}(\mathbf{r} - \mathbf{s}, \mathbf{r}) \, d\mathbf{s}, \tag{8.3}$$

where $f(\mathbf{r})$ is the fluorophore distribution in the sample and $d(\mathbf{r})$ the detected intensity. The dual parametrisation of the PSF $I_{\text{PSF}}(\mathbf{r} - \mathbf{s}, \mathbf{r})$ reflects its spatial variance. Due to the dependence of ϕ on the position in object space this mapping cannot be expressed as a convolution and simply performed. Expanding the $\cos^4$ term in (8.2) gives us

$$
\begin{aligned}
I_{\text{PSF}} \approx I_{\text{env}} \frac{1}{4} \Bigg[&\frac{3}{2} + 2\cos(2\phi)\cos\left(\frac{4\pi nz}{\lambda_{\text{ex}}}\right) - 2\sin(2\phi)\sin\left(\frac{4\pi nz}{\lambda_{\text{ex}}}\right) \\
&+ \frac{1}{2}\cos(4\phi)\cos\left(\frac{8\pi nz}{\lambda_{\text{ex}}}\right) - \frac{1}{2}\sin(4\phi)\sin\left(\frac{8\pi nz}{\lambda_{\text{ex}}}\right) \Bigg].
\end{aligned} \tag{8.4}
$$

Using the linearity property of integration we can now approximate the imaging through:

$$
\begin{aligned}
\mathbf{d} \approx \frac{1}{4}\Bigg(&\frac{3}{2}I_{\text{env}} * \mathbf{f} + 2\cos(2\phi)\left\{\left[\cos\left(\frac{4\pi nz}{\lambda_{\text{ex}}}\right)I_{\text{env}}\right] * \mathbf{f}\right\} - 2\sin(2\phi)\left\{\left[\sin\left(\frac{4\pi nz}{\lambda_{\text{ex}}}\right)I_{\text{env}}\right] * \mathbf{f}\right\} \\
&+ \frac{1}{2}\cos(4\phi)\left\{\left[\cos\left(\frac{8\pi nz}{\lambda_{\text{ex}}}\right)I_{\text{env}}\right] * \mathbf{f}\right\} - \frac{1}{2}\sin(4\phi)\left\{\left[\sin\left(\frac{8\pi nz}{\lambda_{\text{ex}}}\right)I_{\text{env}}\right] * \mathbf{f}\right\}\Bigg)
\end{aligned} \tag{8.5}
$$

where $*$ represents convolution. The five integrals obtained when substituting (8.4) into (8.3) all involve a spatially invariant PSF term, and can thus be expressed as convolutions. These convolutions are easily calculated, for instance by the technique of Fourier space multiplication. Similar reasoning can be applied to the conjugate transpose of the transformation.

The concept is easily extended to the 4Pi-C variants through the introduction of an additional sinusoidal term in I'_{fringes} and a second phase map[2].
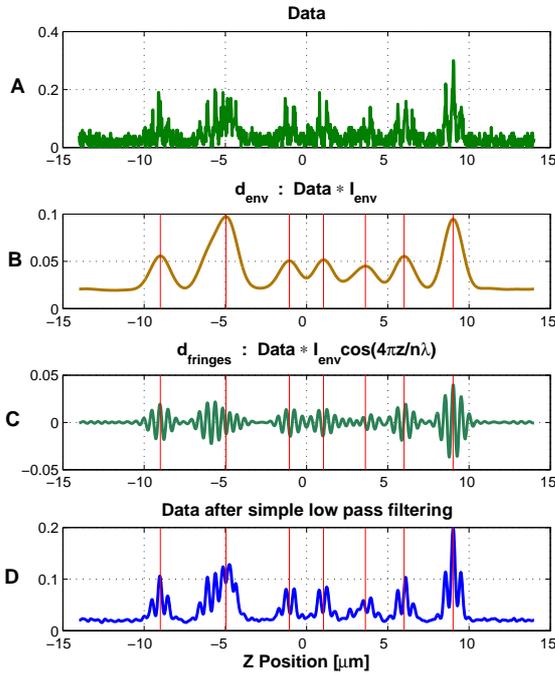
8.3   Phase Estimation

In order to be able to perform a reconstruction based on our variable phase PSF model, as outlined above, we require knowledge of the PSF phase at every point in the image. It is possible to obtain information about the phase during image aquisition using methods such as those presented by Hell et al [49] and Blanca et al [13]. Unfortunately the hardware implementation of the commercial 4Pi microscope does not allow these techniques to be used and one must attempt to extract this information from the resulting images. Given a method for calculating the forward transformation, which maps the fluorophore density and PSF phase in the object space to an image, it is theoretically possible to invert the problem and simultaneously estimate the object and phase through some form of maximum likelihood optimisation. As one must estimate two parameters for each data point and as the vast majority of voxels contain absolutely no information about the phase, being either dark, or part of a larger structure, this problem is both underdetermined (ill conditioned) and poorly posed and it is difficult to envisage a regularisation technique capable of improving the conditioning enough to allow a ML algorithm to produce a sensible result. One alternative is to independently estimate the phase from structures in the image where phase information is present, such as reference beads or thin membranes aligned perpendicular to the optical axis, and interpolate this over the whole image before deconvolving using this phase estimate.

The technique for phase extraction presented here is one of several possible methods, but the requirements for most methods will be similar: structures which are axially small, separated from other structures in the axial direction, and having a good signal to noise ratio. Our method simply attempts to measure the distance between the central (or highest) maximum, and the centre of the envelope. In the case of destructive interference, where both maxima are equally high, the choice of which distance to measure becomes arbitrary. In order to be robust in the presence of noise, the images are filtered before processing. When only the position of an object is important, optimal filtering is obtained by making the cross-correlation of the image with the PSF. In order to extract the centre of each object, one thus performs a cross correlation of the image, $\mathbf{d}$ with the envelope of the PSF, I_{env} giving us an image $\mathbf{d}_{\text{env}}$. The position of the fringes is then extracted using the cross correlation with $\cos\left(\frac{4\pi n z}{\lambda_{\text{ex}}}\right) I_{\text{env}}$ giving us $\mathbf{d}_{\text{fringes}}$, the $\cos\left(\frac{4\pi n z}{\lambda_{\text{ex}}}\right)$ term from the expansion above being a sufficiently good approximation for this purpose. The distance from the peaks of $\mathbf{d}_{\text{env}}$ to the correspondingly nearest peak in $\mathbf{d}_{\text{fringes}}$ then gives us the phase at that point in the image. Because $\cos\left(\frac{4\pi n z}{\lambda_{\text{ex}}}\right) I_{\text{env}}$ averages to zero on a distance scale greater than a few periods of the cos term, the envelope of $\mathbf{d}_{\text{fringes}}$ tends to zero over areas of the image with no high frequency components (corresponding to small structures) in the axial

---

[2]related to the original through a scaling and offset

Figure 8.2: Line scan along z-axis through simulated test image showing the principles of the phase estimation procedure. A cross-correlation is performed between the raw data (**A**) and the PSF envelope (i.e. the appropriate confocal PSF) giving d_{env}, the profile in **B**. From the peaks in this image, the centres of all objects are estimated (as shown by the vertical lines). The raw data is additionally correlated with $I_{\mathrm{env}} \cos\left(\frac{4\pi n z}{\lambda_{\mathrm{ex}}}\right)$ to extract the fringe position (as shown in profile **C**). To verify that the reference point belongs to an isolated object the relative heights of the peaks in the low pass filtered image (**D**) is analysed. In the displayed case, the second object from left would be rejected. Note that the cross correlations / filtering were carried out on the xz image, resulting in a better signal to noise ratio than would be expected from the 1D filtering of the data shown.

direction. A threshold on the envelope of $\mathbf{d}_{\mathrm{fringes}}$ can thus tell us which parts of the image are interesting for phase extraction. This is however not completely sufficient, as fringes will also be visible at the edges of extended structures, and where several small structures are in close vicinity or periodically spaced. Whilst it should theoretically be possible to estimate the phase from such structures, additional information about the structure would be required. In order to remove structures where the assumptions of negligible axial thickness and isolation from neighbouring objects are not fulfilled, we can examine the height of the maxima on either side of the central peak. In the case of an extended or periodic object, these will be elevated. Figure 8.2 illustrates the application of phase estimation using these principles to a line scan through the simulated 2D test image shown in Fig. 8.5.

The method detailed above is limited in accuracy to one voxel. An alternative method of phase estimation, offering sub-voxel phase resolution, extends the above by fitting a suitable model function to the data at each reference point rather than simply looking for the maxima. One could also interpolate d_{env} and d_{fringes} in the z-direction to achieve the same result. In any case, in order to accurately estimate the phase over the entire image the values from the resulting measurement points must be interpolated. As an alternative to the automatic selection of reference points, manual selection could be desirable in poor signal to noise conditions, or when speed/memory usage is an issue. In this case the judgement as to whether a structure is isolated or not is based on the experience of the user.

In the examples presented here, a smoothly varying phase is assumed. It could be advantageous to further develop the phase estimation procedure in order to allow discon-

tinuities in the first derivative, such as would occur at a refractive index step. Such a model would however be likely to require a very large number of control points, or additional information such as could be obtained from Differential Interference Contrast (DIC) images. It is in any case unlikely that the small errors introduced in the assumption of smoothness will be noticeable in the resulting images, although they could be significant for sub-resolution distance measurements.

8.4   Deconvolution Algorithm

Conventional deconvolution algorithms for light microscopy are optimised for either Gaussian or Poisson noise, with algorithms such as Richardson-Lucy [79] which are based on Poisson noise being mathematically superior for almost all cases of confocal imaging. Whether this is also the case for 4Pi images remains to be established. The avalanche photodiodes used in the commercial 4Pi have very limited ($\sim$3 bit) dynamic range. This is due to the fact that they act in photon-counting mode, combined with a limited bandwidth and relatively short ($\sim$1μs) per voxel integration time. In order to overcome the limited dynamic range, it is thus common practice in 4Pi microscopy to significantly oversample ($\sim$20nm voxel size) and to smooth the resulting images with a Gaussian filter. This averaging will alter the noise characteristics, with the central limit theorem implying that the noise should become more Gaussian.

The algorithm used solves a Tikhonov regularised, weighted least squares problem

$$\arg \hat{\mathbf{f}} \min \left[\left\| W(\mathbf{d} - H\hat{\mathbf{f}}) \right\|^2 + \lambda^2 \left\| L(\hat{\mathbf{f}} - \mathbf{f}_{\text{def}}) \right\|^2 \right] \tag{8.6}$$

using a modified conjugate gradient solver [96, 97, 6]. This attempts to find an estimate of the object ($\hat{\mathbf{f}}$) which minimises both a misfit term - $\left\| W(\mathbf{d} - H\hat{\mathbf{f}}) \right\|^2$ - describing how well the object matches the data, and a likelihood term - $\left\| L(\hat{\mathbf{f}} - \mathbf{f}_{\text{def}}) \right\|^2$ - which represents some prior knowledge (most often an assumption of smoothness) about the form of the object. The parameter λ determines the respective weighting of the data misfit and likelihood terms. In contrast to ordinary least squares, which is the solution in the presence of uniform Gaussian noise, choosing $W_{ii} \propto \frac{1}{\sqrt{\mathbf{d}_i}}$, $W_{ij|i \neq j} = 0$ gives a scaled Gaussian noise model where the noise amplitude scales with the square root of pixel intensity. This allows us to approximate Poisson noise, and might indeed be more appropriate than a pure Poisson model in cases where the signal is averaged. These weights are in practice not particularly useful as they give infinite weight to zero valued voxels in the image. A combination of $\sqrt{n}$ and uniform Gaussian noise, $W_{ii} \propto \frac{1}{\beta + \sqrt{\mathbf{d}_i}}$, achieved by putting a constant β in the denominator, yields better results. For a likelihood function, L, a $3 \times 3 \times 3$ approximation to the second derivative is used, as this yields the qualitatively best results on both simulated and real images. The forward mapping (H) and its transpose are evaluated as per section 8.2. A positivity constraint is also introduced by clipping negative values to zero at each iteration.
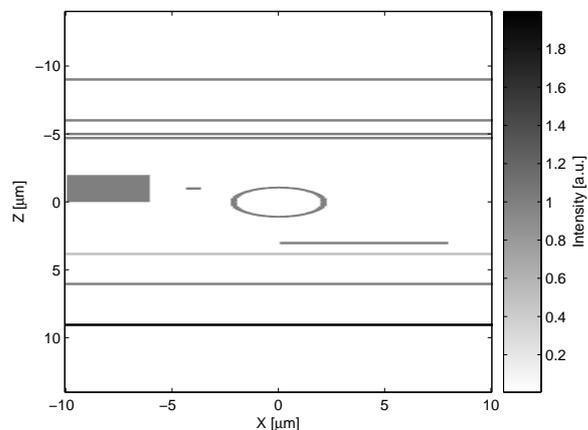
Figure 8.3: Simulated test pattern
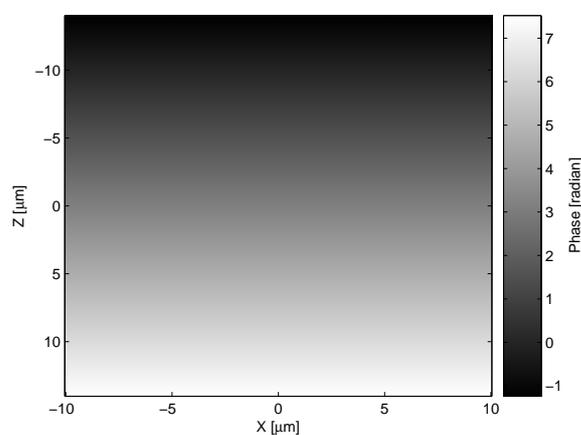

Figure 8.4: Simulated phase map

This deconvolution approach was chosen due to its robustness, fast convergence, the ability to improve the conditioning of the problem through regularisation, and the availability of Matlab code [97] for the conjugate gradient solver. Whilst the Richardson-Lucy algorithm is to some extent regularisable through limiting the number of iterations, or filtering the images between steps [44] it is easier to guarantee convergence with the Tikhonov / least squares approach.

In order to test the proposed methods for image generation, phase estimation and image reconstruction, they have been applied to both a test pattern, and real images as detailed below.

## 8.4.1 Deconvolution of Simulated Data

The test pattern used was the two dimensional arrangement of geometrical shapes shown in Fig. 8.3 and its corresponding phase map (Fig. 8.4). Note that the phase shift applied in this case is fairly extreme - shifting from constructive interference to destructive and on to constructive again as one moves down the z-axis (vertical). A two dimensional image was

Figure 8.5: Two photon 4Pi-A image generated from test pattern and phase map with Poisson noise
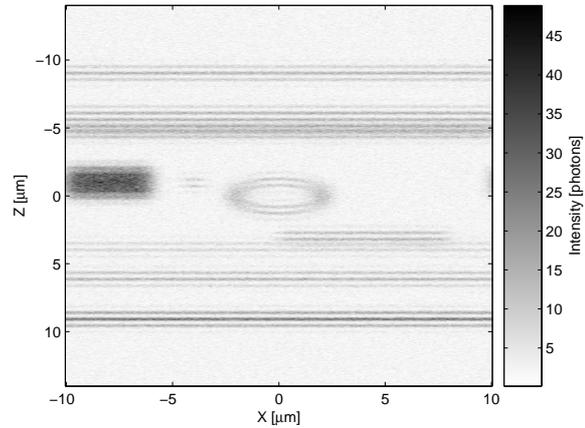

**Figure 8.6:** Reconstructed phase map showing automatically selected control points (crosses)
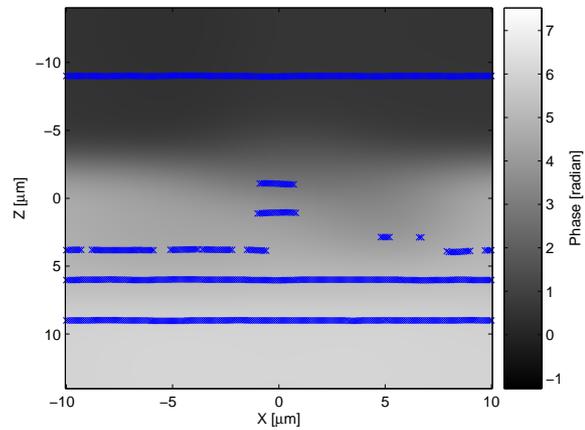
used here as the computation is quicker, and the results are easier to display and interpret than in the 3D case. The algorithm is however the same and is easily generalised to 3D, as will be shown with the biological results. An image, shown in Fig. 8.5, was obtained by applying the procedure in section 8.2 to the test image using the phase map and an approximation to the theoretical PSF envelope, and passing the result through a Poisson noise process. The PSF envelope was approximated with a $\mathrm{sinc}^2(9z/5w)\mathrm{sinc}^4(z/w)$ axial component[3], the width of which was chosen so as to give a sidelobe height of approximately 50 percent, multiplied with a Gaussian lateral component. The simulation was performed for an excitation wavelength of 800nm, a voxel size of 10nm in the axial direction and 100nm laterally, and the noise applied to give a expected photon count of 30 in the brightest voxel of the image. Two counts of background noise were also added to the entire image.

Applying the phase extraction procedure described in section 8.3 gives the reconstructed phase shown in Fig. 8.6. The automatically selected control points are shown as crosses

---

[3]Corresponding to the multiplication of a sinc^2 detection component and a sinc^4 two photon excitation component. The use of sinc functions ensures that the axial PSF remains bandlimited and avoids any artificial resolution improvement.
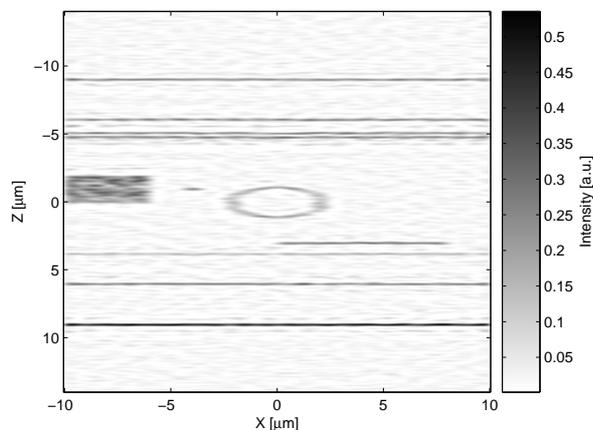
Figure 8.7: Reconstructed image
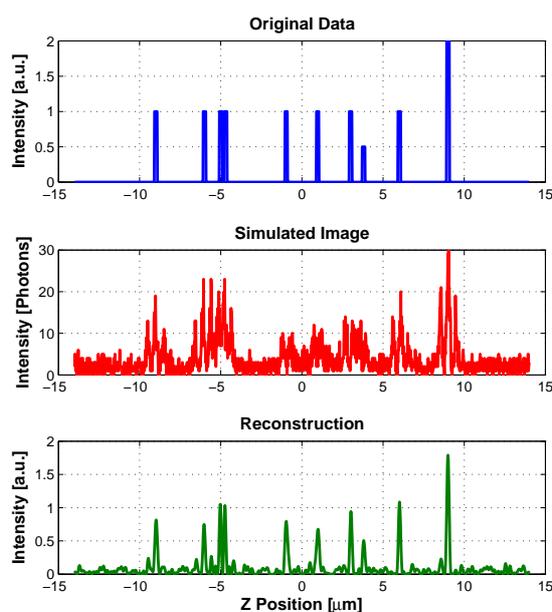
Figure 8.8: Line profiles through the test object, simulated image, and deconvolution result

on the image, and it can be clearly seen that, while not completely perfect, the reconstruction provides a good estimate of the original phase map. The discrepancies that are seen come from a variety of factors including interference from nearby structures, noise, and the interpolation procedure used (nearest neighbour interpolation followed by Gaussian smoothing, the width - $\sigma = 700$nm - of which was chosen to ensure a smooth phase change

	Correlation
Simulated data	0.5937
ML - no phase comp.	0.7187
ML - with phase comp.	0.9188

Table 8.1: Correlation of simulated data (Fig. 8.5) and reconstructions with (Fig. 8.7) and without (not shown) phase compensation to the test pattern (Fig. 8.3)used to generate the simulated data. Note that the values presented are at zero shift and that the correlation coefficients are normalised such that autocorrelation gives a coefficient of one.

between the reference points). It can be seen from the image reconstruction in Fig. 8.7 that it is still possible to obtain an accurate result despite these discrepancies. This image shows the effective phase compensation, sidelobe removal, and noise reduction offered by the restoration procedure. A detailed comparison is afforded by the line profile taken along the z axis close to the centre ($x \sim 0.5\mu m$) of the image through the test object, simulated data, and deconvolution result which is shown in Fig. 8.8. The deconvolution was performed both unweighted (ie. uniform Gaussian noise) and with the intensity dependant weights. As there was no appreciable difference in the quality of the resulting images, and as the exact noise characteristics of the 4Pi images are unclear, the results presented here use an unweighted solution - i.e. a uniform Gaussian noise model. In order to quantify the improvement obtained through phase compensation, the test object was compared with the simulated data and deconvolution results both with and without phase compensation by performing a cross correlation between each of the these images and our original test pattern. The results of these comparisons are shown in Table 8.1. The correlation value shows how similar each image is to the original object, with a value of 1 indicating 100 percent agreement. It is clear that phase compensating deconvolution offers a significantly better match than either the data or the deconvolution performed without phase compensation.

8.4.2   Reconstruction of Biological Images

In order to demonstrate the procedure on biological data, the procedure has been applied to a 3D image stack acquired with the 4Pi microscope from a specimen in which one class of potassium channels in the cell membrane were labeled through the expression of Kir2.1-eYFP (see Appendix for details of preparation and acquisition parameters). The deconvolution was performed on the smoothed data (gaussian blur, $\sigma = 1$ voxel), using a theoretical PSF envelope, and a phase map interpolated from 20 manually selected control points on the cell membrane. The PSF simulation was based on vectorial theory and assumed unpolarised light. The parameter selection was based on the experimental wavelength, voxel size, and pinhole settings and a refractive index of 1.46. In order to obtain the correct sidelobe height, however, a reduced NA of 1.1 was required[4]. Interpolation of the phase map was performed by fitting a linear function in z to the control points.

---

[4]Such a large reduction in NA seems unlikely - one possibility would be a misalignment of the two objectives.

Figures. 8.9, 8.10 and 8.11 show an xz section through the stacks corresponding to the raw data, the result of normal ML deconvolution (using our algorithm with the phase set to zero over the whole image), and the result of ML deconvolution when applying phase correction. In both cases the number of iterations was 20. Fig. 8.12 shows a line profile taken vertically through the centre of the images. The asymmetry in the PSF, and the refractive index mismatch induced phase change through the image are clearly visible. The improvement to the deconvolution obtained through phase correction is easy to see. The residual sidelobes in the deconvolved image can probably be attributed either to over-regularisation or to an inaccuracy in the estimation of the PSF Envelope. Nevertheless they are certainly low enough to facilitate an unambiguous analysis of the images.
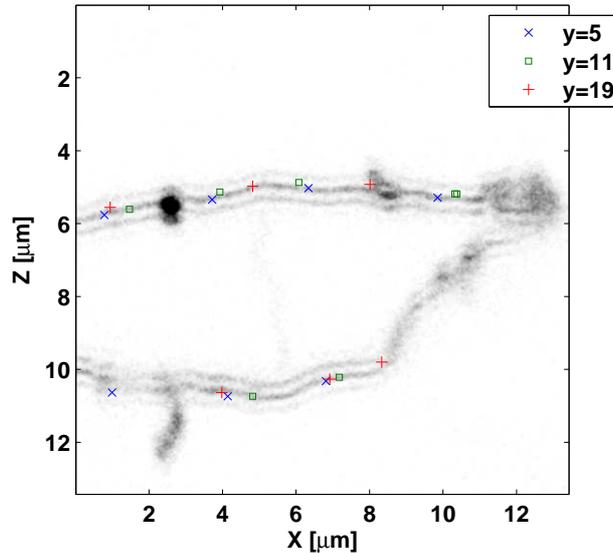
Figure 8.9: XZ-slice through raw Kir2.1 ion channel data after Gaussian blur for noise reduction

8.4.3 Performance of the Algorithm

In its current state, our algorithm is both relatively slow and relatively resource-hungry. On the test machine (AMD AthlonXP 2500+, 1GB RAM, Matlab R14, Windows 2000) the maximum data size that can be computed in one go is around 2×10^6 voxels. On a stack of this size, 30 iterations take around 45 minutes. The large memory usage can be explained by the fact that floating point representation is used for all the data and that it is necessary to keep track of several working variables, including ten (5 forward and 5 transpose) *part OTFs* corresponding to each of the convolutions in Eqn. 8.5, each of which is the full size of the image. The memory usage is thus around five times that of a similar deconvolution algorithm with a spatially invariant PSF. The slow speed is due to the relatively large number of calculations which must be performed at each step. This could be improved somewhat, at the expense of higher memory usage, by pre-calculating
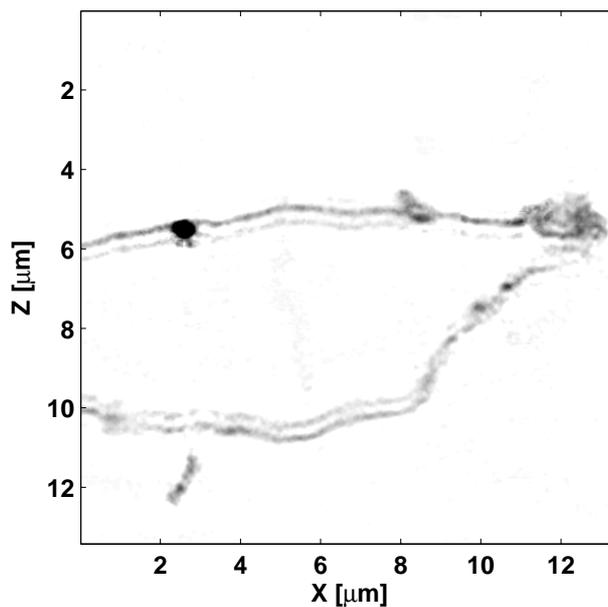
Figure 8.10: XZ-slice through Kir2.1 ion channel image reconstructed with our ML algorithm and the phase set to zero over the whole image. This represents what would be obtained from a conventional ML deconvolution algorithm.
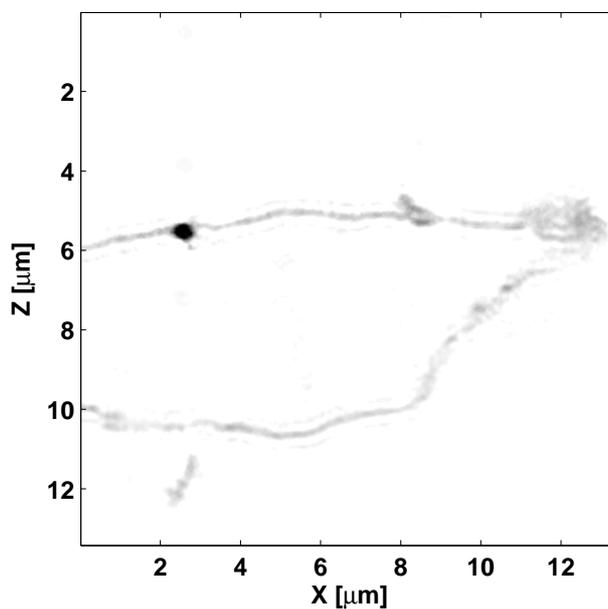
Figure 8.11: XZ-slice through Kir2.1 ion channel image reconstructed with our ML algorithm and a phase map interpolated from 15 control points selected on the cell membrane.
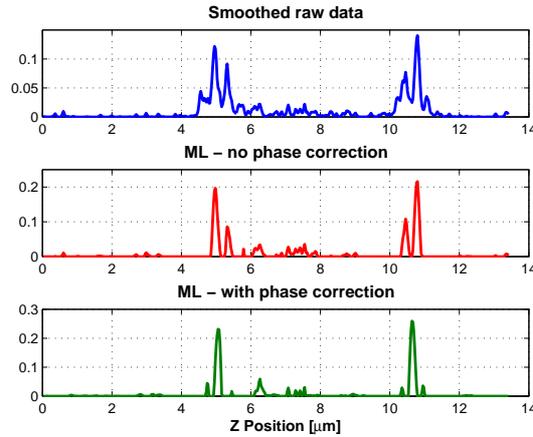
Figure 8.12: Line profiles taken vertically through the the raw data, and deconvolution results with and without phase correction.

the position dependent weights. The relatively large amount of computation involved in the forward (and transpose) transforms makes the fast convergence of a conjugate gradient solver particularly attractive (normally the computation of the forward transform is fast compared with the relatively large per step time required to calculate the new search direction in CG methods, and fast convergence is partly offset by the fact that another algorithm - e.g. steepest descent can make a lot more steps in the same time).

To overcome the memory related limitations on image size, it is possible to segment the image and deconvolve each segment independently. To avoid the edge effect artifacts that would otherwise be present, the segments can be made to overlap and the edge regions of each segment discarded. In theory, due to the $O(n\log n)$ complexity of the FFTs used to compute the forward/transpose mappings, such segmentation should also offer a performance gain although this is likely to be offset by the extra computation introduced by the segment overlaps. It should also be noted that a segmentation based technique lends itself well to parallelisation, falling into the class of embarrassingly parallel problems (those which can be split into many independent tasks with no need for data exchange except at the beginning and end of computation). We are currently developing an algorithm based on these principles, and preliminary tests show an execution time of $\sim$2hr on the machine above for 10 iterations on a $512 \times 370 \times 256 \approx 5 \times 10^7$ voxel image, and a linear decrease with an increasing number of nodes ($\sim$40 minutes on a 3 node cluster).

Part III

Experimental

Chapter 9

Refinements to SMI Hardware

9.1 Black and White Camera

In order to acquire multi-spectral data, the SMI microscope had previously used a colour CCD camera. This approach is not ideal, suffering from a poor effective quantum efficiency (the pigments in the colour mask have a relatively low peak transmission in the order of 50 - 70 percent, and only a quarter of the CCD active area is available to each colour channel). The colour separation between channels is also poor, and whilst marginally sufficient when using the combination of 488 and 647nm excitation, is somewhat wanting when 3 channel measurements are to be made. A black and white camera was thus adopted, resulting in an increase in the the signal level by a factor of approx 5 - 8. The black and white camera used was a PCO sensicam QE, having a peak quantum efficiency of $\sim$ 65 percent, an A/D conversion factor of 2 and a readout noise of around 4 counts.

9.1.1 Single Molecule Sensitivity

With this camera it should be theoretically possible to achieve single molecule sensitivity. A quick, back of the envelope, calculation is as follows:

The laser excitation at 488nm has a typical power, $P = 50$mW within the interferometer. The beam diameter, $2r \approx 2$cm. The area of this beam is thus $A = \pi r^2 \approx 3\text{cm}^2$, and the intensity $I_{\text{interferometer}} = P/A \approx 16\text{mw/cm}^2$. The focal length of the objective is 2mm, that of the lens used to focus into the back focal plane 40mm, giving a magnification $M = 20$. The intensity in the sample plane is thus $I_{\text{sample}} = M^2 I_{\text{interferometer}} \approx 6.4\text{W/cm}^2$. This corresponds to a photon flux of $\Phi = I/E = I/\frac{hc}{\lambda} \approx 2 \times 10^{19} s^{-1}$. The corresponding fluorescence is given by $F = \sigma Q \Phi$, which with an absorption cross-section $\sigma = 3 \times 10^{16}\text{cm}^{-2}$ (fluorescein) and a quantum yield $Q = 0.9$ gives $F \approx 6 \times 10^3 \text{s}^{-1}$. Assuming a collection efficiency of $\sim 30\%$ for a 1.4NA 100x oil objective and losses from the dichroic and blocking filter we can expect around 20% of the emitted signal to fall on the camera. With a camera quantum efficiency $\sim 50\%$ and an A/D conversion factor of 2 electrons per count we should see around 300 counts per second. Given a normal integration time of around 0.1s this would be 30 counts, which is significantly above the noise floor at 4 counts.

Bleaching must also be considered. Under good conditions (no /minimal free oxygen) Fluorescein can undergo something like 360000 fluorescent transitions before photodestruction[1]. This would correspond to a lifetime of approximately 60s under the given conditions. Whilst it is quite likely that in many circumstances background fluorescence will make single molecule detection unfeasible, the sensitivity is such that we cannot rule it out, particularly when considering the causes of non-specific labelling. Several successful single molecule experiments (Juergen Reynmann, Manuel Gunkel and others) have now shown that single molecule detection with the SMI is indeed possible.

9.2 Shutters

The simplest method of achieving colour discrimination when using a black and white camera is time multiplexing, whereby the different spectral signals are sequentially acquired. This can be achieved by switching the excitation, the detection, or both. Switching the detection on its own makes relatively little sense as the other stains will continue to be bleached. Switching both excitation and detection filters is the best solution when considering both crosstalk and signal throughput. It has, however, a few disadvantages when it comes to SMI measurements. The most obvious of these is that it is not possible to change the dichroic mirror without realigning the interferometer. It is possible to chance the blocking filter without realignment, however it is exceedingly difficult if not impossible to do this without introducing an image shift which would be significant for SMI distance measurements (error $\ll 10$nm).

The following compromise was thus reached - colours should be primarily separated through switching the excitation, using a multi-pass dichroic and blocking filter. The blocking filter is however in a filter wheel, allowing single band blocking filters to be switched in for consecutive measurements when the channels cannot otherwise be separated and the distance resolution is not so important. To minimise the effects of drift the excitation should preferably be switched between each slice rather than after a complete z-stack[2].

In order to switch the excitation, I implemented shutters in front of the lasers. One of these was a commercial (although somewhat old) shutter from Uniblitz, the other improvised using an old answering machine cassette deck and a small amount of electronics. Both are controlled through a manual switchbox and/or by computer through the parallel port. The introduction of multiple channels through shutters required the changes to the software described in Appendix A. Another possibility for implementing a shutter, as persued on the 2nd SMI microscope by Hans Mathee uses speakers to which a metal blade is attached. At the time of writing, the Uniblitz shutter is broken. It is intended to replace this and provide a shutter for the new 568nm laser based on the speaker principle.

---

[1]http://www.microscopyu.com/articles/fluorescence/fluorescenceintro.html

[2]This has proved to be less important than originally thought with the current setup showing very small (< 2nm) shifts between two consecutive stacks.

9.3   568nm Excitation

In order to be able to excite the Rhodamine family of dyes (and those with similar spectra) a 568nm Kr laser from Lexel was introduced. To accommodate this the dichroic beamsplitter and the blocking filter had to be changed to triple band versions (F62-568 and F63-488 resp., AHF Analysetechnik). A new dichroic beam combiner was also required. Tests on beads have given the expected results for size measurements. Measurements of biological specimens labelled in the red channel were also successful (replication foci in Cy3 and centromeric CENP-mRFP - Chapters 11 and 13) .

9.4   Transmitted Light Illumination

As with any fluorescence microscope it is useful to have an alternate form of illumination in order to locate and identify the cells to be measured. While it is possible to use a fluorescent counterstain, this reduces the number of available channels for the probes of interest and substantially increases the overall photon dose. An additional illumination should ideally require no staining of the specimen, and operate with relatively low light intensities. The systems conventionally used are phase contrast, differential interference contrast (DIC), or to some extent darkfield. Transmitted light systems are usually rejected due to the poor contrast for transparent specimens, and the supposed reduction in resolution at the low numerical apertures required to obtain any contrast at all with the aforementioned transparent specimens. The construction of the SMI microscope, like that of several other experimental microscopes, does not however lend itself well to any of the alternative schemes. It turns out that when coupled to a high dynamic range CCD camera, a low numerical aperture transmitted light illumination can provide a contrast for transparent specimens similar to that obtainable with DIC or phase contrast.

A form of low aperture transmitted light illumination was thus implemented for the SMI microscope. A super-bright 'white' Luxeon LED (part # LXHL-NWE8) with in-built collimating lens is used as the light source, producing a semi-collimated beam which is coupled into the microscope through the unused port on the beamsplitter (see Fig. 3.1). The beam path is then similar to that of the laser light with the major differences being that a large portion is passed by the dichroic and the blocking filter. This implies approximately collimated light between the objectives, coming largely from the left hand objective (that used only for illumination). The blocking filter in front of the LED is required to prevent the LED from fluorescing as a result of the laser light during the normal measurements. It additionally serves to block the light from the LED in the absorption bands of the fluorescent dyes, thus reducing photodamage from the LED source. As an alternative to a white LED and blocking filter it is also possible to use a single coloured LED in one of the emission bands.
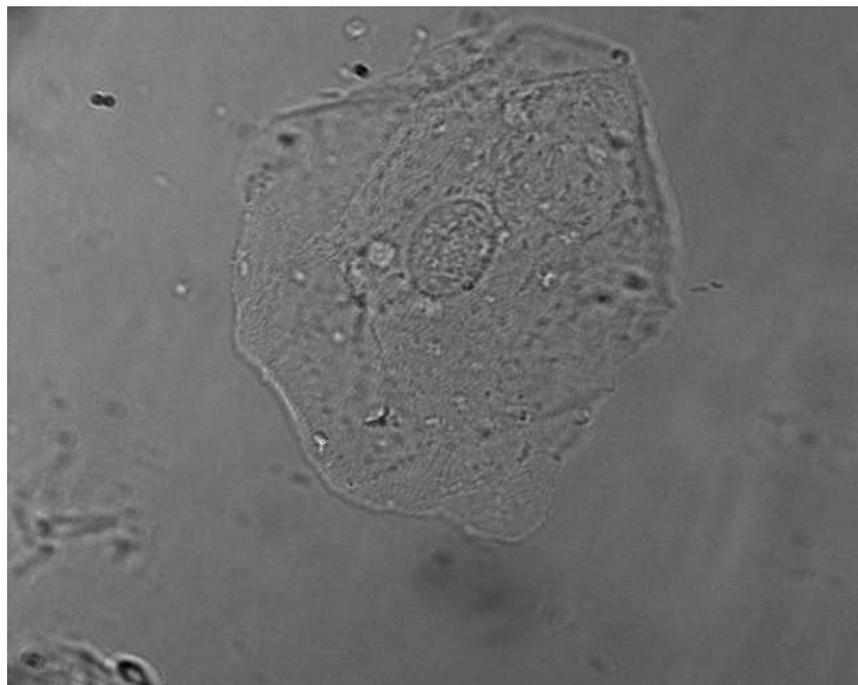
Figure 9.1: Unretouched image of an unstained human epithelial (cheek) cell taken using the SMI transmitted light illumination

9.4.1   Contrast

Figure 9.1 shows the unretouched image of an epithelial cell as displayed on the computer during the process of looking for cells. From this image it is clearly possible to identify the cells, and several internal features. The level of contrast seen here is significantly better than the through-ocular contrast obtained for a similar phase specimen on a conventional microscope, and subjectively similar to that obtained with DIC. This high level of contrast can be explained by the use of very low numerical aperture illumination combined with a CCD camera having a large dynamic range.

When using transmitted light, some contrast can be obtained from phase specimens by stopping the condenser numerical aperture down. This produces semi-coherent illumination, which when diffracted by the object can interfere to produce contrast. In a conventional microscope, this contrast is weak as the intensity in the diffracted orders is significantly less than that which is transmitted through the sample. This results in a very bright background against which the objects are difficult to discern. Various methods such as darkfield and phase contrast overcome this problem by ensuring that the undiffracted illumination light is not imaged onto the detector.

If a uniformly high background is present, the use of a CCD camera opens another possibility for removal of the background intensity - computational subtraction. Such a subtraction is simply performed; in our case the control software already performs a

background subtraction and scaling in order to map the full range of values in the 12bit camera data to an 8bit grey value for display. The mapping used has the form

$$V_{8\text{bit}} = 256 \times \left[\frac{V - \min(V_{12\text{bit}})}{\max(V_{12\text{bit}}) - \min(V_{12\text{bit}})} \right]$$

This simple mapping, currently performed automatically in the control software, has shown itself to be sufficient for all current applications. In theory the contrast could be increased further by applying a high pass filter, or devising an algorithm to mitigate the influence of dirt in the preparation which can cause the background to be non-uniform.

9.5   Minor Changes to Setup and Adjustment Procedure

Several minor changes, mostly simplifications, have been made to the SMI setup. They are quickly summarised below:

9.5.1   Switch to Cube Beamsplitter

The previous wedge beamsplitter was exchanged for a non-polarising cube beamsplitter. This has the advantage that all the rays are at 90 degrees to another, considerably simplifying alignment. It also removes the need for vertical polarisation and the spectral sensitivity of the previous beamsplitter. The resulting experimental layout is that shown in Fig. 3.1.

9.5.2   Introduction of Iris Diaphragms Between Beam Expander and Interferometer

Two co-axial iris blends were introduced in the beam path between the beam expander and the interferometer. These allow the beam to be stopped down to reduce the area within the sample which is illuminated and thus bleached during data aquisition. When used in this function, the second blend can be used to block the some of the light diffracted from the edges of the first. A second, and equally important role for the two blends is to decouple the combination and expansion of the laser beams from the interferometer. After adjusting the beam combination it usually suffices to steer the beam from the beam expander such that it passes through both stopped down irises to reattain interferometer alignment.

9.5.3   Removing all Equipment with Fans from the Optical Table

After introducing the black and white camera the increased sensitivity allowed shorter acquisition times which revealed vibration artifacts. Removing all equipment with cooling fans from the optical table improved the situation.

9.5.4   Moving from Fibre to Free Space Laser Coupling

This resulted in a significant increase in available laser power and a small improvement in beam profile. It was also combined with a switch to beam combination before expansion, considerably reducing the component count. Using achromatic doublets it was possible to provide good collimation of the beams at all three wavelengths simultaneously.

9.5.5   Shifting the Lasers to a Separate Optical Table

In response to unacceptable vibrations derived from the water cooling of the pump laser used with the 2-photon experiment, the decision was made to shift all the lasers onto one table, with the two experimental setups sharing the other table. The two tables were placed approximately 20cm apart and the lasers coupled in free space between them. This arrangement has proved to be remarkably stable, with no adverse effects being observed due to lack of mechanical attachment between the two tables[3]. Indeed the performance of SMI1 as well as that of the 2-photon experiment showed an improvement. Since both fan cooled devices and lasers were moved off the optical table the correction factor introduced by Christian Wagner for unequal beam intensities has been very low and relatively constant between measurements, suggesting that the real explanation for the effect was the mathematically equivalent vibrational blurring of the interference pattern.

In retrospect, when one considers the quadratic dependance of intensity on electric field strength, the relatively high ($\sim .3 - .4$) correction factors measured are difficult to explain through an intensity difference alone. Assuming the beam from one side of the interferometer produces an electric field in the object space of $E_1 = E\cos(kx - wt)$, that from the other side $E_2 = (E + \Delta E)\cos(-kx - wt)$. The resulting electric field is thus

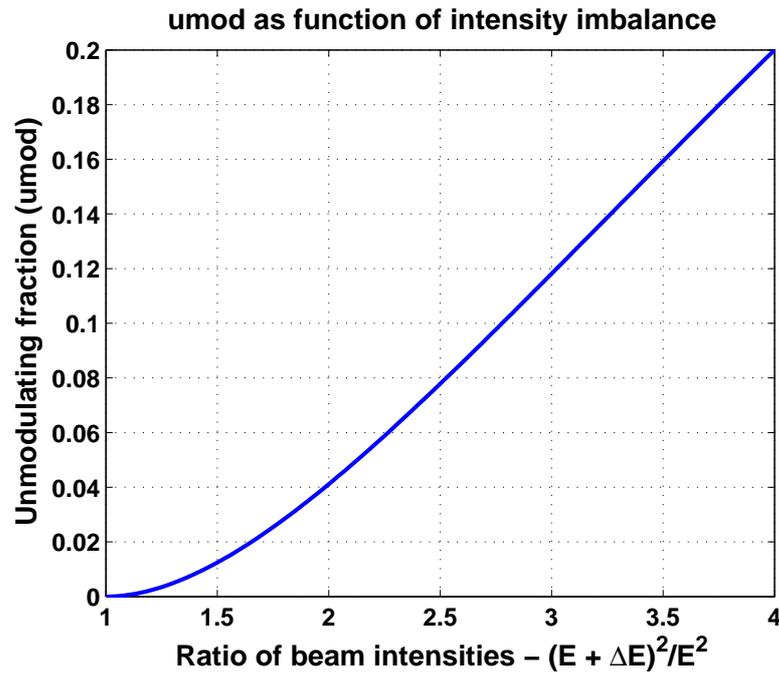$$E_3 = E_1 + E_2 = 2E\cos(kx)\cos(wt) + \Delta E\cos(-kx - wt).$$

The intensity is thus:

$$
\begin{aligned}
I &= \frac{1}{T}\int_0^T (E_3(t))^2 dt \\
&= \frac{1}{T}\int_0^T 4E^2\cos^2(kx)\cos^2(wt) + 4E\cos(kx)\cos(wt)\Delta E\cos(-kx - wt) + \Delta E^2\cos^2(-kx - wt)dt \\
&\propto 4\left(E^2 + E\Delta E\right)\cos^2(kx) + \Delta E^2
\end{aligned}
$$

$$(9.1)$$

From which the unmodulated portion can be extracted;

$$\text{umod} = \frac{\Delta E^2}{4\left(E^2 + E\Delta E\right) + \Delta E^2}.$$

---

[3]It is likely that this is only the case as, due to our fortuitous position in the basement, it is possible to operate the tables without compressed air isolation

Figure 9.2: Dependance of unmodulating fraction on the relative intensities of the two laser beams.

When converted to a relationship between intensities, the graph shown in Figure 9.2 is obtained. A ratio between the intensities of the two beams of more than approximately 1.5 times is highly unlikely in practice. The reason intensity differences have so little effect on the interference pattern contrast is the E^2 dependance of intensity on electric field strength. The E^2 dependance is effectively felt twice, with differences in I corresponding to much smaller differences in E, and through the interference where a significant difference in E produces an even larger difference in intensity.

Chapter 10

Bead Measurements

The aim of this chapter is to give a brief overview of what is possible in terms of bead measurements using the SMI microscope and analysis algorithms developed in this thesis. Much more detailed measurements with several different bead sizes have been presented in several previous works; the theses from Hans Mathée [71] and Juergen Bastert [7] show such measurements with the current setup and using the new analysis algorithms. Similar results, albeit on the *mirror* and *vertico* setups are detailed in [72] and in the thesis from Juergen Reymann (to be completed). A detailed study of the behaviour of the SMI with different bead sizes is thus omitted here, and a couple of example measurements using 110nm TransFluoSphere beads from Molecular Probes[1] (Cat num. T-8878) are presented. These beads are very bright, have the unique property of being able to be excited at all three laser wavelengths currently implemented in the SMI and show the majority of their emission in the 'deep-red' region of the spectrum. In order to achieve this behaviour, a combination of dyes which are optimised for Fluorescence Resonance Energy Transfer (FRET) are used. These beads are significantly brighter than TetraSpec[2] beads and allow the comparison of the properties at the different excitation wavelengths independently from any chromatic shift present in the detection.

10.1 Reproducibility

In order to quantify the reproducibility, and hence the intrinsic measurement accuracy attainable with the SMI microscope and the analysis techniques used, repeated measurements of the same objects were made and the measured sizes and positions compared. It should be noted that such reproducibility measurements are not affected by any systematic errors such as might be produced by tilted or otherwise aberrated wavefronts.

An example measurement and the resulting analysis is shown in Fig. 10.1. A good agreement with the theoretical diameter is obtained with an *umod* of 0. The resulting distribution is somewhat broader than is typically the case, or more specifically shows a

---

[1]now Invitrogen - www.invitrogen.com
[2]also from Molecular Probes/Invitrogen
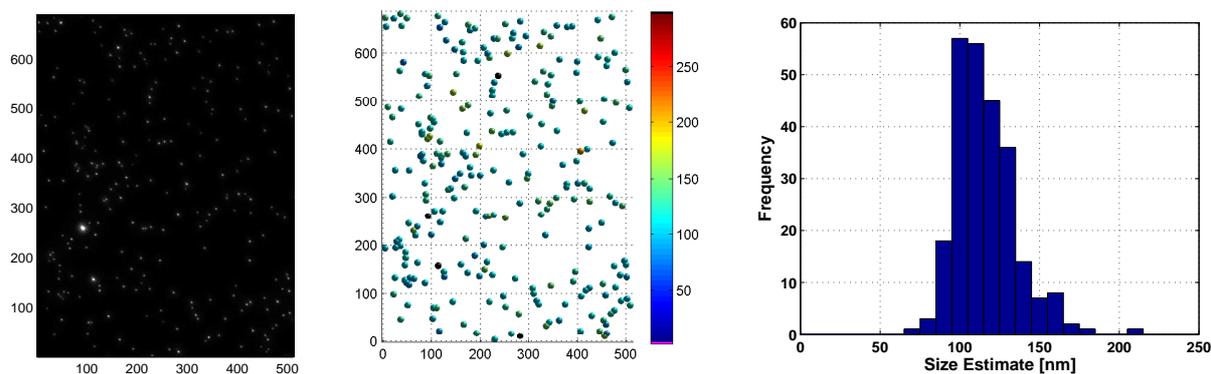
Figure 10.1: Left: Maximum intensity projection from SMI image of 110nm beads taken at 488nm excitation. **Middle:** Evaluated objects with size estimate colour coded. **Right:** Histogram showing the distribution of measured sizes.


Figure 10.2: Difference in the measured size of 110nm beads in two consecutive acquisitions at 488nm.
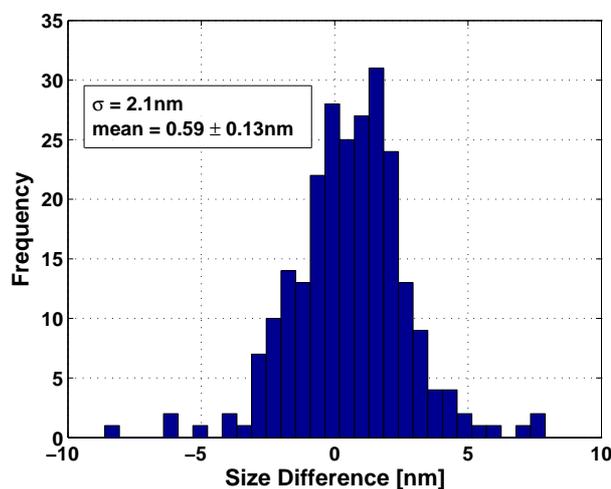
tail towards the large size end of the spectrum. This can most probably be attributed to the presence of bead clusters, a hypothesis that is also supported by the large range of point intensities[3]. Note that such a tail towards larger sizes can also be produced by non-uniform intensities (and hence modulation contrasts) produced by misalignment. Is this the case, a trend in the sizes across the field, or large regions of different size are generally apparent in the colour coded size representation.

Despite the presence of some clustering, when the same region is measured a second time the size estimates obtained are very close to those in the first measurement. The differences between two such measurements are shown in Fig. 10.2. The distribution of these differences is indicative of the statistical measurement accuracy (in this case approximately 2nm). Surprisingly the mean value of 0.59nm is significantly different from zero (standard

---

[3]The obvious large clusters at $\sim (100, 260)$ and $\sim (110, 160)$ were not fitted (black points) and thus are not represented in the size distribution.
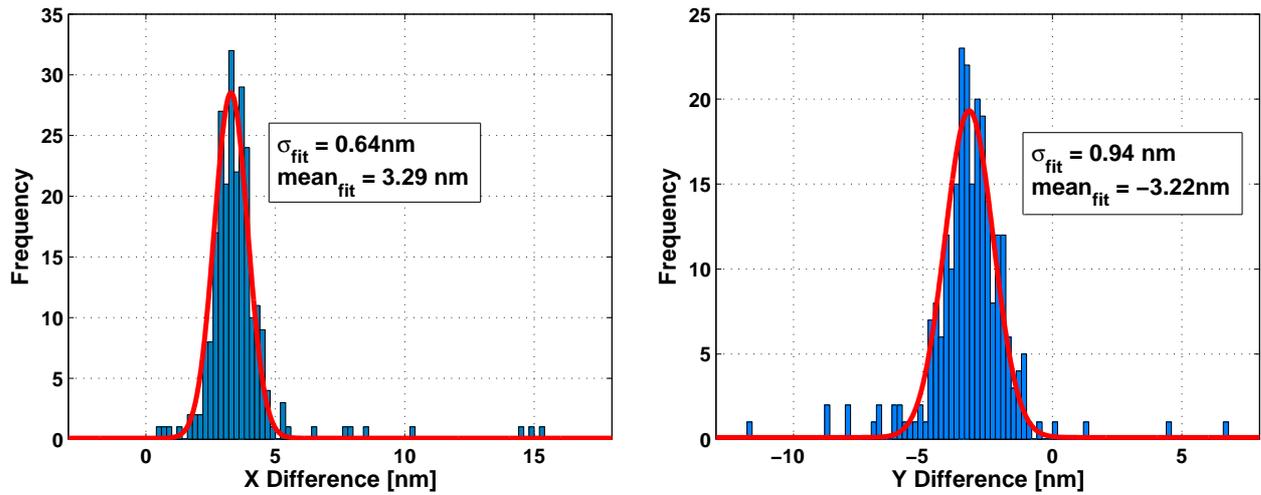
Figure 10.3: Left: Difference in x position between two consecutive measurements as determined by the Gaussian fit. **Right:** Difference in y position.

error of mean$= \sigma/\sqrt{n} = .13$nm). It is unclear what the source of this *shift* in the size estimates is; several mechanisms are possible. The easiest explanation is that the amplitude of some external source of vibration changed between the two measurements, resulting in a small change in *umod*. Another possibility is that the effect is the result of small mechanical drifts (eg: in the phase of the interference pattern) which, when coupled with the approximations (that the widefield PSF has a negligible influence on the modulation depth) implicit in the size estimation process could result in slightly different size estimates. In any case the difference is small. It is however something to bear in mind when thinking about interpreting small (but statistically significant differences) in biological data.

In addition to size measurements, the precision of position measurements is also important. Figure 10.3 shows the differences in the measured lateral positions of the same objects. The mean and standard deviation of these differences were measured by fitting a Gaussian to the histograms. This shows that sub-nanometre accuracy is achievable through use of the lateral Gaussian fit. The mean differences of ~ 3nm can be attributed to instrumental drift. Similarly, as with widefield or confocal measurements, the axial position can be calculated be measuring the centre of intensity (a COI based measure is used for the widefield z-position estimate because the possible asymmetry in the axial profile make a functional fit difficult - see Chapter 7) giving the distribution in Fig. 10.4. The tri-modal distribution is the result of the finite voxel size and that a threshold is required to determine which voxels are included in the COI calculation. This problem can be ameliorated by clipping the images at the threshold [44], this clipping is however absent from the current algorithm[4].

As discussed in Chapter 7, the COI based axial distance measurements can be further

---

[4]When using a clipped COI algorithm the distribution is smoother than in Fig. 10.4, but the overall width is virtually unchanged ($\sigma_{\text{normal}} = 10.3$nm, $\sigma_{\text{clipped}} = 10.8$nm)

Figure 10.4: Difference in the measured z position of 110nm beads in two consecutive acquisitions at 488nm using a *Centre of Intensity* (COI) based method (ie. without using the information present in the modulation). Note that the accuracy of this localisation must be better than $\pm$40nm to allow SMI distance measurements to be made.
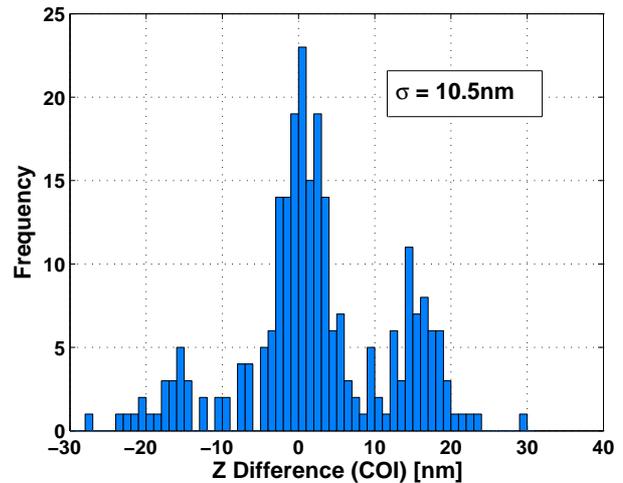
refined using the phase information out of the fit to the axial profile. The resulting distribution in Fig. 10.5 shows a considerable improvement over the COI based measurements. Because the algorithm currently used finds the nearest fringe position to the COI measurement, an ambiguity is present when the relative phase of the interference is such that a minimum occurs in the centre of the profile. This ambiguity results in the two residual peaks in the distribution at $\pm$160nm when the measured positions are subtracted from each other. For true distance (not position) measurements the ambiguity is resolved. It is also resolved for position measurements when the phase changes across the image are measured and corrected for.
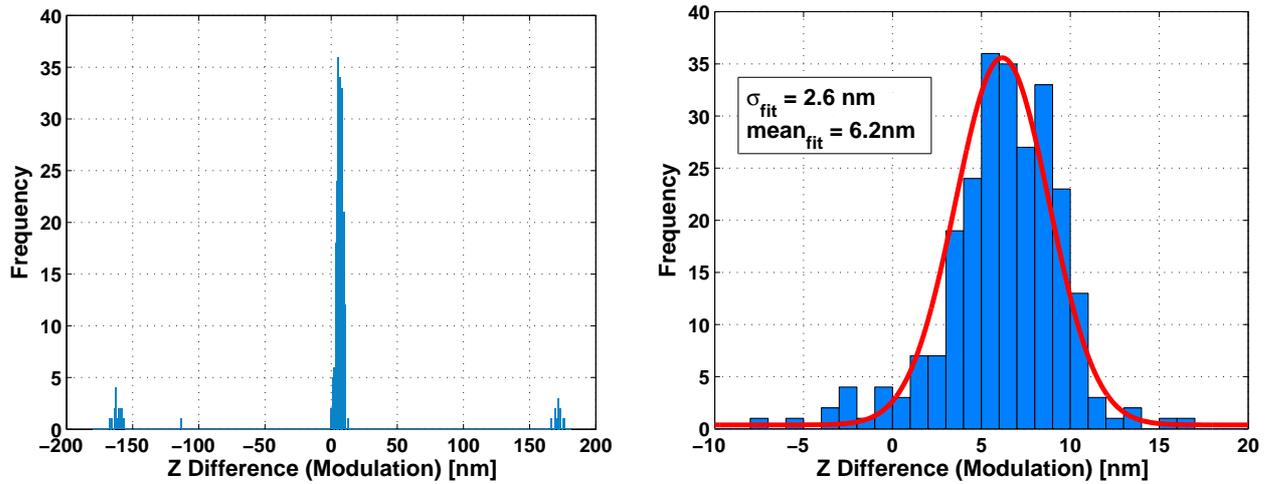
Figure 10.5: Left: Difference in z position between two consecutive measurements using the position of the fringes to refine the COI estimate. The outlying peaks at ± 160nm are due to the fact that no wavefront correction was made, and correspond to beads on which the phase is approximately destructive. **Right:** The distribution obtained when the outliers are folded back to their correct position, indicative of true measurement accuracy when correction for the wavefronts is undertaken.

10.2 Correction for Phase and Focal Plane

10.2.1 Wavefront Tilt and Curvature

Using the phase information in the fit it is possible to obtain a very precise estimate of the distance between an object and the interference fringes. Unless the fringes are absolutely flat and parallel to the focal plane, this does not directly translate into an improvement in localisation accuracy. It is very difficult, if not practically impossible, to obtain wavefronts which are both flat and parallel at all operational wavelengths through mechanical alignment. Even when no fringes are visible over the field of view some residual tilt or curvature is always present. Fig. 10.6 is an example of what is typical for a moderately well aligned SMI. At 488nm excitation (not shown) the field is approximately flat, at 647nm there is a small angle to the illumination resulting in the three stripes across the field of view.

Luckily we can make sensible distance measurements without perfectly aligned wavefronts - all we need to know is **how** they are aligned, and use this knowledge to correct our measurements. In order to establish this we need a relatively dense field of beads on a flat reference surface (as in Fig. 10.6), fitting all the beads allows us to extract sufficient information to characterise the wavefront behaviour. Before we use this information to correct our normal distance measurements, however, we would like to interpolate and smooth the data such that we have information for areas between the bead positions and remove some of the variability. Fig. 10.7 shows first efforts in this direction. The phase
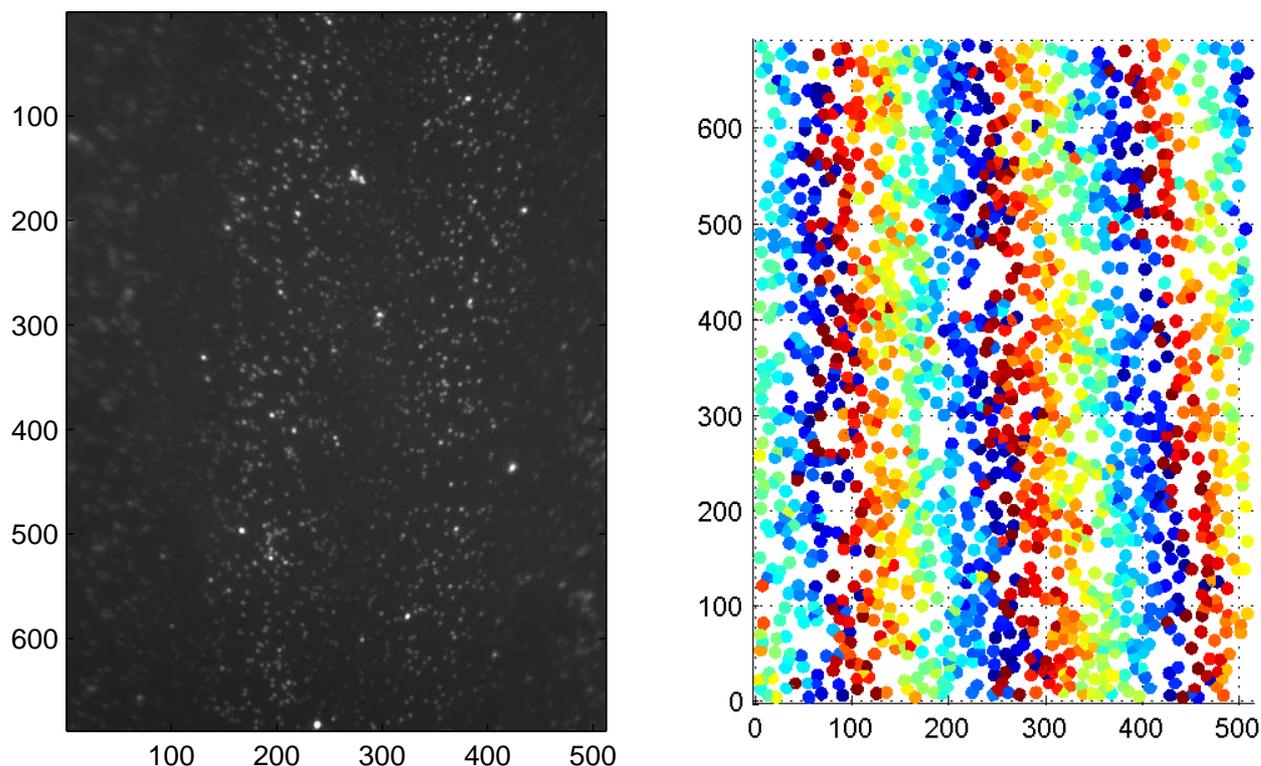
Figure 10.6: Left: In focus image of 110nm beads at 647nm excitation showing three fringes visible across the field of view. **Right:** The phase extracted from the fitted point positions.
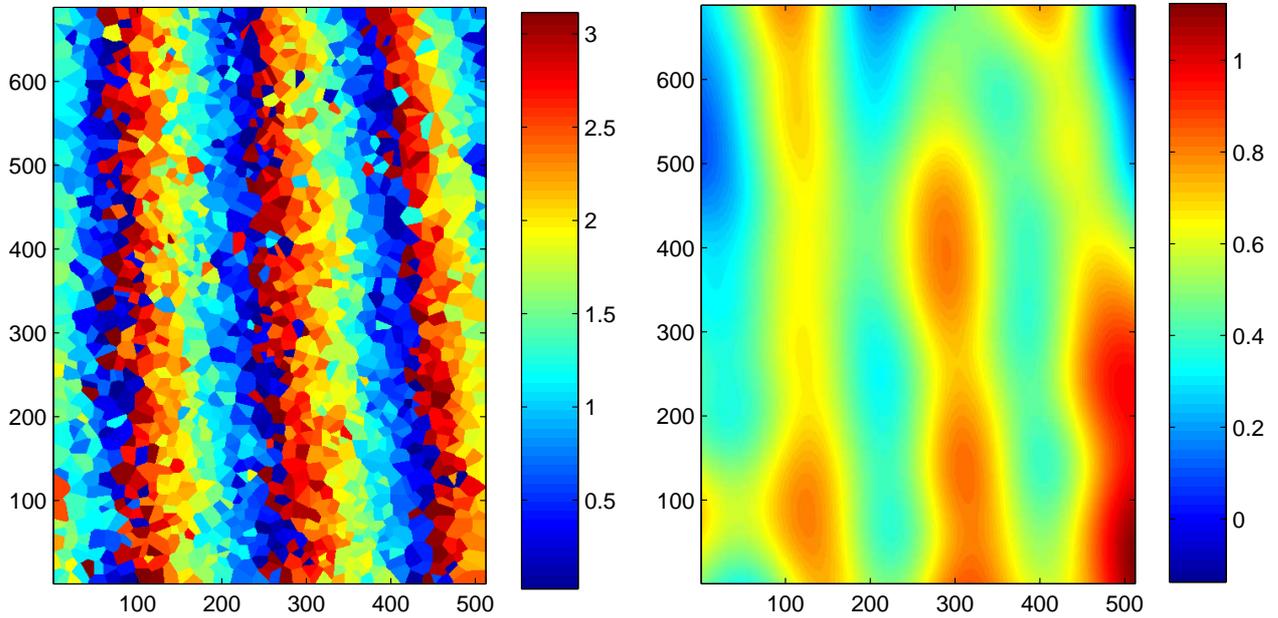
data is interpolated using nearest neighbour interpolation[5], padded to avoid wrap around artifacts and blurred with a Gaussian blur.

In Fig. 10.7 the phase map after interpolation is reasonable enough, that after the blurring however a complete disaster. The reason for this is that we only obtain the phase information modulo π, making a linear phase change across the field of view into a sawtooth pattern as in Fig. 10.8. When we blur over the phase discontinuities thus produced it is not surprising that the result is rubbish. It is thus necessary to remover the phase jumps before performing any smoothing.

Phase Relaxation: Or How to Remove the Jumps

The method I've used to remove the phase jumps is an iterative procedure, roughly depicted in Fig. 10.8. A misfit function is computed in which the phase at each bead location is compared to that of its neighbours. If the phase value is significantly higher than that of the surrounding objects π is subtracted, and the process iterated. The algorithm used

---

[5]One could also use linear or cubic spline interpolation - if we're going to smooth afterwards there is however little advantage to doing this. Nearest neighbour interpolation is faster and provides values near the boundary which would not be possible with a more sophisticated interpolation scheme.

Figure 10.7: Left: Interpolated phase map using nearest neighbour interpolation. **Right:** The result obtained when the phase map is smoothed. Note that due the discontinuities inherent in the phase relationship this phase map has very little to do with the real wavefronts.
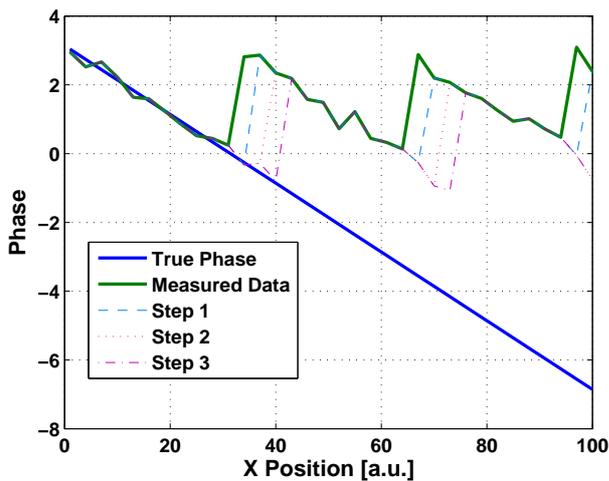
Figure 10.8: A tilted wavefront corresponds to a linearly changing phase across the image. As SMI phase measurements have a modulo π ambiguity, the data obtained is a sawtooth pattern. In order to reconstruct the original phase map the relaxation algorithm iteratively subtracts π from points which are significantly higher than their neighbours. This tends to 'walk' the phase discontinuities out of the image.

has a couple more features such as slowly decreasing the threshold used to decide if a difference significant is (ala simulated annealing), several numerical fudges to try and improve stability, and a mechanism for deciding when to stop. The application of the phase relaxation algorithm to the bead measurements from Fig. 10.6 is shown in Figures 10.9 and 10.10, with the smoothed result shown in Fig. 10.11.

The algorithm in its current state works relatively well on dense beads samples with a moderately shallow phase gradient. If the gradient is too steep or the sampling too sparse, the algorithm fails. When the algorithm fails, one can get periodic behaviour, fractal-like growth of phase steps from defect centres and various other very picturesque but otherwise useless patterns. There is certainly scope for improving the algorithm, but at present a visual control of the end result is highly recommended.

Other Ways of Removing the Phase Jumps

It is also conceivable that one could remove the phase jumps by fitting a function, for instance a plane, to the phase data. This is not as easy as it sounds because most fitting techniques don't particularly like discontinuities in the gradient, let alone the function itself. In the case of the modulo operator, the goal function has several minima, only one of which is right. One way to get around this is to fit a set of parallel functions spaced along the phase axis at intervals of π. This requires an error function which calculates the distance to all lines and takes the minimum. This is also a bit messy due to the fact that *min* is technically a non-linear function, and not necessarily something you want to have in an error function. Nonetheless minimising an error function constructed along these principles has produced sensible results in 4Pi phase evaluation, and should also work for the SMI data.

Assuming that one has a way of coping with the multiple value problem with the data fitting, it is interesting to consider what form of function one should fit. A linear change in one or both axes would seem obvious, but does not help us much when the wavefronts show curvature, as is often the case. It seems likely that the Zernike polynomials[108, 105] might provide a suitable basis for describing the SMI wavefronts, and a scheme based on fitting these warrants further investigation.
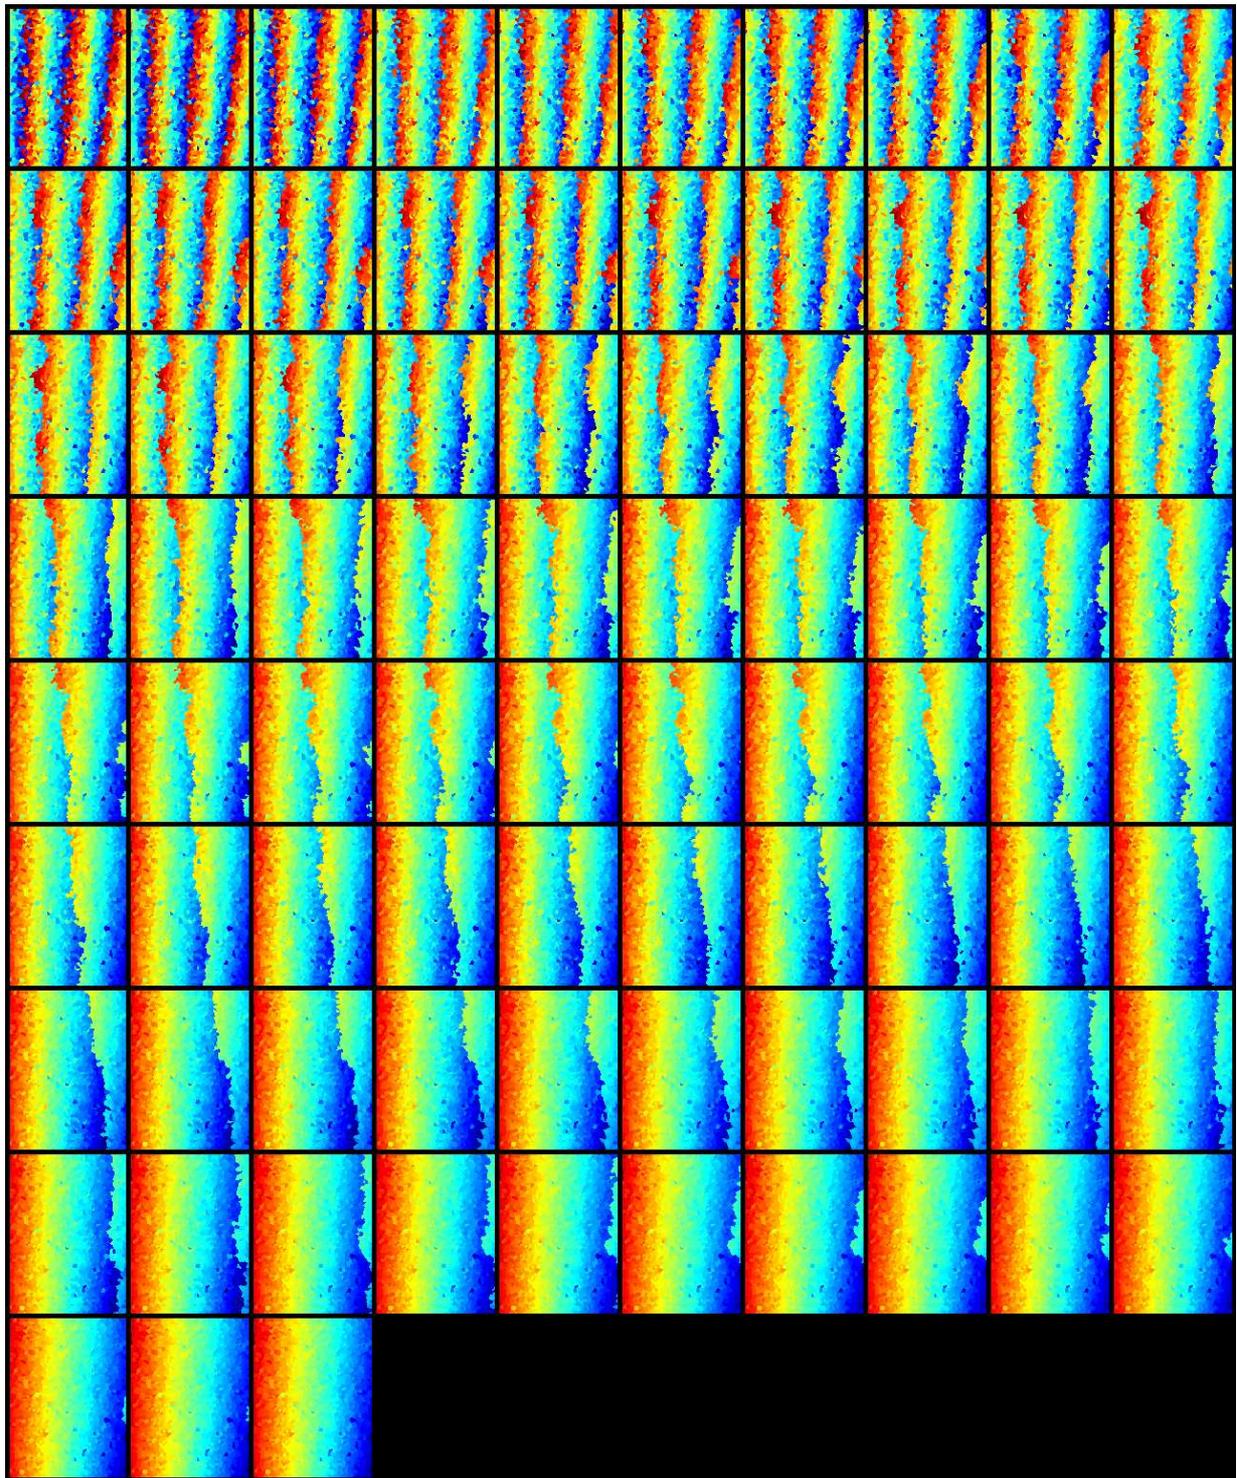
Figure 10.9: Phase progression during phase relaxation process.
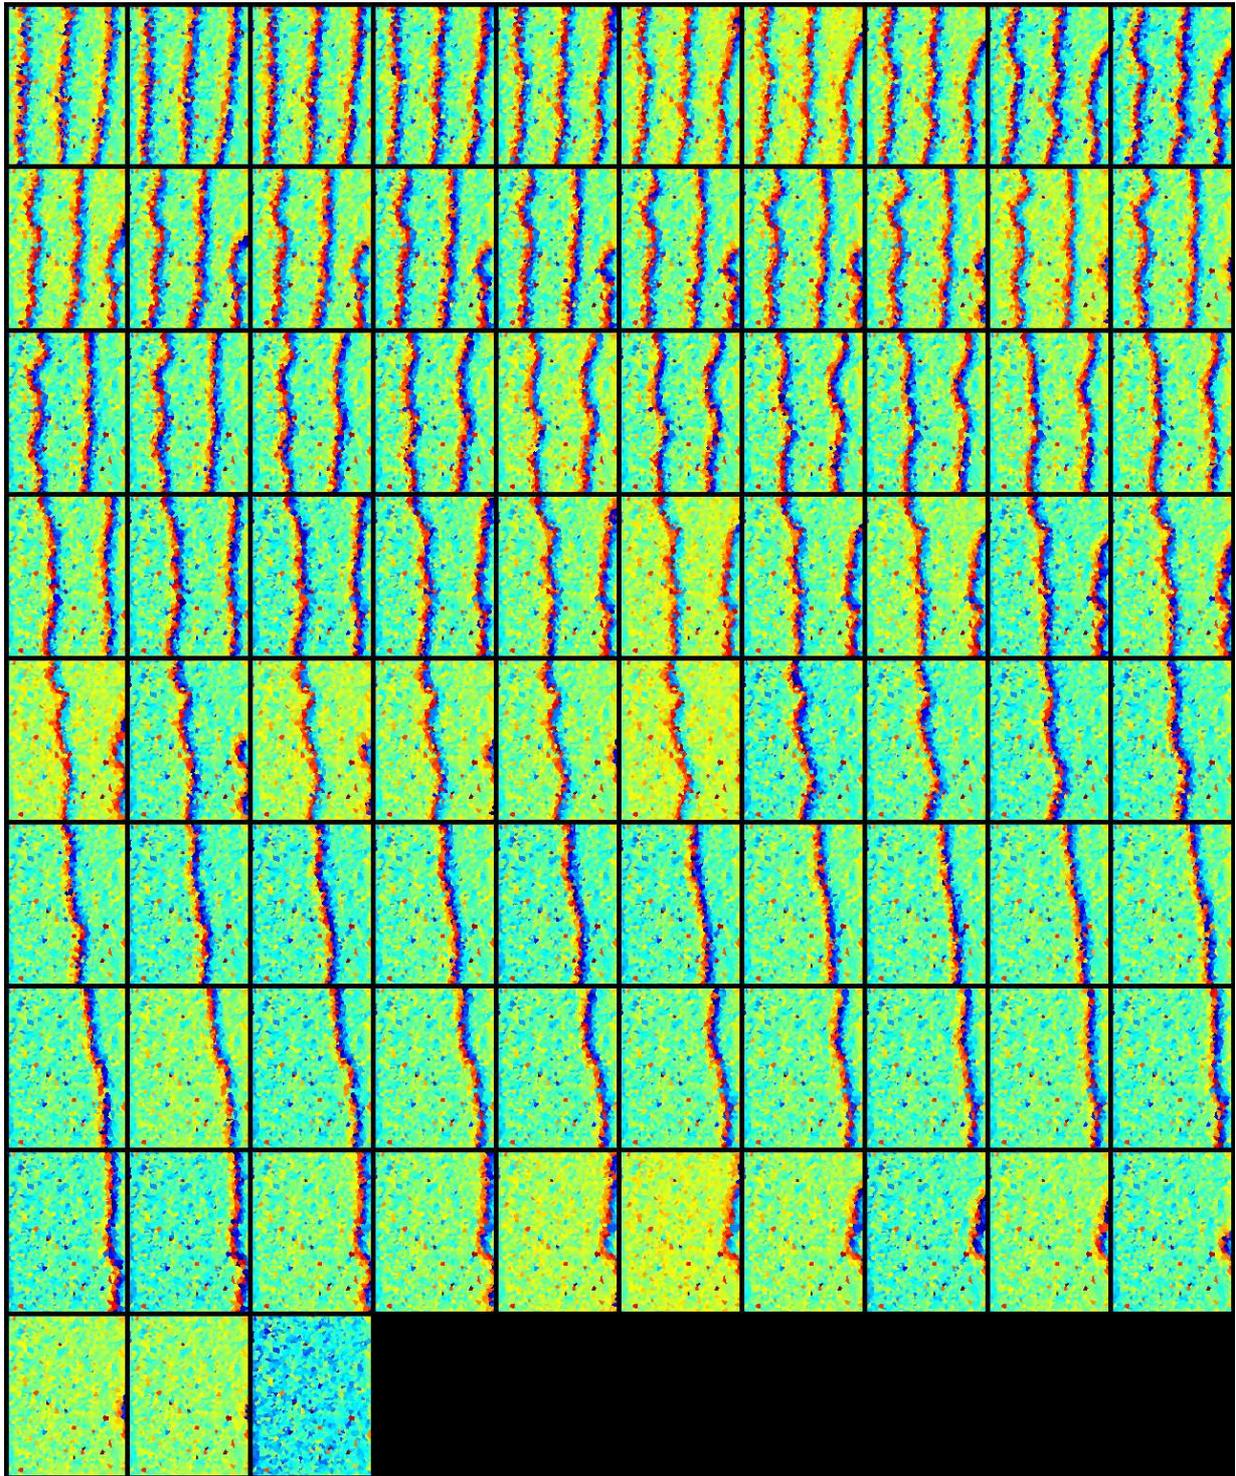
Figure 10.10: Progression of neighbour mismatches during phase relaxation process.
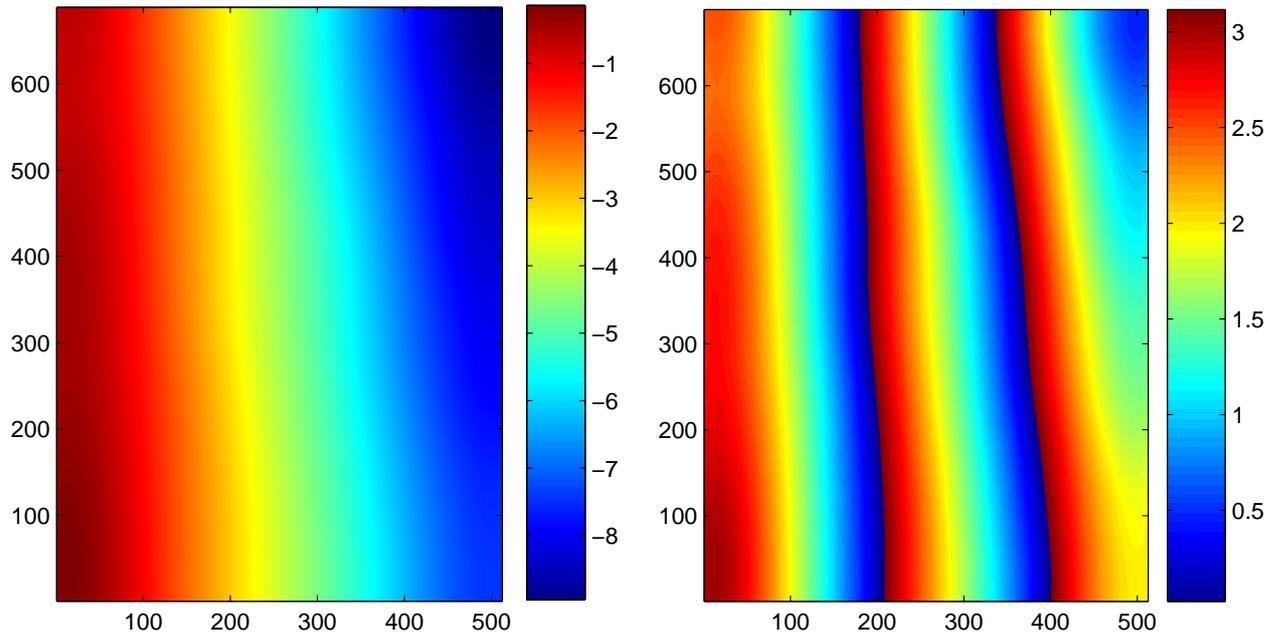
Figure 10.11: Left: Smoothed phase map after phase relaxation process. **Right:** Smoothed phase map modulo π.

10.2.2 Focal Plane Tilt and Curvature

It is difficult to align the slide such that it is perfectly parallel to the focal plane, and, even for *Plan-APO* objectives, the focal plane typically shows some residual curvature. This tilt and curvature can be compensated in a similar way to the wavefronts by measuring the apparent z-positions of a field of beads spread on a flat surface. Unlike wavefront correction no phase ambiguity exists so one may simply interpolate and smooth. An example of a focal plane measurement is shown in Fig. 10.12.

Sviewer Scripts

After normal analysis using `ofind3`, `dofits` etc., the following commands are useful for determining the wavefront and focal surfaces from a beads sample.

`phasesurf`: plots and interpolates the extracted phase information without any unwrapping.

`relax_phase`: performs iterative phase unwrapping to obtain a surface without discontinuities. Relies on some variables calculated by `phasesurf` - call `phasesurf` first.

`focalsurf`: plots the extracted z positions and interpolates these. Call `phasesurf` first.
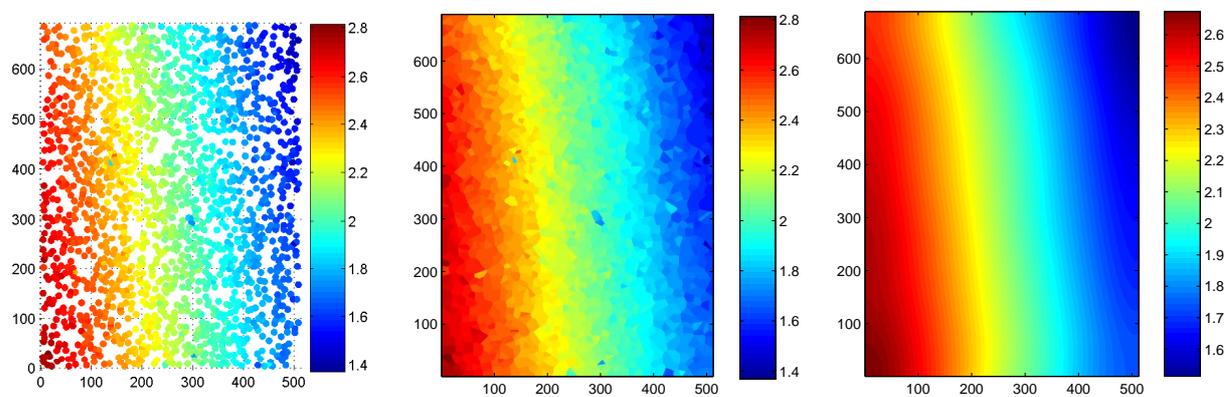
Figure 10.12: Focal plane measurement on 110nm beads: **Left:** Measured z-positions of beads. **Middle:** Interpolation. **Right:** Focal surface after interpolation and smoothing. The colour scale is in μm.

# Chapter 11

Replication Foci

One of the principle biological applications of the SMI microscope during my thesis was the study of the replication machinery in mammalian cells. Replication of the cellular DNA is performed in the S-phase and occurs simultaneously at several (thousand(s)) loci within the nucleus. Previous confocal studies indicate a change in the distribution and number of foci throughout the S-phase. A trend towards larger compound structures is also observed in late S-phase [30]. The time at which particular areas of DNA replicate has been shown to be correlated to gene content and activity with active genes being replicated early and constitutive heterochromatin late [2]. The compound structures seen in late S can be separated into several smaller constituent loci through the application of a hypotonic shock or, to some extent, by serial sectioning. The conglomeration of smaller loci can also be seen when using electron microscopy [61] or advanced light microscopy methods such as OMX[1]. It is however likely that these smaller loci each consist of several replicons[2].

The study in which I was involved[3] focussed on the size of these sub-loci, particularly with the aim of establishing whether a change in size occurs throughout the S-phase. To this end both the replication machinery and the newly replicated DNA were labelled. All experiments were performed in C2C12 mouse myoblast cells. Preliminary experiments indicated that SMI measurements in 3D conserved cells were not possible due to the large number of foci and the lack of optical sectioning. In order to obtain optical sectioning, either cryosections or hypotonically flattened cells were used. For details of the biological preparation see Appendix B.
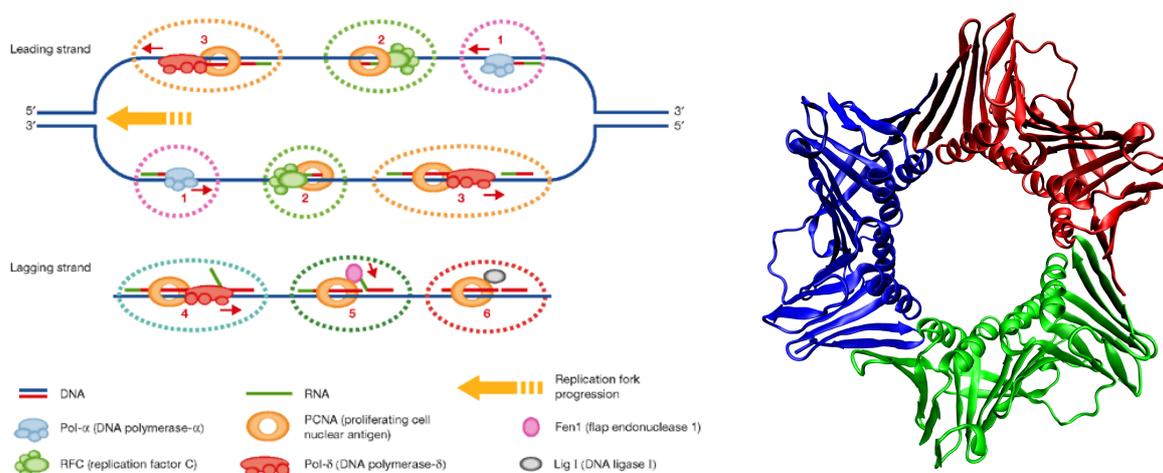
Figure 11.1: Left: Eukaryotic DNA replication. *Step 1: primer synthesis by DNA polymerase-a (Pol-a); step 2: replication factor C (RFC) displacement of DNA polymerase and recruitment of proliferating cell nuclear antigen (PCNA); step 3: elongation by the newly recruited DNA polymerase-d holoenzyme (Pol-d)* [52]. **Right:** Crystal structure of the PCNA trimer (*image source: http://en.wikipedia.org/wiki/PCNA*)

11.1   GFP-PCNA labelling

In the first experiments the Proliferating Cell Nuclear Antigen (PCNA) protein was conjugated to GFP. PCNA [75] is localised at replication foci and is thought to form a doughnut shaped trimer surrounding the DNA to be replicated and to serve as a sliding clamp / loading platform for other proteins (eg: DNA Polymerase) involved in the replication process. A schematic of the replication process is shown in Fig. 11.1, along with the X-ray crystal structure of the PCNA trimer.

Cyrosections were made from cells expressing the GFP-PCNA labelling, and could be detected satisfactorily in the SMI microscope. The image of a typical cell is shown in Fig. 11.2 and the size distributions obtained in Fig. 11.3. There is very little difference between the sizes at different S-phase positions. Due to the relatively poor signal to noise ratio, only around 50% of the foci could be analysed.
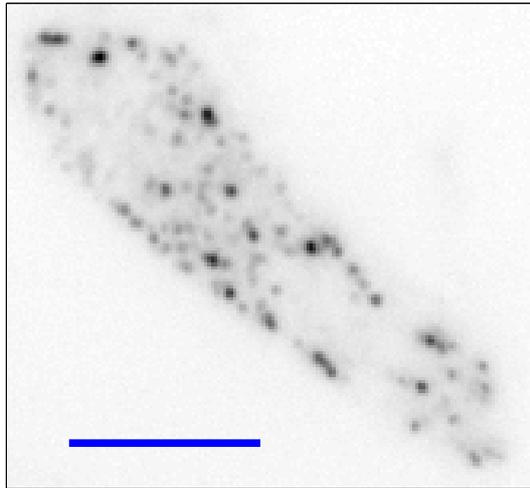
11.2   GFP-PCNA with antibody amplification

In order to improve the signal levels the GFP-PCNA labelling was amplified by incubating with Alexa488 labelled antibodies against GFP. This resulted in a slight improvement in

---

[1]Priv. Comm. L. Schermelleh

[2]Priv. Comm. C. Cardoso

[3]a collaboration with C. Cardoso, V. Chagin, A. Gahl - MDC Berlin and A. Pombo, S.Martin - MRC London

Figure 11.2: In focus slice from typical SMI image of the GFP conjugated PCNA component of the replication machinery in a cryosection through a C2C12 mouse myoblast. The scale bar is 5μm, and the gray scale inverted for better contrast.
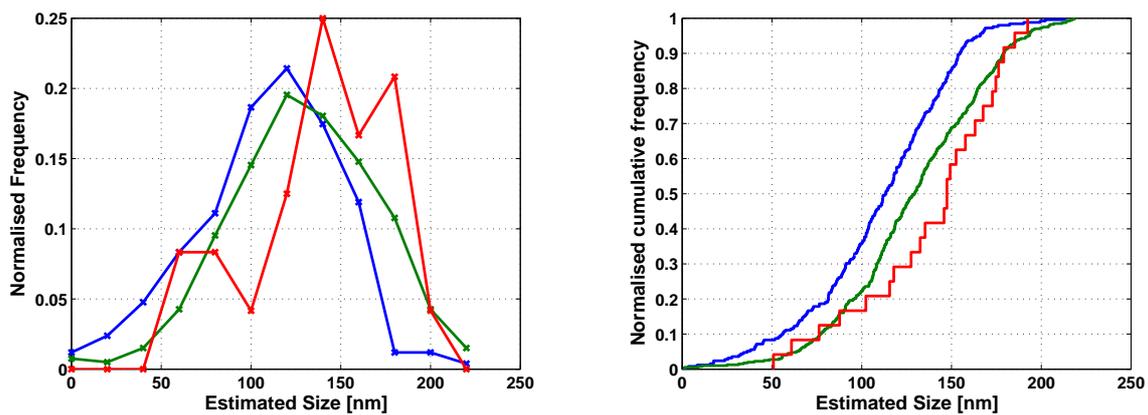
Figure 11.3: Distribution of sizes in GFP-PCNA labelled foci when classified according to S-phase position (blue = early, green = mid, red = late).

signal, with around 60% of foci being able to be fitted. A typical image is shown in Fig. 11.4 and the results of the size measurements on all cells in Fig. 11.5. Again the difference between different S-phase positions is slight.

11.3 Labelling through BrdU incorporation

An alternative method for labelling the replication structures is to focus on the newly replicated DNA. In order to visualise this, an analogue of thymidine (in this case Bromodeoxyuridine (BrdU)) is added to the medium in a concentration which is considerably higher than that of the endogenous, unmodified nucleotide. This analogue is then pref-

Figure 11.4: In focus slice from typical SMI image of antibody amplified GFP-PCNA. The scale bar is 5μm.
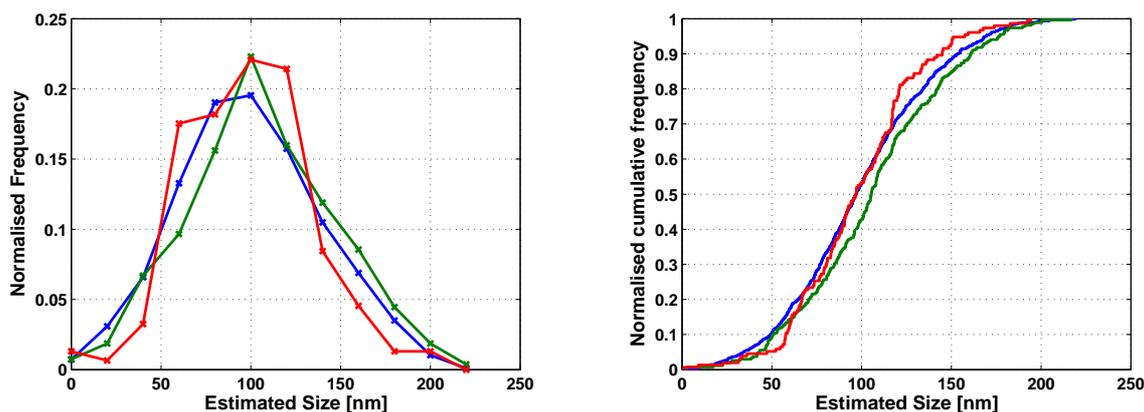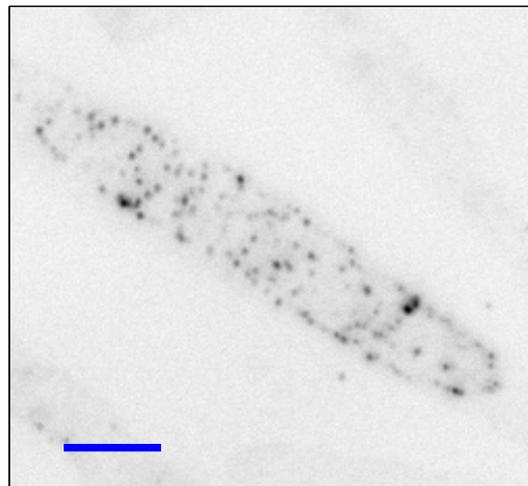
Figure 11.5: Distribution of sizes in GFP-PCNA labelled and antibody amplified foci when classified according to S-phase position (blue = early, green = mid, red = late).

erentially incorporated into the newly replicated DNA. After a certain incubation time the nucleotide substitution can be stopped by changing the medium or fixing the cells. The newly replicated DNA can then be recognised by an immunostaining using antibodies against BrdU.

The cells in which BrdU was incorporated were subsequently hypotonically flattened and spun down before antibody labelling. This treatment method separated the larger late S-phase conglomerations into individual loci, and resulted in a preparation that was very thin (less than approximately 200nm, estimated from the spread of z-axis positions of the individual foci), and thus circumvented the need for optical sectioning. These *whole cell projections* allowed all foci within a cell to be investigated rather than a small sub-
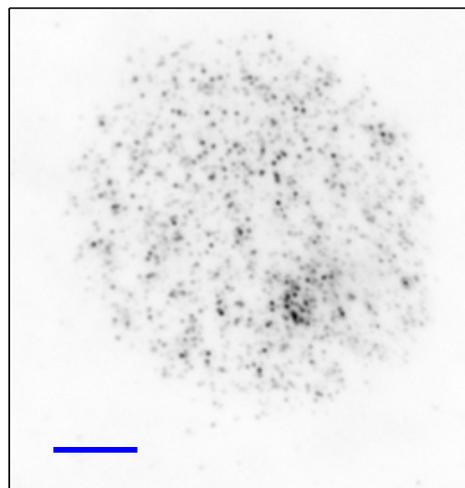
Figure 11.6: In focus slice from typical SMI image of antibody labelled foci of BrdU incorporation. The scale bar is 5μm.
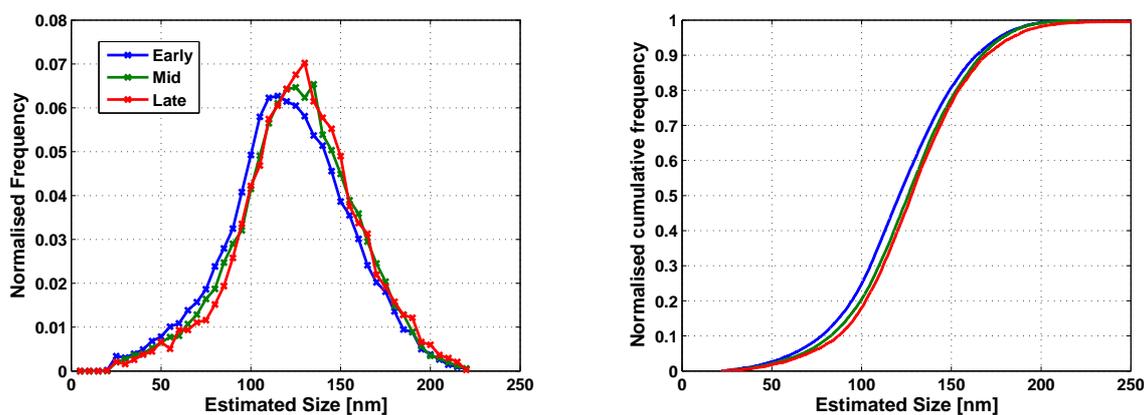
Figure 11.7: Distribution of sizes in foci of BrdU incorporation when classified according to S-phase position.

population as was the case for the cryosections, resulting in a much larger number of analysed foci and better total statistics. A two layer antibody protocol with Alexa488 coupled to the secondary antibody was used, resulting in an excellent signal. An example SMI image of the BrdU labelling is shown in Fig. 11.6, with the results from all cells being shown in Fig. 11.7. As with the PCNA labellings, the difference between different S-phase positions is slight.

11.4   Comparison of Labelling Strategies

A summary of the results of all labellings is shown in Fig. 11.8 and Table 11.4. Whilst the results obtained for the BrdU incorporation and the GFP-PCNA labelling are similar, a significant difference is seen when using antibodies against GFP to amplify the GFP-PCNA signal. As the antibody labelling is present in addition to the GFP labelling, one would expect a slight increase in the measured size. The antibody labelled loci are however smaller than those with only GFP.

One possible explanation is that crosslinking due to or between the antibodies leads to a smaller structure. A similar effect has been seen in EM studies where additional layers of antibody labelling can lead to smaller structures[4]. An example in which the addition of antibodies can be shown to induce a structural change is the induced clustering of membrane proteins [101]. In this case, the addition of antibodies was performed before fixation. Whether antibodies are capable of inducing structural changes after fixation remains unclear.
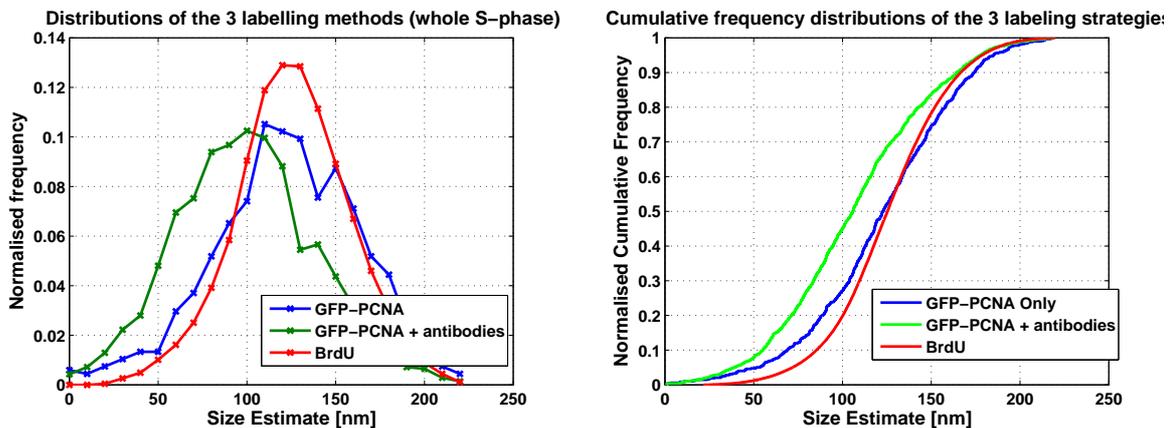
Figure 11.8: Comparison of foci size distributions obtained from different labelling procedures.

It is also interesting that the size of the BrdU labelling (marking the replicated DNA) is such a good match to that of the PCNA (a component of the replication machinery). This would suggest both relatively slow progress of the foci and a large number of individual replicons within one focus. The broad distribution of sizes also speaks for a relatively large number of individual replicons within the foci. Given the breadth of the distributions, it is however difficult to draw any more specific conclusions.

When exposed to statistical tests (Kolmogorov-Smirnov test, T-test) most of the differences in the distributions are highly statistically significant. These differences are however small compared with the width of the distributions and when compared with the variability seen between individual cells. It is thus difficult to interpret these differences in a biological context. It is indeed possible, or even likely, that the differences between early and late foci sizes are related to systematic effects due to the density of foci.

---

[4]Priv. Comm. A. Pombo

		Parameters				KS Test p-value		
		Mean [nm]	s.d. [nm]	# foci	# cells		Mid	Late
	Early	123.5	32.2	17572	29		5e-26	3e-38
BrdU	Mid	127.1	31.5	13041	24			0.0062
	Late	128.6	31.1	10897	18			
	Early	109.9	39.1	252	6		7e-6	0.0013
GFP-PCNA	Mid	128.8	40.2	399	14			0.1576
	Late	138.3	39.3	24	2			
	Early	99.0	40.1	972	14		0.0117	0.2381
abGFP-PCNA	Mid	106.2	40.6	269	6			0.0415
	Late	98.0	34.5	154	6			

Table 11.1: Table summarising the results from all three labelling techniques: Statistical testing was performed using the Kolmogorov-Smirnov test and the corresponding p-values are shown. Green colouring indicates that the null hypothesis that the distributions are equal can be rejected on a 95% confidence interval, red that it is not possible to reject the null hypothesis. Due to the very large sample size, the small differences present in the BrdU labelling are all statistically significant. The GFP labelling shows somewhat larger differences, which are also significant when comparing the early with the mid and/or late distributions. It however suffers a little from the small sample size in terms of late foci.

11.5 Cell to Cell Variability

Whereas the differences between the different S-phase positions are slight, there is significant cell to cell variability, as shown in Fig. 11.9. This variability does not seem to depend on either the S-Phase state or number of foci detected within the cell as shown in Fig. 11.10. The distribution of the per-cell mean sizes is shown in Fig. 11.11. This is notably much broader than would be the case if the foci within each cell were drawn randomly from the population of all measured foci (with an average of ~ 800 foci/cell one would expect a standard deviation of $30/\sqrt{800} \approx 1nm$). Another fact which points to multiple underlying distributions is the non-Gaussian form of the distribution of mean sizes (see cumulative histogram in Fig. 11.11). According to the central limit theorem, the distribution of the means of samples taken from any distribution should be Gaussian, even if the underlying distribution is non-Gaussian.

The division into classes according Early, Mid, and Late replicating status did not result in any distinguishable difference in the size distributions. The variability data however indicates that there are significant differences between cells. This leaves three possibilities; that we messed up the categorisation[5], that there is some other biological parameter which would result in a categorisation, or that the natural variability found in several biological

---

[5]One would hope that this is not the case. The categorisation is however performed by visually examining the r.f. patterns and is not trivial in either cryosections or hypotonically flattened specimens. This resulted in several cases ($\sim$30%) in which experts from different labs produced contradictory categorisations.

Distributions of foci sizes within early s−phase cells

Figure 11.9: Box-plot show-
ing variability within early S-phase
cells (BrdU labelling). The mean
sizes and to some extent the
spread of sizes vary significantly
between individual cells. Note
that the spread of sizes within a
single cell is typically smaller than
in the distribution of all foci.
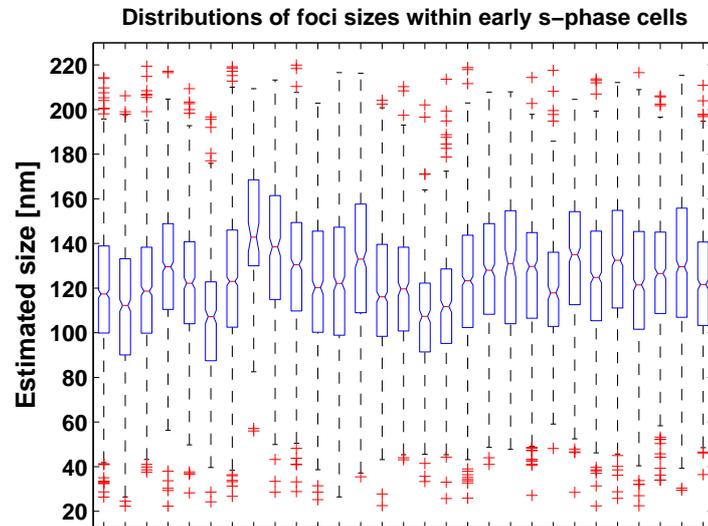
Figure 11.10: Mean size as func-
tion of number of points detected,
plotted for all S-phase positions.
A large spread in mean sizes can
be seen which does not appear to
correlate with either the S-phase
position or the number of detec-
ted foci.
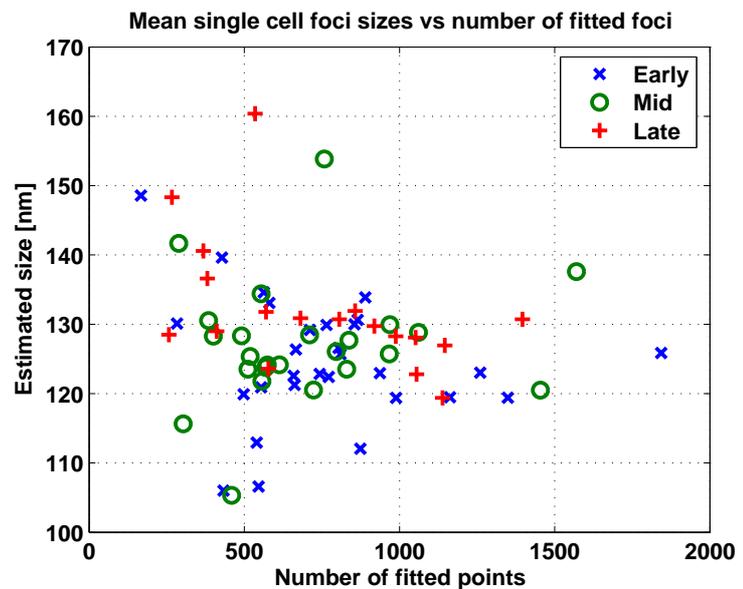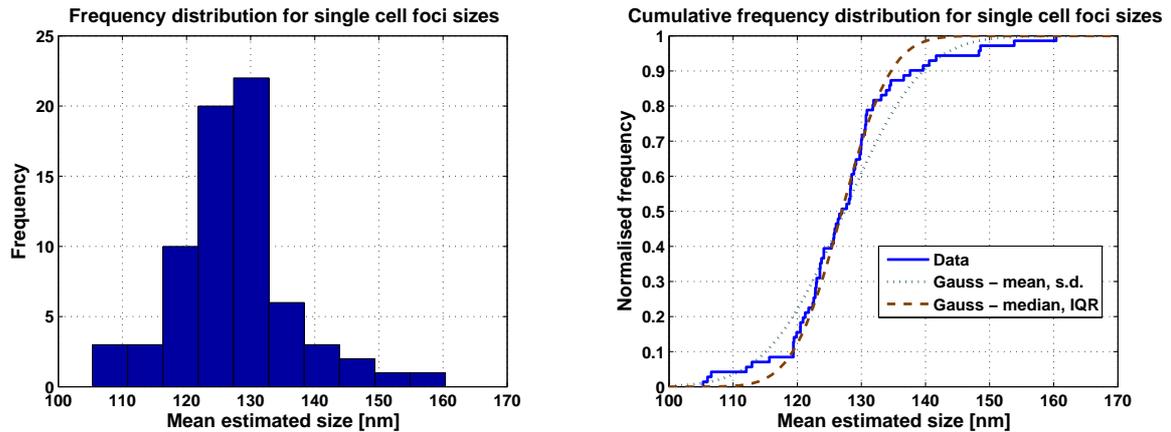
Figure 11.11: Distribution of mean sizes. Non-Gaussian shape is illustrated in the cumulative histogram.

systems is at work. In any case the spectrum of mean sizes is continuous enough that it seems unlikely that there are discrete sub-populations which could be identified from the data itself[6] without the use of an additional parameter.

11.6 Outlook

In addition to the experiments outlined above, a couple of further experiments were started which could be worth following up.

11.6.1 Multi-colour Labelling

In an effort to clarify the categorisation of the cells according to S-phase position, preliminary experiments were made with multi-colour labelling in which foci were labelled by two periods of pulse labelling, one with BrdU, the other with IdU. This enables two time points in S-Phase to be captured and the comparison of earlier and later foci within the one cell. The areas of BrdU and IdU were recognised with corresponding antibodies and fluorophores at 488 and 568nm excitation. Unfortunately the signal level in these samples was rather low, and the labelling inconsistent with few cells in which both labellings were useful. This is an area which should be followed up with additional experiments. It would also be useful to have both the GFP-PCNA and BrdU labellings in the same cells. One hinderence to this was that the GFP-PCNA labelling did not survive the hypotonic treatment used to obtain the 'whole cell projections'.

---

[6]This is not really to be expected - replication is a continuous process.

11.6.2 4Pi Measurements

The 4Pi, at least theoretically, has the potential to allow high precision size measurements within entire, 3D preserved, nuclei. As previously mentioned (Section 7.7.2) our attempts at size measurements using 4Pi microscopy have not yet been able to provide reliable results. Even if the sources of the errors in our size estimates cannot be reasonably controlled, high resolution 4Pi images of replication foci would be of considerable interest.
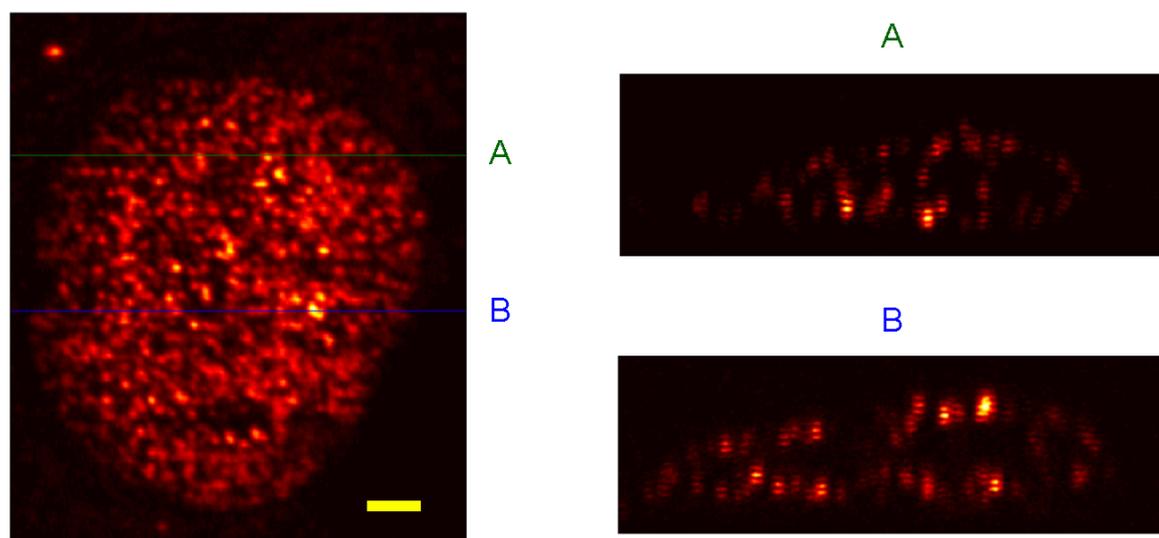
Figure 11.12: Left: Mean intensity xy-projection of 4Pi image of replication foci. The scale bar is approximately 2μm. **Right:** xz slices through the stack at the locations **A** and **B**.

Preliminary 4Pi measurements were made, and a 3-layer antibody protocol with alexa488 was shown to provide a good signal level. Unfortunately the measurements were made using the early oil objectives and with Vectashield as the embedding medium. This resulted in sidelobes which were too high (50-60%) for deconvolution and a significant refractive index mismatch induced phase change through the image. An example measurement is shown in Fig. 11.12.

Although the conditions were not ideal for deconvolution, the axial fit described in section 7.7 was able to be applied to each of the foci giving information about the positions of the loci and various parameters related to the PSF and object. One of the parameters, the PSF phase is shown in Fig. 11.13
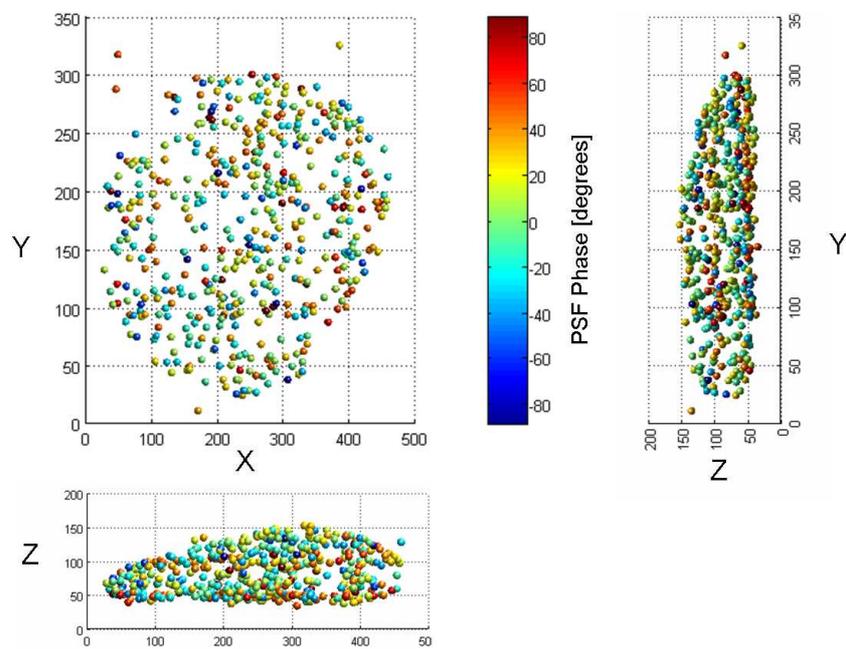
Figure 11.13: Phase information obtained from applying the 4Pi fit function to each locus. A phase progression along the z-axis is clearly visible.

Chapter 12

Virtual SMI Microscopy of Gene Regions

12.1 Numerical Models for Chromatin Structure

The theoretical and numerical description of chromatin structure is a complex field, a full description of which is beyond the scope of this work. The comments here thus represent a minimalist overview of those concepts used for the virtual microscopy simulations. The interested reader is advised to consult additional sources for further details and derivations e.g. [76, 62, 77]

12.1.1 Polymer Basics

As a first approximation, chromatin structure, be it bare DNA, 10nm or 30nm fibre, can be described as a flexible rod. The appropriate model from polymer physics is thus the worm like chain model. We can thus parameterise a chain $\mathbf{r}(s)$ of length L according to the partial length $s \in (0, L)$. The tangent vector $\hat{\mathbf{t}}$ is then

$$\hat{\mathbf{t}}(s) \equiv \frac{\partial \mathbf{r}(s)}{\partial s}$$

It can be shown that the *orientation correlation function* has the following form;

$$\left\langle \hat{\mathbf{t}}\left(0\right) \cdot \hat{\mathbf{t}}\left(s\right) \right\rangle = e^{-s/P} \tag{12.1}$$

where P is the *persistence length*, or alternatively

$$\left\langle \cos\left(\theta\left(s\right)\right) \right\rangle = e^{-s/P}$$

where θ is the angle between the tangent at 0 and s. Another important result is that the Kuhn segment length for a worm like chain is twice the persistence length.

	Persistence length [nm]	Length per kpb [nm]
Pure DNA	50	$.34 \times 10^3$
30nm fibre	150	10

Table 12.1: Values used for simulation

12.1.2   Segmented Chain

On a length scale $L \gg P$ a worm like chain can be approximated as a freely jointed chain with a segment length equal to the Kuhn segment length. A freely jointed chain shows no correlation between the directions of adjacent segments and can thus be modelled as a random walk. This is trivial to implement as illustrated by the following Matlab code:

```
steplength = 300;
length_per_kbp = 10;

kbp = 240;

numsteps = round(length_per_kbp*kbp/steplength);

xs = 1 - 2*rand(numsteps, 1);
ys = 1 - 2*rand(numsteps, 1);
zs = 1 - 2*rand(numsteps, 1);

nrm = sqrt(xs.^2 + ys.^2 + zs.^2)

xs = steplength*xs./nrm;
ys = steplength*ys./nrm;
ys = steplength*ys./nrm;

rx = cumsum([0;xs]);
ry = cumsum([0;ys]);
rz = cumsum([0;zs]);
```

Using the literature values for length per base pair and persistence length it is thus possible to generate representative gene structure with minimal effort. The values of persistence length and genomic length used in the work presented here are presented in Table 12.1. Typical structures for a 240kbp region using the segmented chain model are shown in Figure 12.1.

12.1.3   Worm Like Chain

The segmented chain model, whilst a good approximation, results in a piecewise linear model with several jagged corners, and considerable high frequency content. As SMI
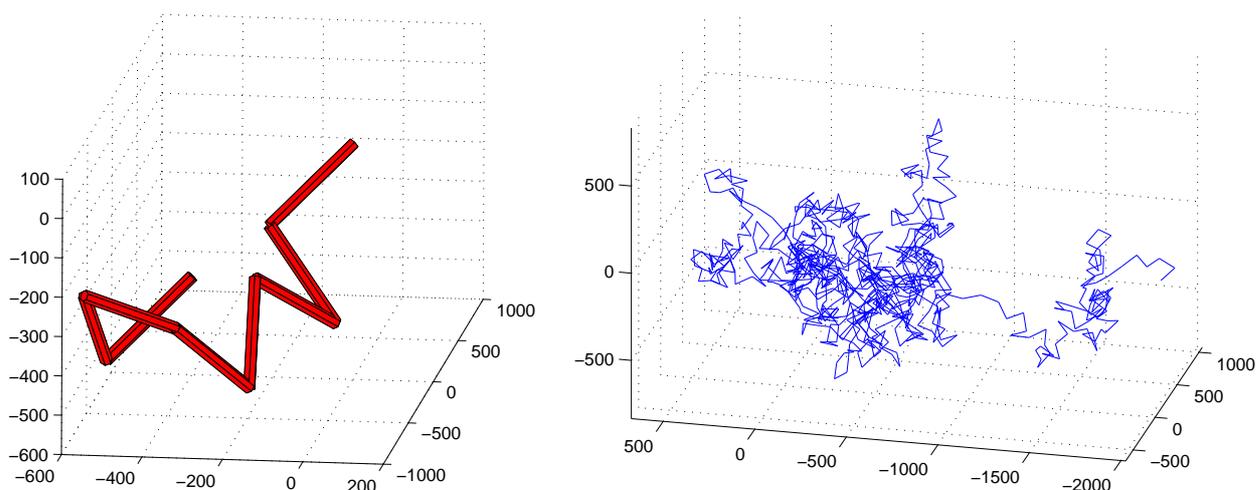
Figure 12.1: 240kbp of DNA as simulated using a random walk. Left: 30nm fibre, Right: bare DNA. The scale is in nm in both cases.

microscopy extracts its information on a length scale from around 20 to 200nm, the assumption that $L \gg P$ cannot be assumed to be valid. To see what kind of errors are introduced through use of a segmented chain model for image generation, it would be useful to compare the results with those obtained from a true worm like chain model. This can be simulated by using a shorter segment length and choosing the angle between subsequent segments from a random distribution with expectation value derived from Eqn. 12.1. Typical structures obtained using the worm like chain model are shown in Figure 12.2. When compared to the segmented chain structures, those derived from the worm like chain model are both qualitatively and quantitatively similar in terms of overall extension and in the presence of substructures.

Neither of the models have excluded volume or self-avoiding properties which are present in real chromatin. In visual analysis of several simulated gene structures no overlap between different parts of the chain could be determined. It thus unlikely that the inclusion of excluded volume interactions would significantly alter the results obtained when using these models. A further factor which is not included in these simple models is any form of higher order chromatin structure.

12.1.4 Odenheimer Model

The model developed by Jens Odenheimer in his doctoral thesis attempts to simulate higher order chromatin structure such as rosette formation. This higher order structure is created by assuming regularly spaced attractive segments along a stretch of 30nm fibre. The underlying model is a segmented chain model with excluded volume interactions, and the structures are solved by performing Monte-Carlo simulation. This model is fully described in [77].
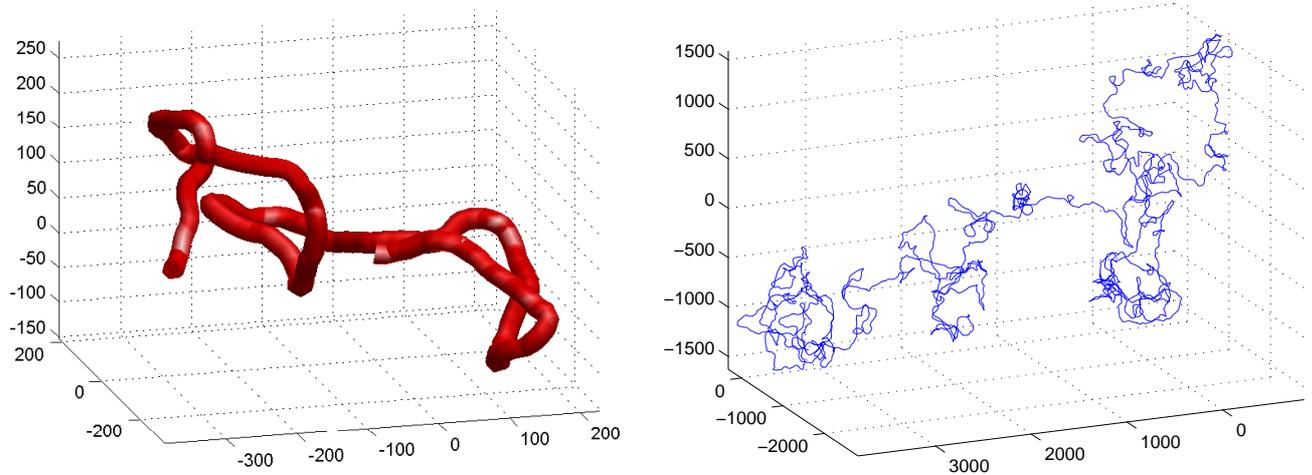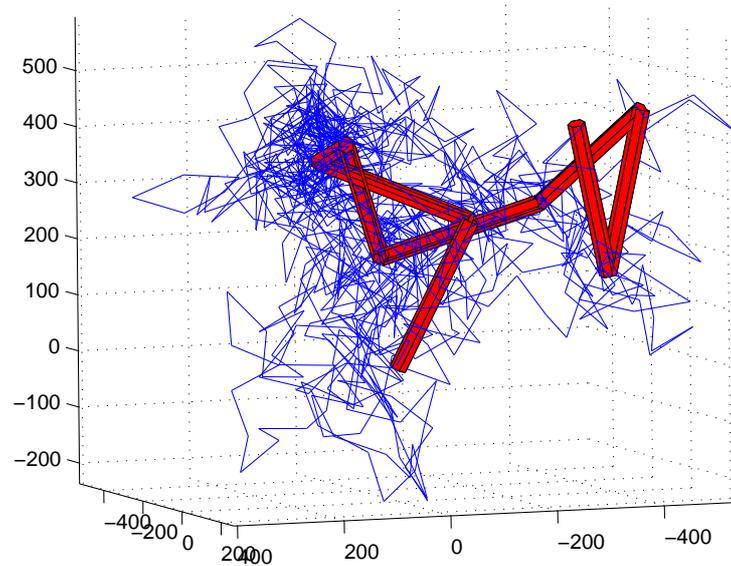
Figure 12.2: 240kbp of DNA as simulated using a worm like chain model. Left: 30nm fibre, Right: bare DNA. The scale is in nm in both cases.

12.1.5   Denatured 30nm Fibre Model

Comparison with confocal and SMI images of real, FISH labelled gene regions showed a reasonable (but nonetheless with significant differences) agreement between larger probes and both 30nm random walk and Odenheimer models. The results from smaller regions could not however be brought into agreement with any form of 30nm fibre model. One possible explanation of this is that, due to denaturation and loss of histone proteins through fixation and FISH, a previous 30nm fibre structure collapses to a somewhat larger structure consisting of bare DNA. This collapse is however hindered by the surrounding DNA, resulting in a degree of structural preservation at larger scales. In an attempt to model such a situation, a 30nm fibre was simulated by the random walk method. A potential surface was then created which depended on the distance from this 'template' structure. This potential surface was then explored with a modified random walk of bare DNA fibre, in which steps were made in a Monte-Carlo fashion with the probability of accepting the next step being dependant on the change in 'energy' it would bring. An example structure as created by the denatured 30nm fibre model is shown in Figure 12.3

12.1.6   Wormlike Chain With Extra Compaction

The simple random walk and wormlike chain models presented above have no mechanism for higher order compaction, and are thus potentially insufficient for the description of real chromatin structures. The Odenheimer model includes a compaction mechanism, but due to the underlying 'freely jointed chain' model only provides a realistic representation of structures on length scales significantly longer than the Kuhn segment length (≈ 300nm). The extent to which such discrepancies affect the results of the virtual microscopy simulations is unclear. In order to rule out such effects a re-implementation of the Odenheimer
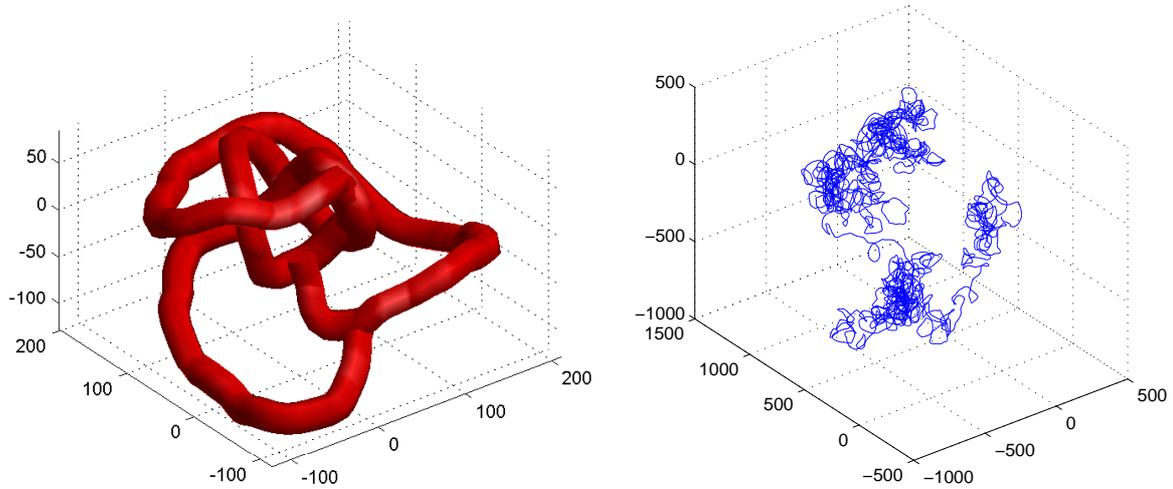
Figure 12.3: 240kbp of DNA as simulated using denatured 30nm fibre model. Red: 30nm fibre template, Blue: bare DNA. The scale is in nm. cases.

model with a wormlike chain as the underlying structure would be desirable. This would entail a significant amount of work, and could be an interesting project for a 'Diplomarbeit'. It would also place appreciable demands on the computer hardware used.

A simple model which nonetheless offers a way of modelling compaction can be constructed with a small modification to the wormlike chain model discussed above. Instead of assuming free rotation between the individual segments, a certain torsional stiffness is assumed. In order to include this stiffness in the model, the rotational angle between segments can be chosen from a narrower distribution (in this case Gaussian) instead of a uniform distribution on $[0; 2\pi)$. The stiffness is then determined by the width of the Gaussian. In order to induce compaction the fibre can now be twisted, causing it to coil up. This twist can be introduced by shifting the mean of the Gaussian away from zero. Whilst the biophysical foundations of this model are somewhat shaky (I have been unable to find data on the torsional stiffness of 30nm fibre and it is unclear whether there is significant chromatin coiling on the 30nm fibre scale), it is physically plausible - as anyone who has had to battle with their phone cord will testify. Despite the fact that the compaction mechanism is most likely incorrect, it can expected that the structures obtained using this model will be qualitatively similar to those produced through other compaction mechanisms. As such the model should allow an impression of the effects of higher order compaction to be obtained.

Example structures generated using the twisted wormlike chain model are shown in Fig 12.4. At this level of compaction the absence of an excluded volume constraint begins to have an effect on the structure. It is however still relatively small and unlikely to have a significant effect after the virtual microscopy image generation and analysis.

Figure 12.4: 240kbp of DNA as simulated using a worm like chain model with additional compaction due to twisting. Left: 30nm fibre, Right: bare DNA. The scale is in nm in both cases.

12.2   Image Generation and Evaluation

For the simulations made with the Odemheimer model, the data was in form of 'fluorophore' positions with 'fluorophores' spaced at 60nm intervals along each segment. In this case an image was created by adding 1 to the voxel nearest to each 'fluorophore', and convolving the result with a theoretical PSF. This PSF was simulated using a Khoros PSF simulation plugin written by R. Heintzmann to generate a widefield PSF and multiplying the result by $\cos^2(2\pi.n.z/\lambda)$ to construct the SMI axial illumination pattern. The voxel size used was $108 \times 108 \times 40$nm, the values typically used in real SMI images. Figure 12.5 depicts the simulation process.

For the other 30nm fibre simulations, a similar approach was taken, with the most important difference being that a much larger number of virtual fluorophores was used and that these were distributed randomly around the centre of the structures to give the fibre its 30nm width. Unlike the simulations using the Odenheimer model, which did not include any noise, the later simulations assumed a Poisson noise with 100 counts per fluorophore. Whether the resulting signal to noise ration accurately represents that obtained in real images is however debatable due to the large number of fluorophores. In all cases the resulting images were evaluated as if they were normal SMI images[1] using the standard SMI evaluation procedure (`ofind`, `dofits` etc...).

12.3   Results

12.3.1   Odenheimer Model

---

[1] albeit small ones
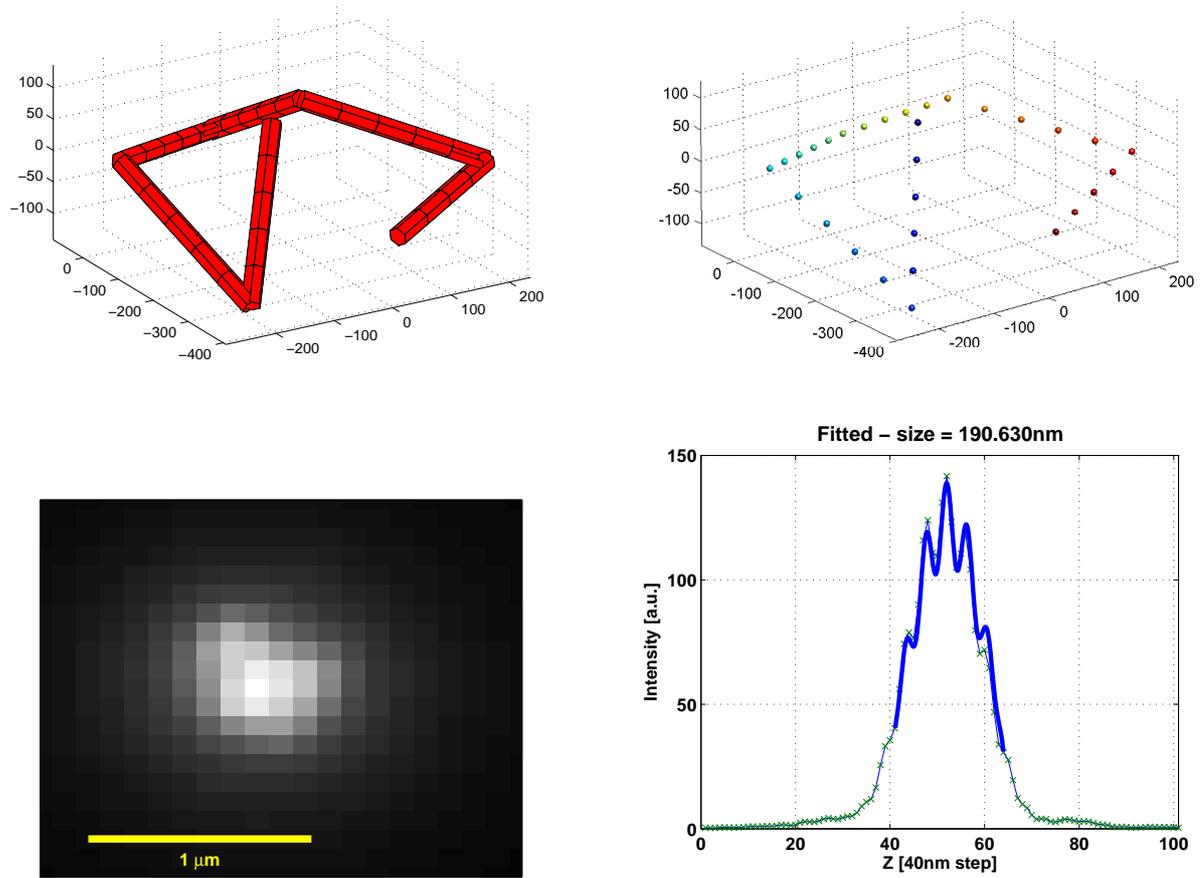
Figure 12.5: Schematic of simulations using data from the Odenheimer model. A 180kbp portion of a simulated locus (**upper left**) is taken, and 30 'fluorophores' are evenly distributed along its length (**upper right**). These fluorophore positions are used to generate an image through convolution with the SMI PSF (mean projection - **bottom left**). This virtual image is then analysed using the normal SMI analysis routines (axial profile and fit - **bottom right**). The axes of the two 3D images are in nm.

The size distribution obtained from a 180kbp segment simulated with the Odenheimer model is shown in Fig. 12.6. A broad distribution is obtained with a mean at around 150nm. As can be seen from both the scale on the axes of the 3D representation of the structure and the non-diffraction limited nature of the image in the $x - y$ plane (Fig. 12.5), the structures are however significantly larger than the 200nm upper limit for SMI size measurements. This is a clear indication that the size estimates obtained for such structures can not be unambiguously[2] interpreted as sizes. The extent to which the modulation present may provide information about other aspects of the structure and/or whether useful size information could be extracted from smaller structures of a similar complexity is discussed

---

[2]or potentially at all

Figure 12.6: Measured size distribution obtained from 50 Odenheimer model configurations (each 180kbp). Note the broad distribution and the fact that despite a size significantly larger than 200nm, some very small size estimates are obtained.
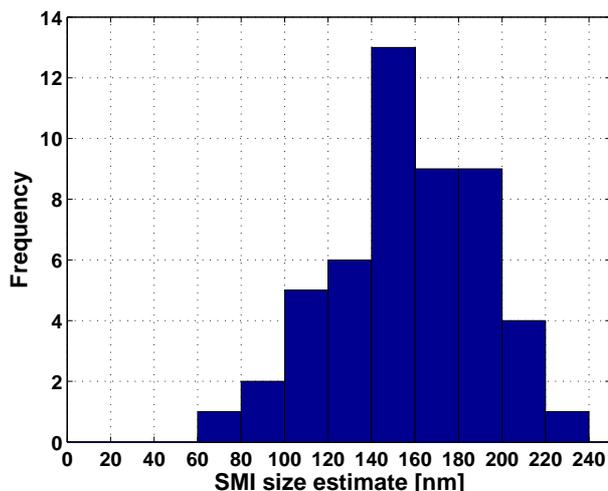
Figure 12.7: Comparison of the SMI size estimates to the true axial size of the underlying structures (calculated by multiplying the axial standard deviation by 4 to allow it to be compared with the diameter of a sphere). The red points correspond to the original size structures, and show no correlation between measured and actual sizes. The blue and green points are scaled down copies of the structures, computed with the aim of establishing whether a measured size vs actual size correlation could be obtained for such objects if they were in principle small enough.
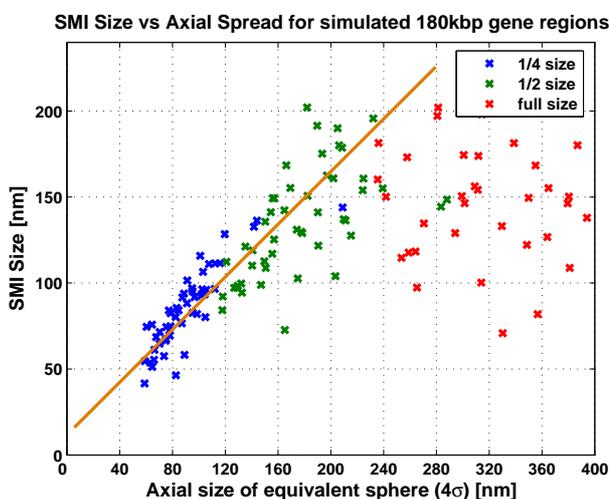
in the following sections.

Comparison to actual z-extensions

Figure 12.7 shows the correlation (or lack thereof) between the SMI size estimates from simulated gene structures and their true size as calculated from the distribution of the virtual fluorophores used to simulate the images. For the original size loci (red crosses) no correlation can be seen. Scaling the structures to half their original size (green crosses) results in a weak correlation with a large spread and halving the size again results in a reasonably good correlation. This implies that, if one can prove that the structure is small enough, a reasonable correlation between the sizes measured with SMI microscopy and the actual sizes should exist, even for relatively arbitrary objects.

It is however uncertain to what extent this relationship will be valid for smaller gene structures, as the persistence length and attractive centre spacing have also been scaled down in the scaled structures from Fig. 12.7. The scaled structures are thus compact when compared with the likely structure of true small gene regions. If the region is small enough

(compared to the persistence length, $\lesssim$ 15kbp), and consists of 30nm fibre one can consider it to be a rigid rod with 30nm diameter. If one considers all possible orientations of such a structure a broad size distribution covering most of the SMI accessible size range would be expected.

Figure 12.7 also gives us some idea of what we might expect for arbitrary objects for which size measurements are useful. If the mean size is less than $\sim$ 150nm and the distribution is relatively narrow, then there is a good chance that the SMI size estimates correlate with the real object size.

12.3.2   Effect of Labelling Efficiency

The initial motivation of the simulations discussed here being to try and explain the results of Connie's measurements [35] (COMBO-FISH probes), I also tried to simulate the effect of a poor labelling efficiency on the measurement results. This was done by randomly choosing (with a given probability) whether or not to include each fluorophore in the image generation process. A reduction in the labelling efficiency leads to a lower estimated size and a larger spread in the size estimates, as illustrated in Fig. 12.8.
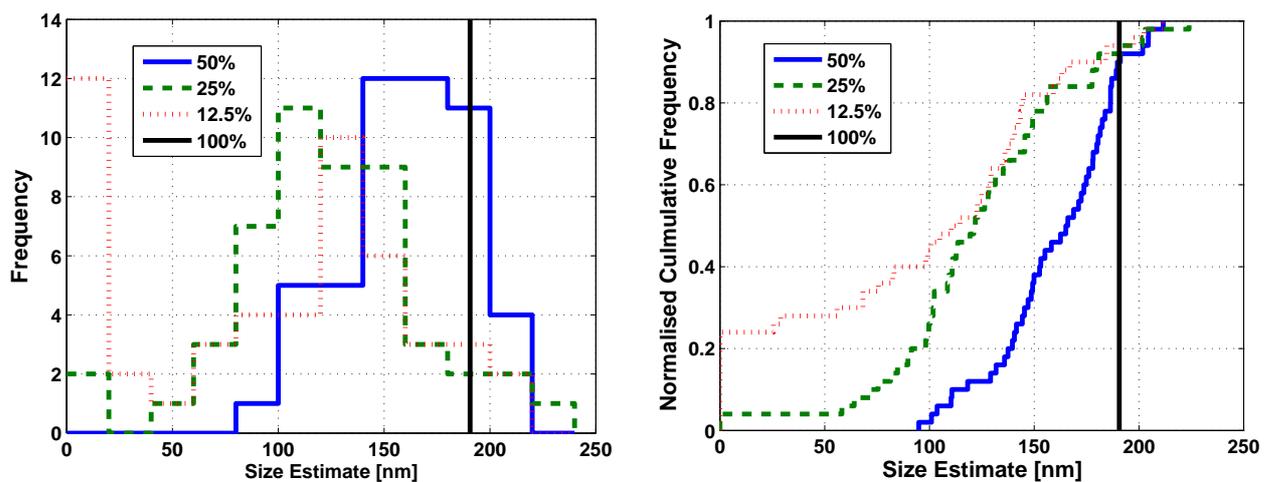
Figure 12.8: 'Size' distributions obtained from the same 180kbp gene segment with different labelling efficiencies (**left:** distributions, **right:** cumulative histogram). A lower labelling efficiency generally leads to a lower mean size and a larger spread.

12.3.3   200kbp 30nm Fibre - Wormlike Chain

The wormlike chain model was used to simulate several configurations of a 200kbp length of 30nm fibre. 200kpb was chosen to approximate the length of the BAC probes (and of the regions over which the oligo-FISH probe sets were distributed) used in the majority of the biological experiments performed within the group during my thesis. In order to

Figure 12.9: Distribution of modulation depths acquired from virtual SMI images of 200kbp of 30nm fibre simulated using the worm like chain model. The simulation was performed for excitation wavelengths of both 488 and 568nm.
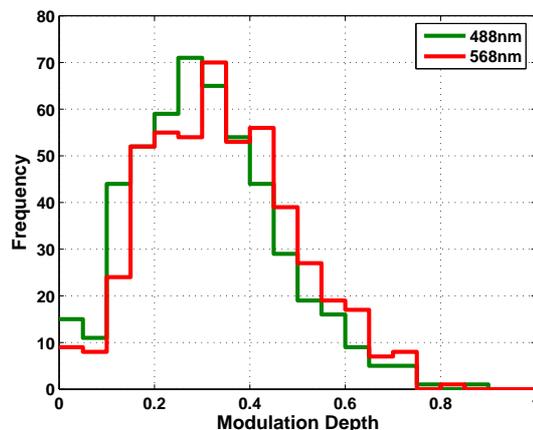
Figure 12.10: Distribution of size estimates using a spherical object approximation aquired from virtual SMI images of 200kbp of 30nm fibre simulated using the worm like chain model. The simulation was performed for excitation wavelengths of both 488 and 568nm.

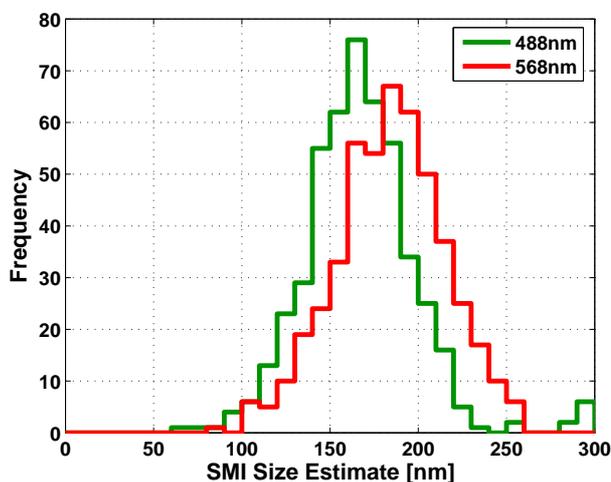evaluate the potential of multi-colour measurements for the determination of structural parameters, the simulation was performed for excitation wavelengths of 488nm and 568nm, corresponding the the FITC and TRITC channels. Figure 12.9 shows that significant modulation was observed despite a size (see 'decondensed' curve of Fig. 12.13) which is considerably larger than the typical 40-200nm range accessible to SMI microscopy.

Translating the modulation depths into sizes based on the spherical model results in the distributions seen in Fig. 12.10. Whilst, given the known size distribution of the regions, the validity of these estimates is more than questionable, most existing data is presented in this form, and the representation is thus useful for comparison, and to highlight the possible pitfalls of the method.

When comparing the distribution of modulation depths/sizes between the two excitation wavelengths there is surprisingly little difference in the modulation depth, with the size estimates being simply scaled in approximately the ration of the wavelengths.

Figure 12.11: Distribution of modulation depths aquired from virtual SMI images of 200kbp of 30nm fibre simulated using the worm like chain model with extra compaction. The simulation was performed for excitation wavelengths of both 488 and 568nm.
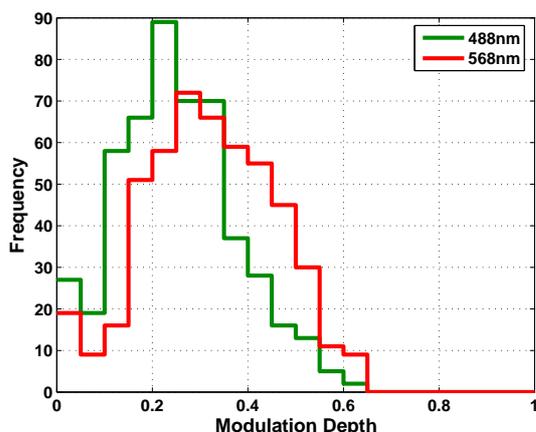

Figure 12.12: Distribution of size estimates using a spherical object approximation acquired from virtual SMI images of 200kbp of 30nm fibre simulated using the worm like chain model with extra compaction. The simulation was performed for excitation wavelengths of both 488 and 568nm.
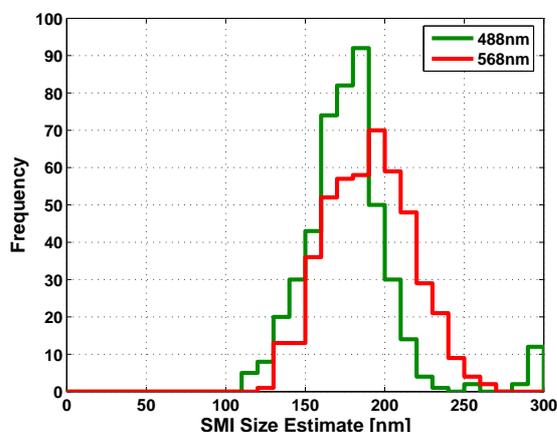
12.3.4   200kbp 30nm Fibre - Wormlike Chain With Extra Compaction

When compared with confocal measurements, the structures generated with the simple wormlike chain model are significantly larger than those of real gene loci (Baddeley, Weiland, Birk, & Cremer - submitted to JoM). Real gene regions are thus presumably exposed to some additional method of compaction. It was thus interesting to ascertain whether a more compact structure would show different a dependance on excitation wavelength. The wormlike chain model with extra compaction was thus used and the previous experiments repeated. Figure 12.11 shows the resulting distribution of modulation depths, and Fig. 12.12 the corresponding size distributions.

For the more compact structure, a difference in the modulation depth between the two excitation wavelengths can be seen, although the difference between the compact and normal wormlike chain models is not nearly as large as one might expect.

12.3.5   Comparison of Compact and Normal Wormlike Chain

To further investigate the differences between the compact and normal wormlike chain models it is instructive to view them side by side. A comparison of the true sizes of

Figure 12.13: Distribution of lateral FWHM estimates (ie 'real' sizes) obtained from the fluorophore positions of 200kbp structures simulated with the normal and compact wormlike chain models.
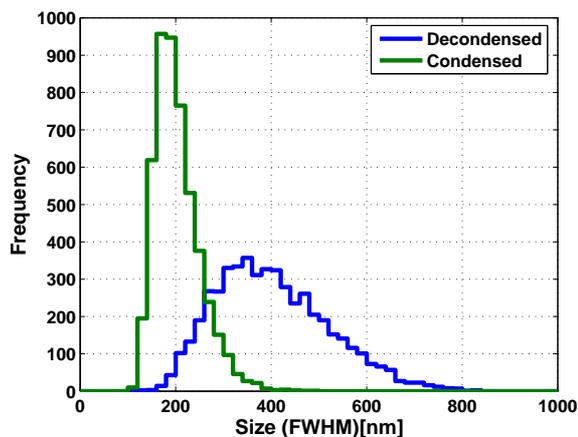
the compact and normal wormlike chain structures is shown in Fig. 12.13. Note that a two-fold difference in diameter corresponds to an 8-fold difference in effective volume, or condensation. Figures 12.14 and 12.15, show the corresponding SMI results at 488nm. Although there are differences, these are suprisingly small given the significant differences in compaction and hence locus substructure.
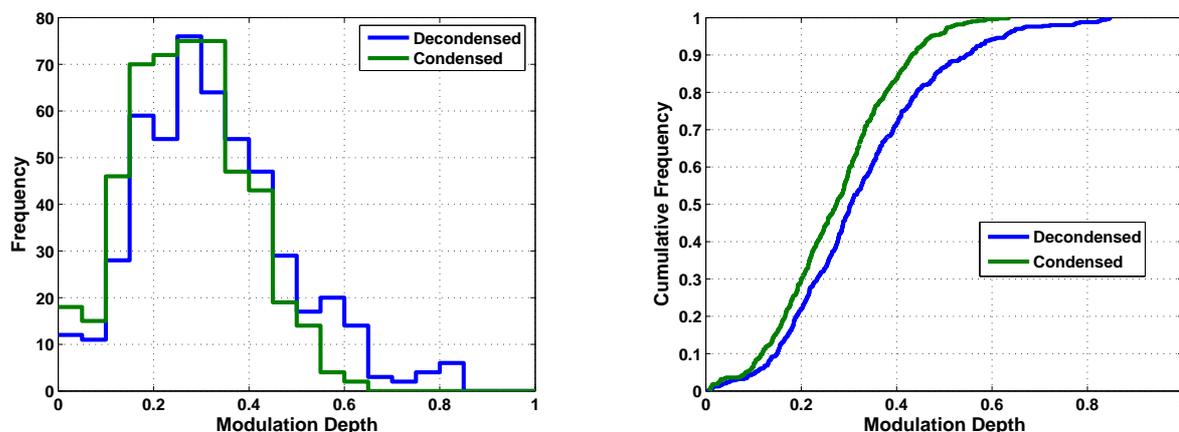
Figure 12.14: Comparison of the modulation depths obtained for 200kbp of 30nm fibre simulated using the worm like chain model with and without extra compaction.

Figure 12.15: Comparison of sizes obtained for 200kbp of 30nm fibre simulated using the worm like chain model with and without extra compaction.

12.3.6 Effect of 'Denaturation'

When compared with structures generated by applying the 'denaturing' procedure to the compact 30nm fibre structures, a significant difference is seen, as shown in Fig. 12.17. Because a much smaller (~50%) fraction of the loci fitted, and that with a very small modulation depth, the estimated size distribution has been ommitted. The true size distributions of the compact and 'denatured' structures are shown in Fig. 12.16.

Figure 12.16: Distribution of lateral FWHM estimates (ie 'real' sizes) obtained from the fluorophore positions of 200kbp structures simulated with the compact and denatured models.
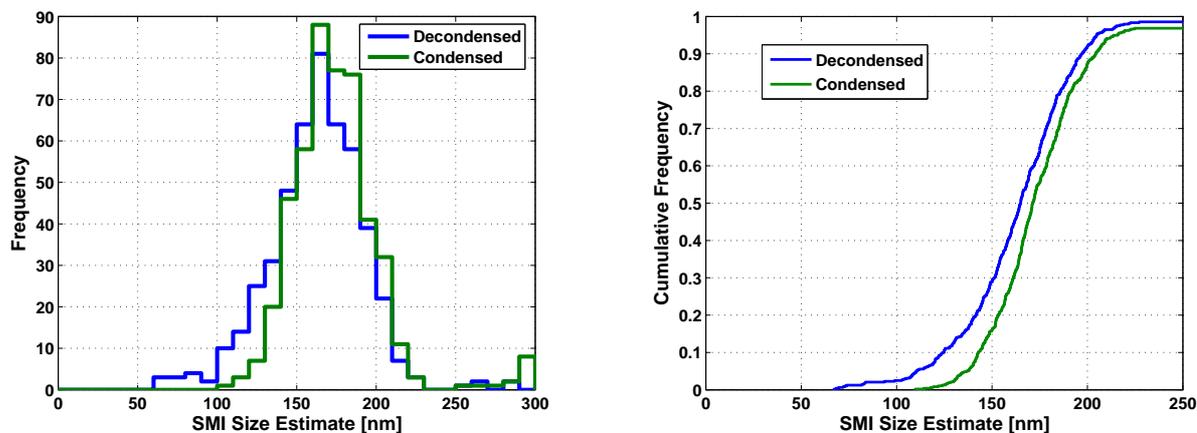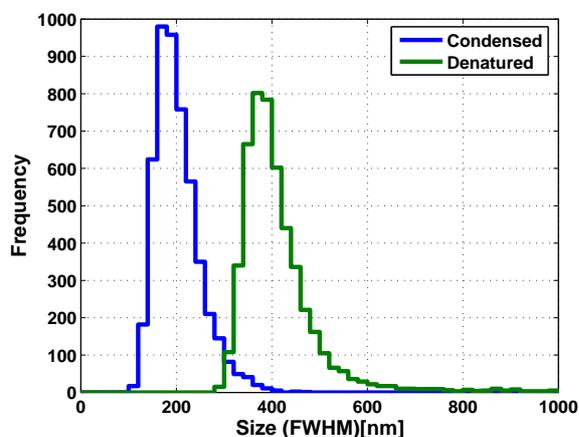
Figure 12.17: Comparison of the modulation depths obtained for 200kbp of 30nm fibre simulated using the worm like chain model with extra compaction before and after 'denaturation'.
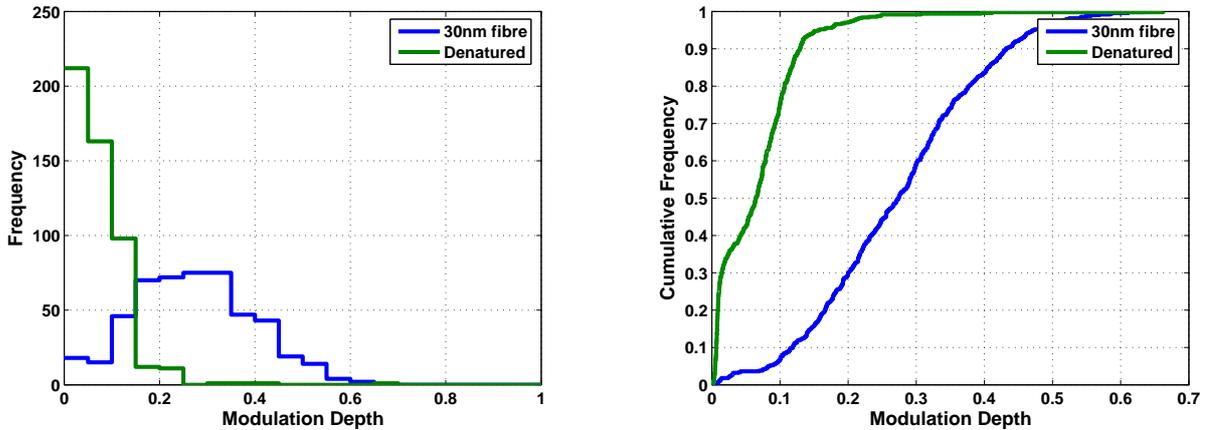
12.4 Summary

In brief, the simulations show that a correlation between modulation contrast and size is by no means assured for gene-like structures. Despite the disturbing nature of this revelation, the scaled structures in 12.3.1 indicate that a scaling relationship is present for arbitrary structures which are relatively compact and *small enough*. Small enough is in this case (for 488nm excitation) smaller than around 150nm. Whether structures fall into this category or not should be able to be determined by evaluation of the lateral size, either in the SMI or confocal microscopes.

For structures which are too large, the modulation depth is still capable of delivering some information about the object, although the interpretation of this is not trivial. The simulations indicate significant differences in modulation depending related to the underlying object structure, although the only clear trend is that sparser object tend to show more modulation than more homogeneous ones. It is also important to note that the differences are only visible in the distributions obtained from the measurement of several structures, the measurement of a single loci is unlikely to provide usable information about its form. Due to the difficulties in interpretation, we are unfortunately limited to rather qualitative statements, and these are best made by the comparison with the simulated results for hypothetical object structures.

Chapter 13

Inner Kinetochore Proteins

13.1 CENP Proteins

The kinetochore is a complex of proteins which form at the centromere. In metaphase the kinetochore is a multi-layer structure, to which the microtubuli responsible for separating the sister chromatids attach. Electron microscopic images of metaphase centromeres imply a disc or plate like structure with two layers, the inner and outer kinetochore, each about 30-40nm thick and separated by a similarly sized gap. The lateral width of this plate is typically 300-500nm ([80, 81] also Priv. Comm. S. Diekmann).

In interphase, only the inner kinetochore proteins are present and very little is known about its structure. The CENP family of proteins localise to the kinetochore, with CENP-A replacing the histone H3. There is evidence which suggests that this histone substitution, along with the binding of the other CENP proteins results in a more rigid structure of the centromeric DNA [102].

SMI measurements of centromeric proteins were performed in co-operation with Stephan Diekmann and Sandra Orthaus (IMB Jena). Labelling of the CENP-A and CENP-C components of the inner kinetochore was performed by expression of GFP and mRFP fusion proteins. From a technical point of view, these measurements represent the first measurements performed with the SMI on biological samples using the new 568nm excitation, and
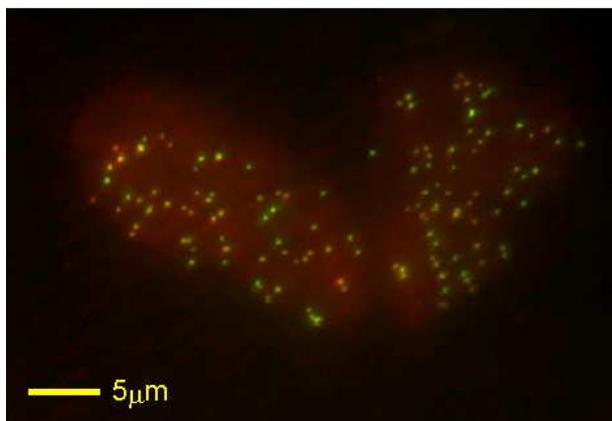
Figure 13.1: Typical SMI image ($x-y$ projection) of a pair of cells expressing both GFP-CENP-C and mRFP-CENP-A.

the first in which a reliable double labelling was present in a biological specimen allowing a comparison of measurements in different channels and hence the possibility of obtaining form information, as suggested in [104] by Christian Wagner.

13.2 Size/Form Measurements

13.2.1 GFP, mRFP double labelling of CENP-A

As a control, and in order to allow the analysis of object form, both GFP-CENP-A and mRFP-CENP-A were expressed within the same cells, resulting in a double labelling of the CENP-A distribution within the kinetochores. Some cells showed a very low expression, some a clear over-expression and several had very different expression levels for the two different colours. All in all around 25% of the cells gave a good signal in both channels, and could be acquired and analysed.

From the acquired data, around 25% of the cells were discarded due to the defects mentioned above (i.e. insufficient signal, over-expression, or large imbalance between channels). These were generally early measurements where every cell was acquired - for later measurements, the selection was made prior to data acquisition. From the resulting data set the distribution of modulation depths shown in Fig. 13.2 was obtained. It can be seen that the modulation depth is relatively low for the GFP labelling and somewhat larger for the mRFP labelling. This is what one would expect when the modulation depth reflects a scaling relationship, as the illumination pattern at 568nm is larger than that at 488nm. The distribution is broad in both cases.
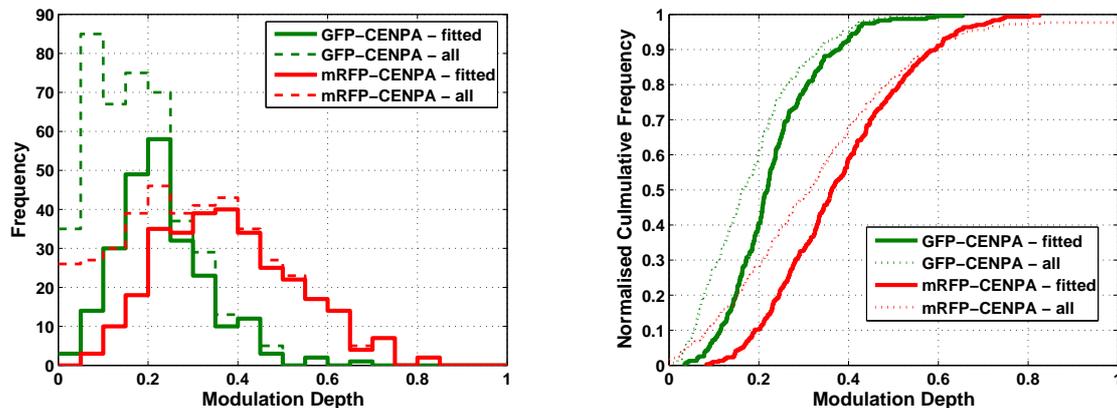
Figure 13.2: Distribution of measured modulation depths for the GFP (488nm excitation) and mRFP (568nm excitation) labelling of CENPA. **Left:** histogram, **Right:** cumulative histogram. Note that the data not categorised as *fitted* should not be considered to be reliable, it is however useful to obtain an impression as to what the trend would be if these points were included.

When translated into a size estimate using a spherical object model, the distributions shown in Fig. 13.3 are obtained. These are very similar, with the width of the mRFP

labelled distribution being somewhat broader. The fact that the same mean size estimate was obtained at both laser wavelengths could imply that the assumption of a spherical object form is not too bad, and the broader distribution of the mRFP labelling could be explained by an inferior measurement accuracy (too be expected, as the signal level was generally poorer and the wavelength longer).
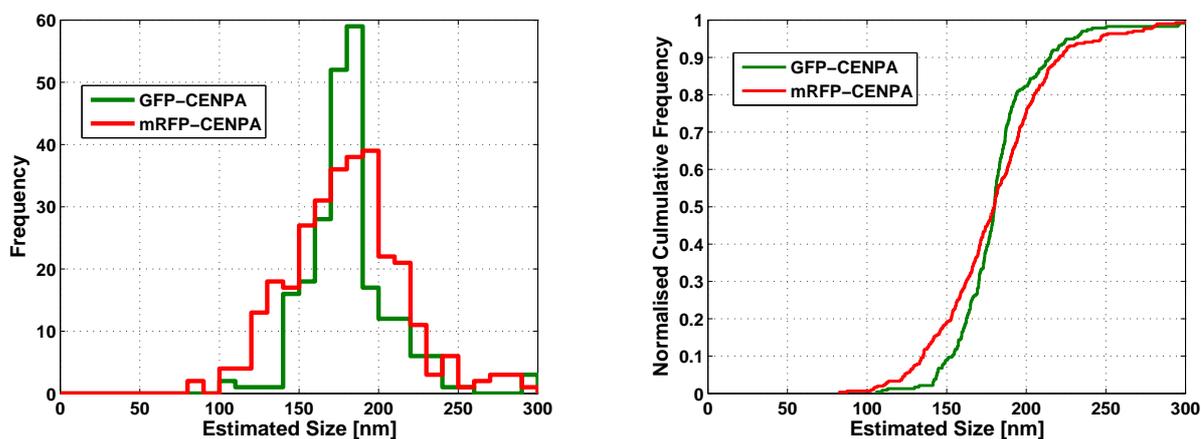
Figure 13.3: Distribution of size estimates based on a spherical model for the GFP (488nm excitation) and mRFP (568nm excitation) labellings of CENPA. **Left:** histogram, **Right:** cumulative histogram.

If the object form was approximately spherical one would expect a correlation between the individual modulation depths/ size estimates. From the scatter plot in Fig. 13.4 it is clear that such a correlation is not present. Measurement inaccuracies and subtle variations in object form would tend to broaden any correlation, but it is unlikely that they could account for the complete loss of a visible trend. A factor which could have a significantly greater effect is a variable labelling efficiency which would result in a different effective size and shape in the two colours even if the underlying form was the same. Given the stochastic nature of incorporation (the mRFP, GFP and endogenous forms will all have a certain, unique, probability of being incorporated into the kinetochore complex) and the observed variations in signal strength, this is not at all unlikely.
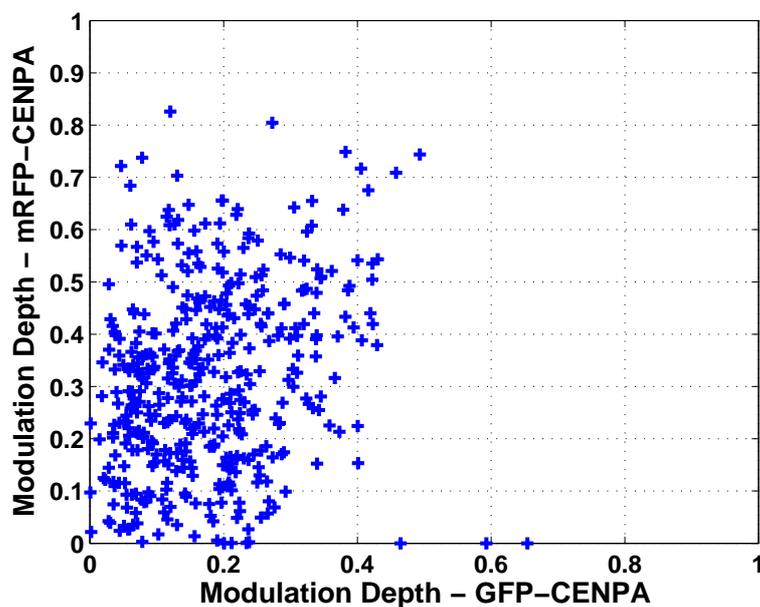
Figure 13.4: Scatter plot showing the relationship between the measured modulation depth of individual loci in both channels. For a uniform shape (e.g. spherical) and a varying size distribution one would expect a strong positive correlation, such a correlation is absent in the data.

GFP-CENP-C and mRFP-CENP-A

In the hope of detecting differences between the distribution of CENP-A and CENP-C in the kinetochore complex, cells were analysed in which CENP-C was conjugated with GFP and CENP-A with mRFP. Qualitatively the signal and expression levels were similar to those in the CENP-A, CENP-A control. The distributions obtained in quantitative analysis are shown in Figs. 13.5 and 13.6. A scatter plot comparing the individual kinetochores in the two channels is show in Fig. 13.7. The results are very similar to those of the control measurements, with no evidence for a difference in the distribution of the proteins.
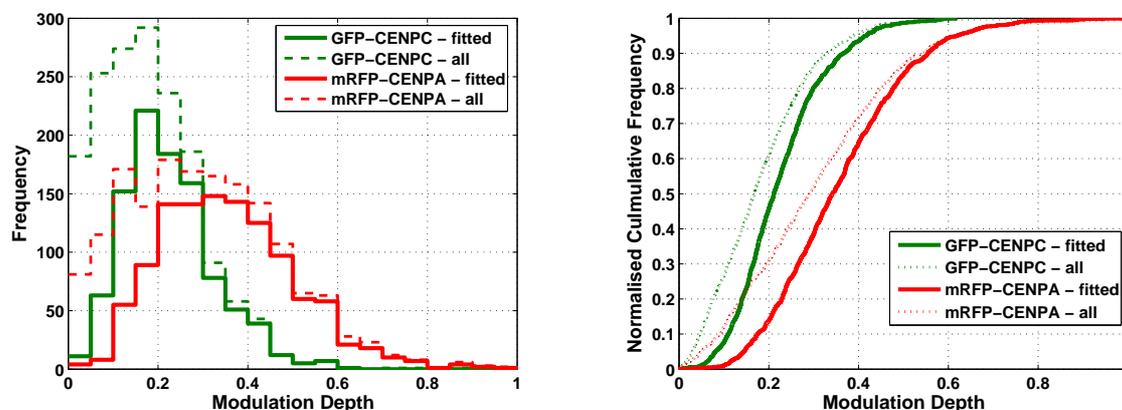
Figure 13.5: Distribution of measured modulation depths for the GFP-CENPC (488nm excitation) and mRFP-CENPA (568nm excitation) labellings. **Left:** histogram, **Right:** cumulative histogram. Note that the data not categorised as *fitted* should not be considered to be reliable, it is however useful to obtain an impression as to what the trend would be if the points were included.


Figure 13.6: Distribution of size estimates based on a spherical model for the GFP-CENPC (488nm excitation) and mRFP-CENPA (568nm excitation) labellings. **Left:** histogram, **Right:** cumulative histogram.
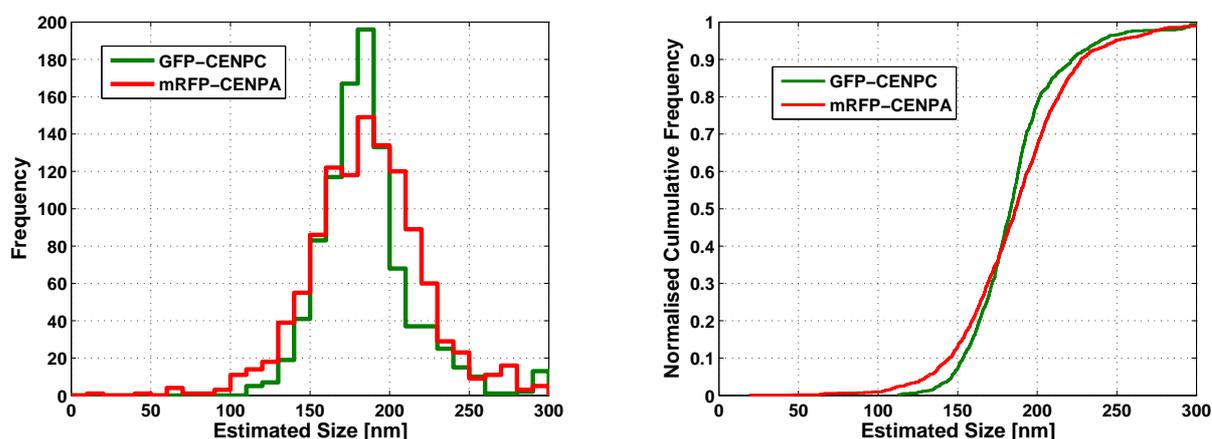
Figure 13.7: Scatter plot showing the relationship between the measured modulation depth of individual loci in both channels for the CENP-C vs CENP-A measurements. For a uniform shape (e.g. spherical) and a varying size distribution one would expect a strong positive correlation, such a correlation is absent in the data.
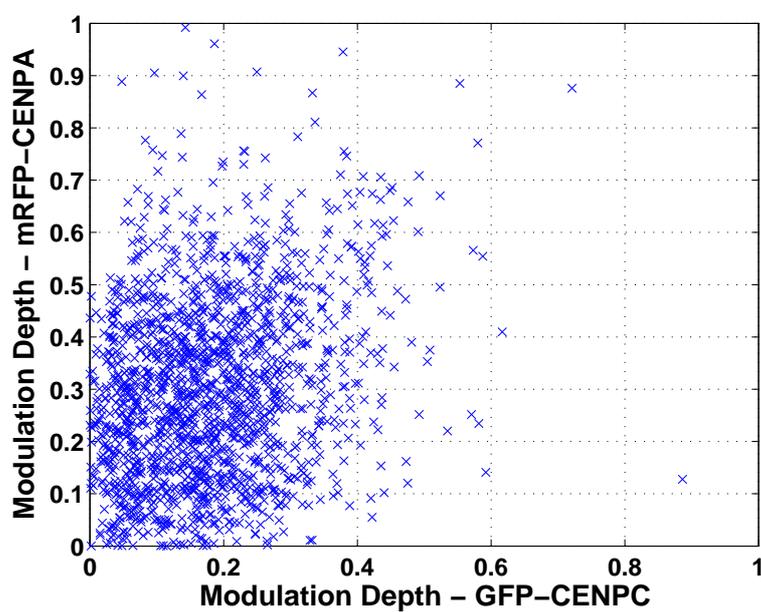
13.3 Simulation of 300nm disc

Whilst the observed results would be consistent with a relatively homogeneous, approximately spherical dye distribution, many other distributions could be possible. As the EM images of the metaphase chromosomes indicate a disc or plate like structure for the inner kinetochore a comparison to simulation results of such a structure should be instructive. A disc with a diameter of 300nm and a thickness of 30nm was simulated, rotated to a random orientation, convolved with the SMI PSF, and evaluated with the SMI image analysis algorithms as per the gene structure simulations in Chapter 12. The resulting modulation depth and size distributions are shown in Fig. 13.8.
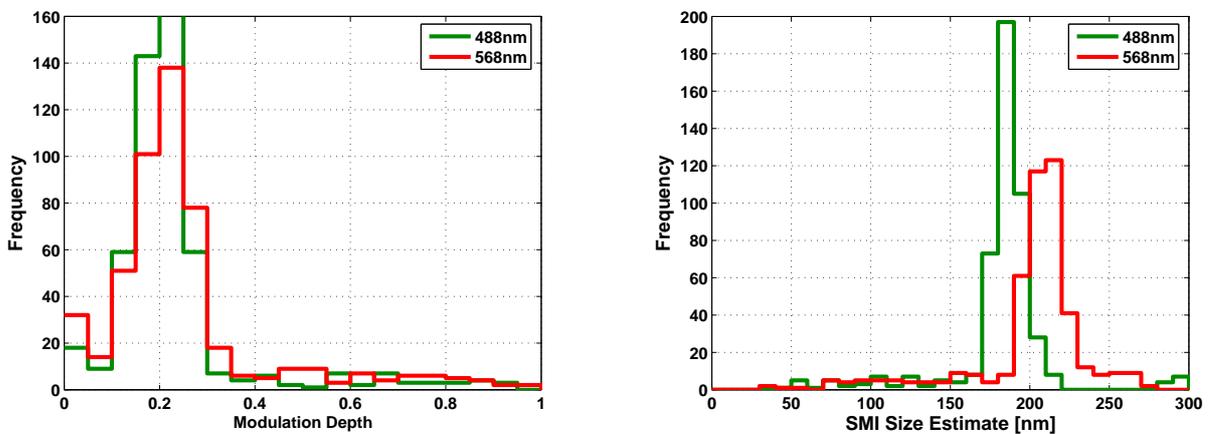
Figure 13.8: Left: modulation depth at 488nm and 568nm for a simulated 300nm disc with a random orientation. **Right:** corresponding 'size' estimates.

Despite a significantly non-spherical geometry, the disc structure leads to distributions which are similar to the measured ones. That said, there is a significant difference between the size estimates at the two wavelengths, making the flat disc model an unlikely explanation of the experimental results, and implying that the true structure could be more compact. One cannot, however, rule out a crumpled disc, such as might be encountered when the extra stiffness given by the outer kinetochore layer during metaphase was lost. When the experimental data is compared with the gene structure simulations (Chapter 12), the results are also similar enough to be a little disconcerting.

13.4 Distance Measurements

In addition to the 'size' measurements detailed above, distance measurements were made between the two kinetochore components. A Gaussian fit was used to obtain the lateral object positions (with the modification including a linear background term - this improved the reliability of the fit) and an intensity centre computed from the axial profile to obtain the z-position. The SMI information was not used in the distance determination due to

Figure 13.9: Distribution of measured 3D distances between the GFP and mRFP labelled signals in the CENP-A double labelling experiment. As the same protein is labelled in both colours, the 'true' distance should be zero and the distances seen here are either the result of measurement error or the small structural differences due to the stochastic nature of the labelling. This distribution is thus indicative of the overall sensitivity of the distance measurements.
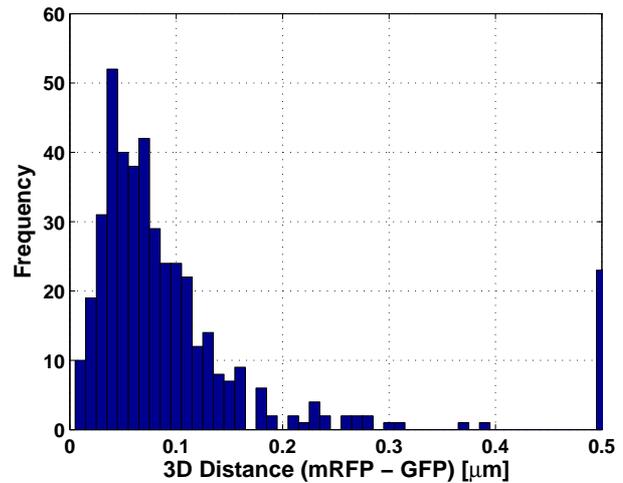
Figure 13.10: Distribution of measured 3D distances between the GFP-CENP-C and mRFP-CENP-A signals. As the form of the distribution is approximately the same as in Fig. 13.9, one can say that no detectable difference in the positions of the CENP-C and CENP-A protiens is observed.
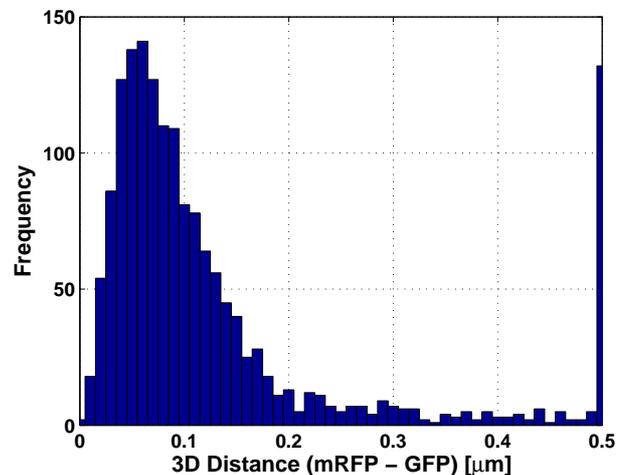
the absence of suitable calibration objects in the sample preparation. In retrospect it seems unlikely that this would have brought any additional information due to the limited accuracy in the lateral direction.

The presence of several, multiply labelled objects in the cell nucleus allowed a in-situ correction of the mean chromatic shift. Although we were interested in the distance, if any, between the two components, one can assume that the orientation of the individual kinetochores is random and thus the true distances will average to zero. It is thus possible to auto-calibrate by subtracting the mean distance of all loci in the cell from each of the individual distances.

When each cell is processed in this way the results seen in Fig. 13.10 are obtained. These are indistinguishable from those obtained for the calibration measurements as shown in Fig. 13.9. We can thus conclude that any distance, if present, is smaller than our measurement accuracy. With reference to the calibration measurements, an accuracy of around 50nm can be assumed.

Chapter 14

Evaluation of Confocal and Widefield Images

14.1   Lateral Size Estimates Using the Gaussian Fit

As many, if not most, labelled gene loci are technically too large to give sensible size estimates using the SMI microscope,[1] and yet not large enough to be well resolved in the confocal we need a method which is capable of extracting the limited size information that can be obtained from the confocal measurements. One such candidate is the lateral Gaussian fit introduced in Chapter 7, originally in the context of SMI lateral position measurements. In order to obtain size estimates we however use the width of the fitted Gaussian. In [92], it is shown that the width of a Gaussian fitted to bead images does not show a one to one correspondence with the size of the measured object. Indeed the smallest width measured corresponds approximately to the size of the PSF. The results are compared to SMI measurements, showing a nice one to one correspondence between measured and actual sizes[2]. Whilst an exemplary piece of propaganda for the SMI microscope, the comparison was hardly fair. A very naive interpretation of the confocal data was compared with a sophisticated method of analysing the SMI data, incorporating knowledge of both PSF and object form. What is significant, but not mentioned, in the confocal data is that, at least for the larger sizes, the relationship between the measured and real size was both monotonically increasing and smooth, implying that it should be possible to correct the measurements to arrive at a somewhat better approximation to the correct object size.

14.1.1   How to Correct Lateral Size Measurements

The reason that the measured FWHM is not equal to the object FHWM is that the object is convolved with the PSF and thus appears larger. If both object and PSF are assumed to be Gaussian (not a bad assumption regarding the PSF, especially seeing as our Gaussian

---

[1]This is a very bold assertion, for a thorough discussion please read Chapter 12

[2]Despite an incorrect object model and the failure to compensate for a significant systematic error in the SMI microscope. ([104])

is fitted to the average of several slices around the focus - implying that the original Bessel function shape will not be conserved - indeed through the central limit theorem it should tend towards a Gaussian[3]. As regards the object, in the absence of any other information a Gaussian is as good a model as any.) this correction is easily performed as the convolution of two Gaussians is another Gaussian with a variance equal to the sum of the two variances. One can thus perform the correction using the following simple formula:

$$\sigma^2_{\text{real}} = \sigma^2_{\text{measured}} - \sigma^2_{\text{PSF}} \tag{14.1}$$

The PSF width can in turn be determined using calibration beads and reversing the above calculation:

$$\sigma^2_{\text{PSF}} = \sigma^2_{\text{measured}} - \sigma^2_{\text{real}} \tag{14.2}$$

14.1.2   Relationship Between σ and Size

The standard deviation of an object is not a parameter that we can easily associate with size. For a Gaussian, the obvious size parameter is the Full Width at Half Maximum (FWHM). Setting $e^{-\frac{x^2}{2\sigma^2}} = \frac{1}{2}$ gives the textbook relationship between σ and the FWHM:

$$\text{FWHM} \approx 2.355 \times \sigma \tag{14.3}$$

When comparing to (current) SMI measurements, or bead sizes, it is important to note that FWHM is not equal to diameter. The FWHM of a sphere, when projected onto one axis, is:

$$\text{FWHM}_{\text{sphere}} = \frac{d}{\sqrt{2}} \tag{14.4}$$

where d is the diameter of the sphere.

In order to estimate a volume from the standard deviation (σ) of an object we must assume some object form. A Gaussian object is in this case not particularly useful, as it has infinite extent, and hence infinite volume - regardless of width. The obvious remaining candidate is a sphere, giving

$$V \approx \frac{4}{3}\pi \left(\frac{\sqrt{2} \times 2.355 \times \sigma}{2} \right)^3 \tag{14.5}$$

14.1.3   How Accurate is our Size Estimate?

Once we know how to correct for the PSF size, we should, assuming a perfect knowledge of the PSF and infinite photon number, be able to measure arbitrarily small objects. It is however interesting to know what the practical limitations on lateral size measurements imposed by both the accuracy of our PSF knowledge and signal to noise characteristics

---

[3]At a first glance, this would appear to be a misuse of the central limit theorem. If one however regards each slice of the PSF as a probability distribution ...

Nominal Diameter [nm]	Estimated Diameter $\pm$ s.d. [nm]
44	136 ± 75
100	122 ± 45
190	188 ± 48
390	357 ± 13

Table 14.1: Size estimates of beads obtained from confocal images using the lateral Gaussian fit

are. It turns out that for objects smaller than the resolution limit, errors in both PSF and in the uncorrected size are amplified. These relations are easily demonstrated using Gaussian error propagation.

From Eqn. 14.1 we obtain trivially the expression

$$\sigma_{\mathrm{real}} = \sqrt{\sigma_{\mathrm{meas}}^2 - \sigma_{\mathrm{PSF}}^2} \tag{14.6}$$

The error, $\Delta\sigma_{\mathrm{real}}$, is thus given by

$$\|\Delta\sigma_{\mathrm{real}}\|^2 = \left\|\frac{\partial}{\partial\sigma_{\mathrm{meas}}}\left[\sqrt{\sigma_{\mathrm{meas}}^2 - \sigma_{\mathrm{PSF}}^2}\right]\Delta\sigma_{\mathrm{meas}}\right\|^2 + \left\|\frac{\partial}{\partial\sigma_{\mathrm{PSF}}}\left[\sqrt{\sigma_{\mathrm{meas}}^2 - \sigma_{\mathrm{PSF}}^2}\right]\Delta\sigma_{\mathrm{PSF}}\right\|^2 \tag{14.7}$$

Which, on evaluation of the partial derivatives, gives

$$\|\Delta\sigma_{\mathrm{real}}\|^2 = \left\|\frac{\sigma_{\mathrm{meas}}}{\sqrt{\sigma_{\mathrm{meas}}^2 - \sigma_{\mathrm{PSF}}^2}}\Delta\sigma_{\mathrm{meas}}\right\|^2 + \left\|\frac{\sigma_{\mathrm{PSF}}}{\sqrt{\sigma_{\mathrm{meas}}^2 - \sigma_{\mathrm{PSF}}^2}}\Delta\sigma_{\mathrm{PSF}}\right\|^2 \tag{14.8}$$

As we are really interested in how the error depends on object size, we substitute $\sigma_{\mathrm{meas}} = \sqrt{\sigma_{\mathrm{real}}^2 + \sigma_{\mathrm{PSF}}^2}$ giving

$$\|\Delta\sigma_{\mathrm{real}}\|^2 = \left\|\frac{\sqrt{\sigma_{\mathrm{real}}^2 + \sigma_{\mathrm{PSF}}^2}}{\sigma_{\mathrm{real}}}\Delta\sigma_{\mathrm{meas}}\right\|^2 + \left\|\frac{\sigma_{\mathrm{PSF}}}{\sigma_{\mathrm{real}}}\Delta\sigma_{\mathrm{PSF}}\right\|^2 \tag{14.9}$$

Looking at the contributions of the individual terms, it is clear that in both cases a very small object will be associated with a large error. In the limit of large object sizes the contribution of the error in the PSF estimate tends to zero, and that of the fit to one. The error is hence simply that of the fit. These relationships are shown in Fig. 14.1.

14.1.4   Testing and Calibration on Beads Measurements

Susanne Fenz made and analysed several confocal measurements on 44, 100, 200, and 390nm green fluorescent beads. Calibration of the PSF width was performed using one of the 190nm bead series. The most essential results are summarised in Table 14.1.

All results except those for 44nm beads show a relatively good agreement with the manufacturer specifications. The accuracy of the measurements is however severely limited through the poor signal to noise ratio offered by the confocal microscope. These
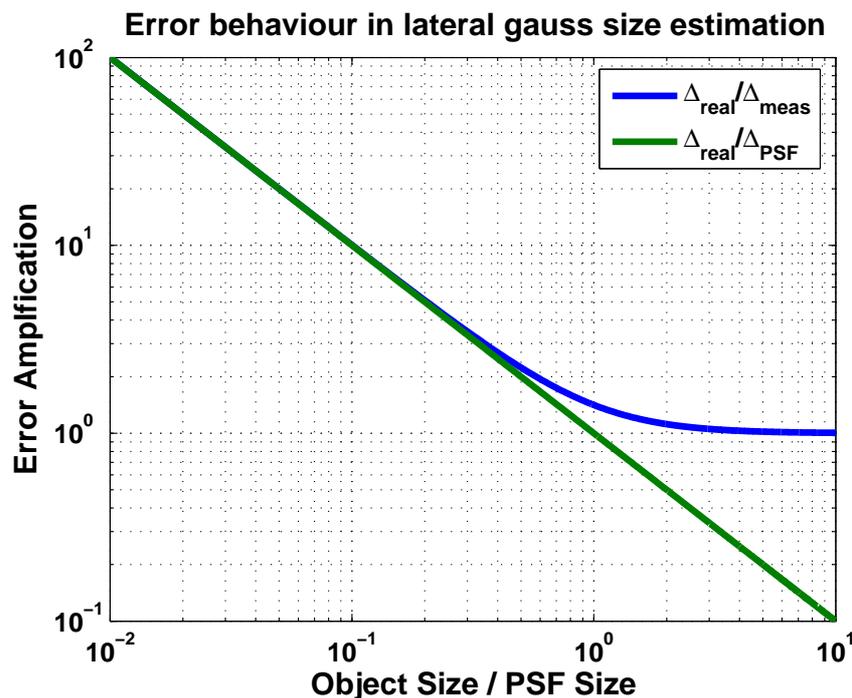
Figure 14.1: Contribution of error in fit and error in PSF estimate to the error in corrected size. One can see that the effect of the measurement errors are very large for small objects and decrease with increasing object size. At large object sizes the error starts to be dominated by the error in the size measurement and is not as reliant on an accurate PSF estimate.

results indicate that, although not as accurate as SMI measurements (capable of $\pm \sim 5$nm - see Section 10.1), confocal size measurements are nevertheless capable of detecting size differences considerably below the optical resolution. At the larger (390nm) bead sizes, a small bias in the estimated sizes is observed. The best explanation of this bias is that the assumption of Gaussian object form is starting to have a noticeable effect.

When applied to widefield images, or in the following example the lateral widefield information present in an SMI image, similar results are possible. This is shown in Figures 14.2 and 14.3. The large difference between the reproducibility and the distribution width additionally suggests that it should be possible to considerably improve the measurements by addressing further systematic errors.

14.1.5   Application to Biological Measurements

The lateral Gaussian fit has been applied successfully to several biological specimens, including FISH labelled gene loci and regions of centromeric DNA. These measurements are largely the work of Yanina Weiland and Claudia Batram, and I do not intend to discuss the results in detail here. It is however interesting to touch on a few general points, and to provide a few comments on accuracy and calibration. The most striking observation is that all gene loci have measured significantly larger than the 200nm upper limit for SMI
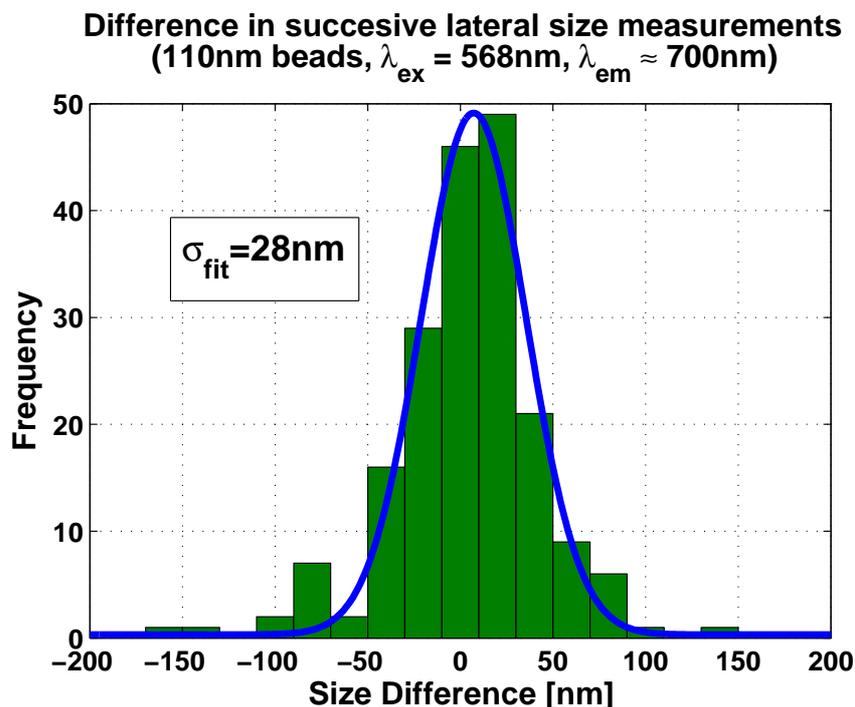
Figure 14.2: Reproducibility of lateral size measurements using widefield detection in the SMI microscope. The measurements were carried out on 110nm diameter deep red fluorescent beads from Molecular Probes (cat. num. T-8878). These beads are excitable throughout the visible spectrum and fluoresce in the deep red. Note that only the emission wavelength is relevant in widefield measurements, indicating an accuracy of ∼1/20th of a wavelength.

microscopy.

Accuracy and Calibration

The accuracy of the size estimates obtained with the lateral Gaussian fit is very sensitive to the quality of the PSF measurement. This must be performed using the same objective, at the same excitation and emission wavelengths, and, in the case of confocal measurements, with the same pinhole settings as the sample to be measured. It is well known that the PSF degrades significantly with increasing distance from the coverslip, thus when attempting to perform quantitative measurements it is highly desirable, if not essential, to prepare the specimen on the coverslip. When, for whatever reasons, this is not practical, an alternative would be to prepare beads on the slide with the sample. Note that as the distance between slide and coverslip and the exact composition of the mounting medium is difficult if not impossible to control, the preparation of a separate slide with beads is generally unable to give a useful PSF estimate. This is well illustrated in the comparison between two beads slides prepared under similar conditions giving PSF widths of 230 and 280nm for beads. PSF predictability is unlikely to be improved in biological specimens.

From the analysis of the error characteristics it is clear that the dependance on an

Measured lateral size of 110nm dark red beads
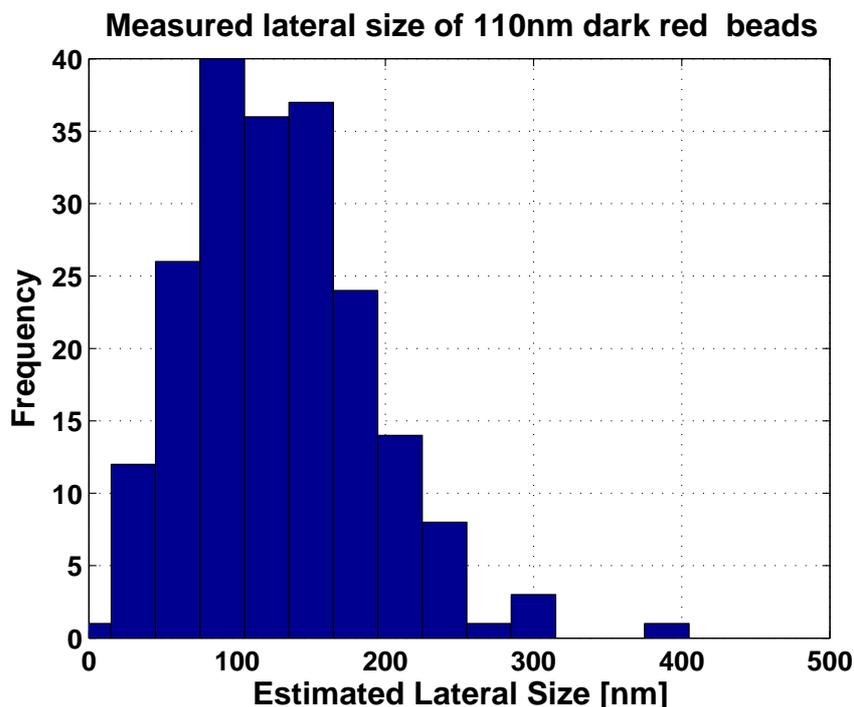
Figure 14.3: Distribution of lateral size measurements, from the same measurements as in Fig. 14.2. Note that the distribution shows a much larger spread than the statistical limit derived through the reproducibility measurements. This suggests the presence of a systematic error, for example a non-uniform PSF or a susceptibility of the fit to neighbouring structures. It is also possible that the measured distribution represents the real size distribution of the beads, although the accompanying SMI measurements render this highly unlikely.

accurate PSF is significantly reduced for larger object sizes. When the objects are thus approximately as large as the PSF or larger, some error in the PSF measurement can be tolerated. The error will be passed on to the measurements, but will not be amplified as it would for smaller objects. Using, for example, a PSF estimate of 220nm when the real PSF was 250nm would thus result in the following errors (calculated by generating 'measured' sizes with a σ_{PSF} corresponding to 250nm FWHM and converting back to real sizes using a σ_{PSF} corresponding to 220nm; the sizes and their errors are expressed in terms of their FWMHs):

Real object size [nm]	Measured size [nm]	Error [nm]
10	119	109
50	129	79
100	155	55
200	233	33
400	417	17
800	809	9

In other words, whilst one should try and obtain the best possible PSF estimate, the re-

quirements can be relaxed somewhat when one is dealing with larger objects. It is however important to bear the error behaviour in mind when comparing different measurements, especially when different preparation or mounting procedures have been used.

14.2 Ten Point Model

Whilst a lateral Gaussian is capable of delivering good size estimates for relatively small objects, the applicability of the model becomes limited as soon as the object cannot be considered diffraction limited. Threshold based measurements, however, require that the size of the object is significantly ($\sim$ 5 - 10x) larger than the resolution. Most gene loci, at least when using BAC based probes, fall into the gap between these two methods, with visible yet poorly resolved lateral structure (for example images of gene loci see Fig. 14.9). This lateral structure is then completely ignored when fitting a lateral Gaussian. In order to extract the largest amount of information from these structures we need a technique that is aware of the PSF, yet does not inflict unreasonable constraints on the form of the structure.

One possible approach would be to deconvolve the image and then apply threshold based methods. Whilst this is certainly better than using threshold based method on the original images, it is still not ideal as the resolution is still limited and, when iterative procedures have been used, unlikely to be constant over the entire image. Another disadvantage of threshold based methods, with or without deconvolution, is that they are typically highly subjective with the result strongly dependent on the selected threshold.

Another approach is to fit a model which allows the object to have more structure than a simple Gaussian. One such model would approximate the object as a finite number of equally bright points. The obvious analogy here is to fluorophores - if one can determine the position of every individual fluorophore in a sample, one can 'image' with a resolution which is theoretically only limited by photon statistics and labelling density. The concept is however more general, each point can represent many fluorophores, in which case we loose some structural resolution in exchange for better signal levels.

14.2.1 Implementation of a multiple point model

What does our image look like?

When we assume that our object consists of many points, it follows that our image will be the sum of many individual PSFs at the locations of the points, weighted with the intensity of a single point.

$$I(\mathbf{r}) = \sum_i \delta(\mathbf{r}_i) * I_{\text{PSF}}(\mathbf{r}) \tag{14.10}$$

In order to fit this model to the image, especially when we want sub-voxel accuracy, we need a functional description of the PSF. This function should however be computationally inexpensive as it will typically be evaluated thousands of times during a fit. For confocal measurements the PSF can be reasonably approximated with a 3D Gaussian, which allows

a reasonable computation time. The model is thus:

$$I(x, y, z) = A \sum_i e^{-\frac{x_i^2}{2\sigma_{xy}^2} - \frac{y_i^2}{2\sigma_{xy}^2} - \frac{z_i^2}{2\sigma_z^2}} + c \tag{14.11}$$

with σ_{xy} and σ_z chosen to give sensible lateral and axial FWHMs[4] . The normalisation contained in the factor A, a constant offset c, and the positions of the individual point (x_i, y_i, z_i) are thus the free parameters in the fit.

How many points?

The choice of the number of points to fit is a compromise. The more points, the more accurately the model can represent arbitrary objects. However speed and stability of convergence decrease drastically with an increasing number of points and the influence of noise becomes more pronounced. It is also not sensible to use more points than one has fluorophores, and indeed one must be careful when one is in the same order of magnitude - fitting 5 points to 7 fluorophores is unlikely to be very successful (in confocal microscopy the signal to noise ratio kills you long before you have to worry about this). If there are either many fluorophores per point (or, though not highly sensible, many points per fluorophore) these quantisation problems are not so much of an issue.

In practical terms it is reasonable to choose the minimum number of points with which one can obtain a reasonable estimate of the structure. This number depends on the size of the structure. For the smaller gene regions and for beads, five points is enough to obtain a reasonable estimate of structure and size. Ten points allow slightly larger structures to be investigated and appear sufficient for most if not all gene loci. Tests have been performed with up to 30 points, with dramatically increasing fit time. Whole chromosome territories would require significantly more points[5], and it is unlikely that convergence could be achieved in a reasonable amount of time. These structures however are large enough for threshold based methods.

14.2.2   Results

Tests on simulated data

A ten point version of the multi-point model was tested against images of 200kbp gene domains simulated using the worm like chain model. Thirty virtual fluorophores were placed on the simulated structure and the resulting image was convolved with a theoretical widefield PSF (Figure 14.4). Several representative structures are shown in Fig. 14.6. A good correlation is observed between the standard deviation of the 'fluorophore' positions and that of the points in the fitted model, despite the difference between the number of virtual fluorophores and the number of points fitted. This relationship is shown in Figures 14.7 and 14.8.

---

[4]250nm and 600nm respectively are reasonably sensible values for a typical confocal when using 488nm excitation.

[5]at a guess approximately 100
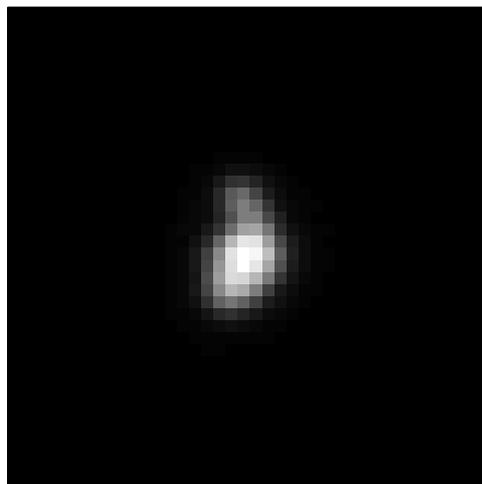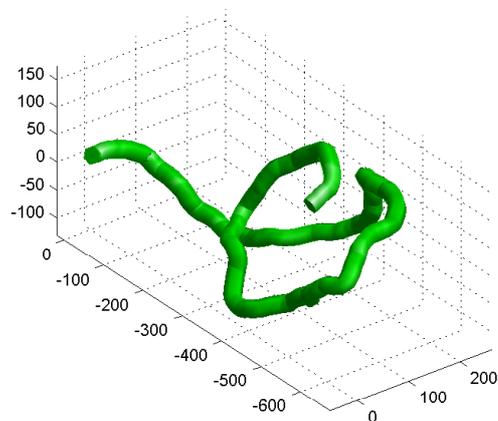
Figure 14.4: Left: 200kbp of DNA simulated as 30nm fibre using the worm like chain model. The axes are in nm. **Right:** Mean intensity projection after convolution with a theoretical confocal PSF (1 pixel = 78x78nm).
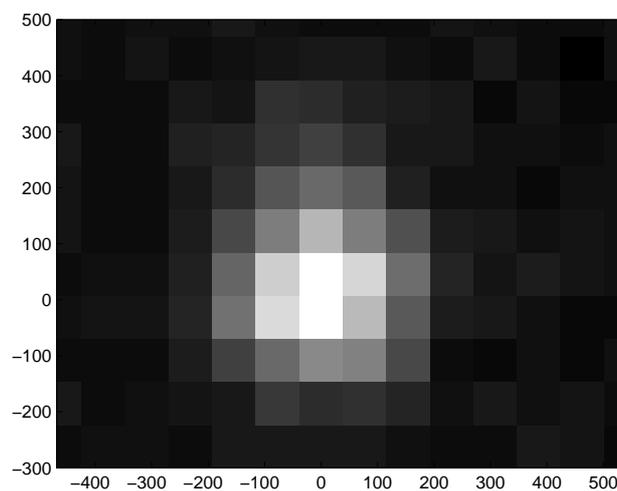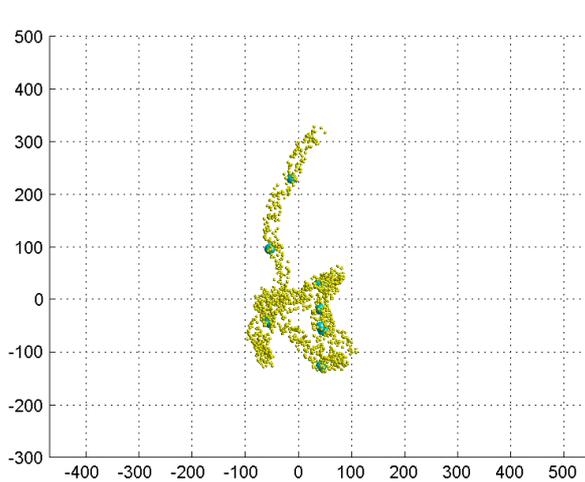
Figure 14.5: Left: Top down view of one 200kbp 30nm fibre segment showing positions of the virtual fluorophores (yellow) and of the fitted 10 point model positions. **Right:** mean intensity projection of corresponding confocal image. The axes are in nm.
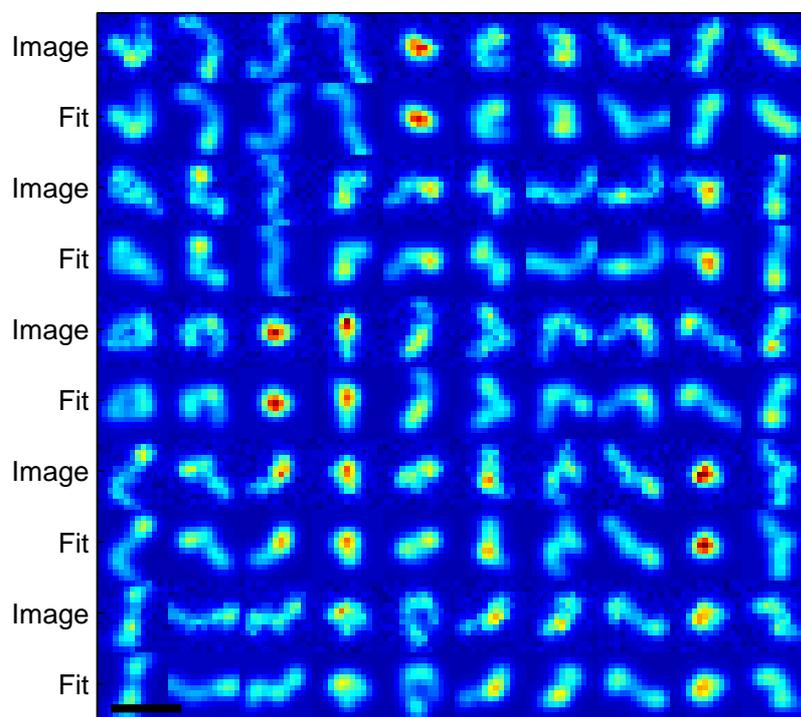
Figure 14.6: Mean projections of 50 simulated confocal images of a 200kbp 30nm fibre segment and the accompanying 10 point fits. One can see a good match between the fit results and the original image. The scale bar is 1μm.

Application to experimental data

The ten point model has been tested on various confocal data sets, both from beads[6] and gene loci[7]. The mean intensity projections shown in Fig. 14.9 show that the ten point model is capable of accurately reproducing the structure of gene loci which are large enough to show some lateral structure.

14.2.3   From Point Positions to Size, Volume or Form

The positions of the fitted points contain a lot of information about the object of interest, but not however in a form which is easily interpretable. In order to obtain structural parameters which we understand, we must further massage the data. Several parameters are interesting and we choose to concentrate on extension and volume, these being the most

---

[6]kindly provided by Susanne Fenz

[7]kindly provided by Claudia Batram
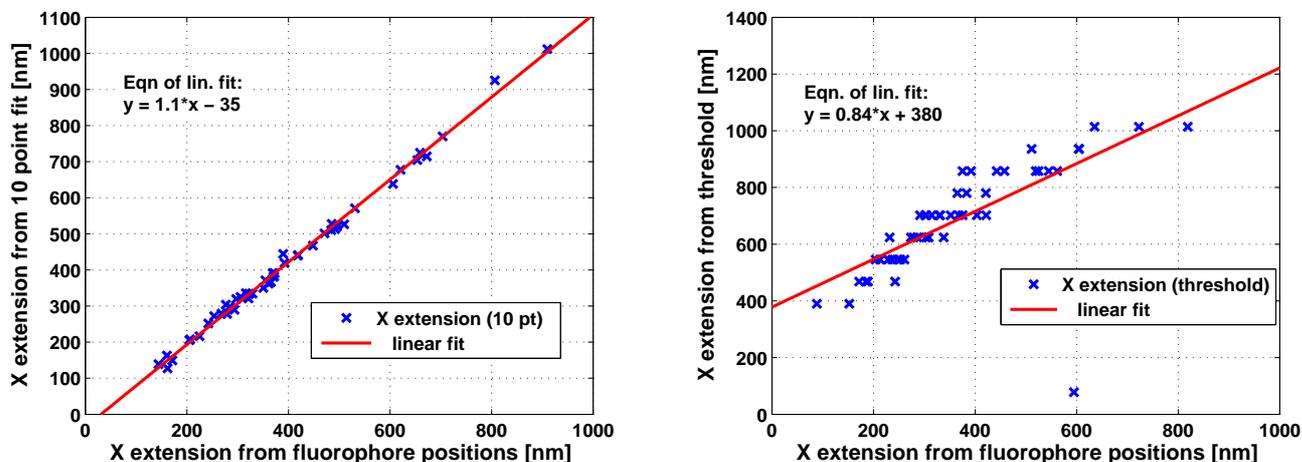
Figure 14.7: Extent (FWHM) of simulated gene structure in X direction as measured from the virtual microscope images compared to the extent calculated from the locations of the virtual *fluorophores*. **Left:** Using the 10 point model. **Right:** Using an automatic thresholding procedure.
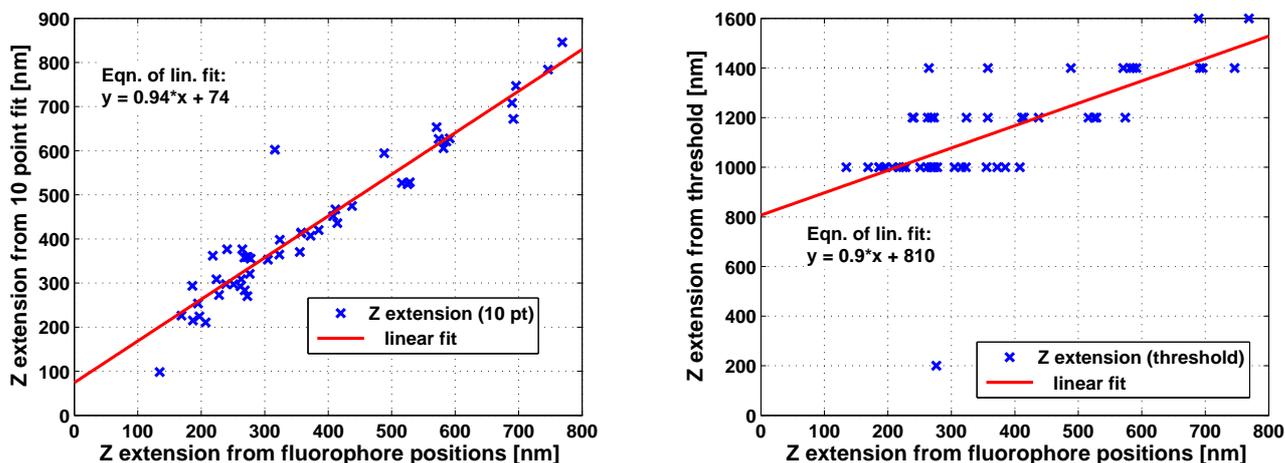
Figure 14.8: Extent (FWHM) of simulated gene structure in Z direction as measured from the virtual microscope images compared to the extent calculated from the locations of the virtual *fluorophores*. **Left:** Using the 10 point model. **Right:** Using an automatic thresholding procedure.
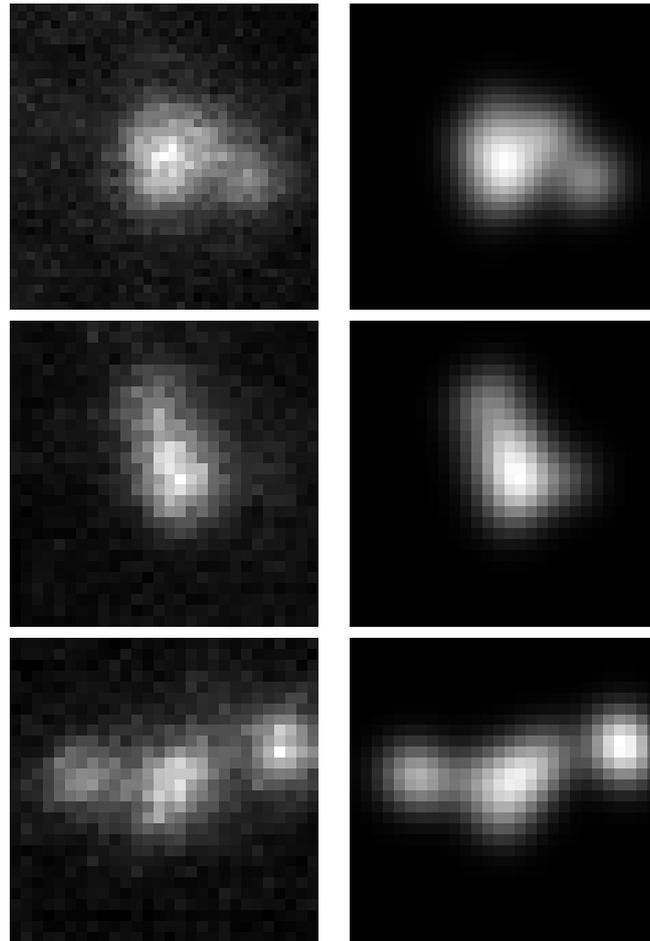
Figure 14.9: Mean intensity projection of three representative gene loci (**left**) and of the fitted 10 point models (**right**), showing that both size and shape of the labeled loci are captured in the fitted model (1 pixel = 78x78nm).
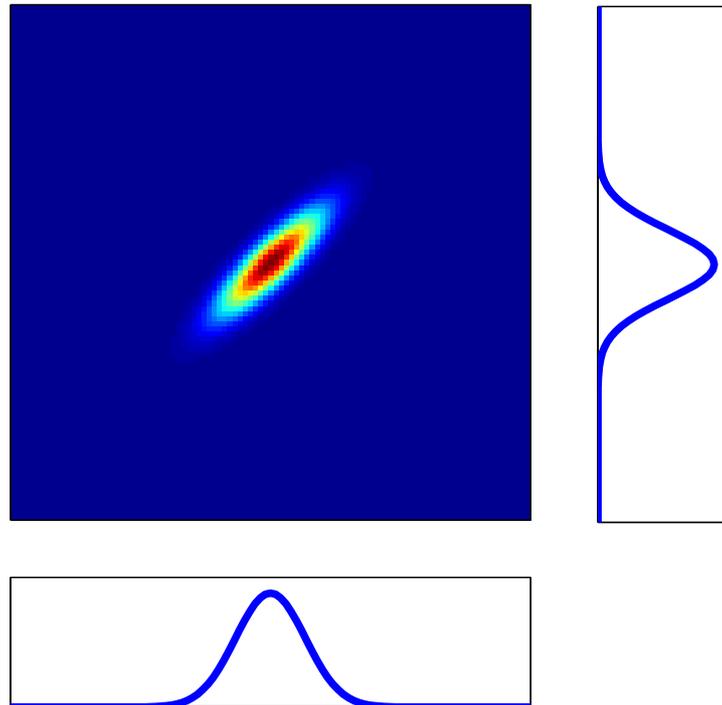
Figure 14.10: Non uniform fluorophore distribution with projections onto x and y axes.

obvious parameters for the study of gene compaction. Unfortunately these parameters are not unambiguously contained within our data and there are several possible ways to obtain a mathematically sensible estimate of them from our data, all delivering slightly different values.

Extension

The extension, or spread, of an object is essentially what is delivered by the size estimates delivered by the lateral Gaussian and SMI measurements. Given simple objects, for example spheres or cubes, the extension of an object is easily described by the diameter or side length respectively. For more general objects it is difficult to find a sensible measure, indeed, perhaps the only measure of spread that is defined for an arbitrary object being its standard deviation.

The most obvious way to extract the standard deviation of the object from our fitted data is simply to take the standard deviations of the x, y, and z components of our point positions. This gives us the spread of the object along the three axes separately. A 2D or 3D spread can be estimated by averaging the respective single axis spreads. The results obtained using this method are mathematically the most similar to those obtained from the fit of a lateral Gaussian. This method, however, is not without its disadvantages. Given a fluorophore distribution which is non uniform, such as in Fig. 14.10, the standard deviations of the projections on the axes provide limited information about the object itself.

One approach to non-uniform objects is to establish the major and minor axes of the object and project the point positions onto this new coordinate system. In addition to a general extension, this method gives a measure of object anisotropy. Another, and mathematically equally good, method is to base the estimate of object spread on the expectation value of the distances from the individual points to the centre of the object.

Object extensions which are expressed as either standard deviations or average distances are quite different to the size parameters such as diameter or FWHM with which one is usually familiar. It is thus useful to express the extensions in terms of a familiar quantity such as the FWHM. The method I have used to convert between standard deviation or mean distance based measures and equivalent FWHMs is as follows: As we know nothing about the underlying distribution, the transformation is calculated for a Gaussian distribution[8]. For standard deviation based measures, the mapping is, as stated earlier, $FWHM_{\text{equiv}} = 2 \times \sqrt{2\ln(2)} \times \sigma$. For the mean distance measurements, the expected, or mean, distance, $< |\mathbf{r}| >$ for a given Gaussian can be calculated from the following integral[9]:

$$< |r| > = \frac{\int_0^\infty r \cdot e^{-\frac{r^2}{2\sigma^2}}\ dr}{\int_0^\infty e^{-\frac{r^2}{2\sigma^2}}\ dr} \tag{14.12}$$

This gives the relationship

$$< |r| > = \frac{2\sigma}{\sqrt{2\pi}}$$

or

$$FWHM_{\text{equiv}} = 2\sqrt{\pi\ln(2)} \cdot < |r| >$$

Volume

When the underlying question is the compaction of gene regions, volume and not extension is arguably the relevant parameter. A naive, although not necessarily invalid, approach to the estimation of object volume is to transform the object extension into an effective diameter (for a sphere the diameter is approximately $\sqrt{2} \times$ FWHM, with FWHM $\approx 2.355 \times \sigma$), and to compute the volume using the normal formulae for sphere volume. It is likely that this method will overestimate the volume of non-uniform objects as in Fig. 14.10. A slightly more sophisticated method is to obtain the convex hull of our fitted points (ie the smallest convex polyhedron which contains all points). A diagrammatic representation of the convex hull method is shown in Fig. 14.11. This method will perform better on arbitrary shapes than the one based on object extension, although the volume of objects with indentations will still be overestimated.[10] In addition to a volume, the convex hull also provides an estimate of object surface area.

---

[8]the standard 'default' distribution

[9]one can assume the distribution is centred at the origin without loss of generality.

[10]It is uncertain whether this will be a problem for compaction measurements - the volume of the DNA itself is known, and not what we are really interested in measuring - we want to know the space it occupies. There will necessarily be lots of free space (at least on an atomic resolution scale), and the convex hull volume is thus not an unsuitable measure.
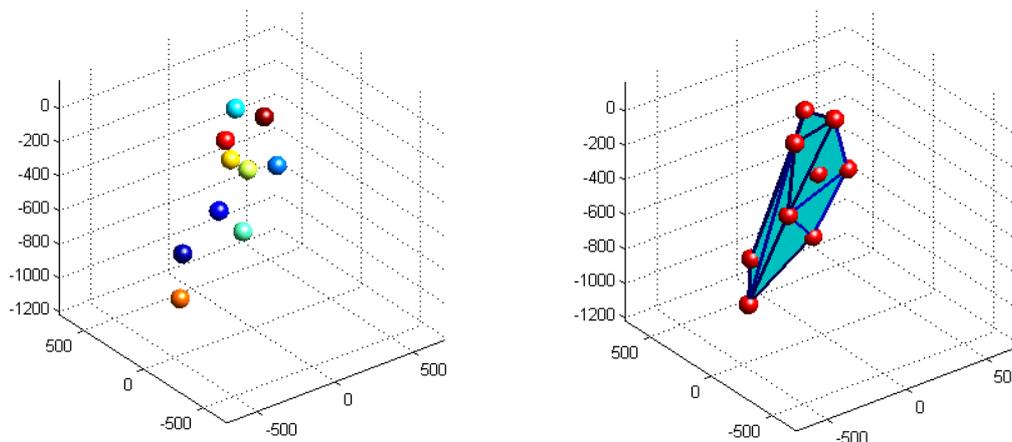
Figure 14.11: Convex hull based volume estimation. **Left:** Point locations obtained from fit. **Right:** The convex hull is the smallest convex polyhedron which encloses all points.

14.2.4   2D or 3D?

One of the ever returning questions in the quantitative analysis of microscopic data is whether to perform a 2D or 3D analysis. As the objects themselves are 3D, three dimensional data analysis is highly desirable. As the resolution along the z axis is however significantly worse than that in the x-y plane, the accuracy of any measurements along the z axis, and thus of 3D measurements, will be significantly worse than in two dimensions. The ten point model is intrinsically 3D, providing the 3D position of each point in the model. The question here is thus whether it is advisable to use the z-component of the position data when computing sizes and/or volumes. As there is no clearly correct answer here, both values are calculated. To obtain a volume estimate from the 2D convex hull calculation, the volume of a sphere with an equatorial plane of the same area is used.

The comparison of the two methods yields some interesting results. Taking the mean of individual axis standard deviations, the spread estimate based on the full 3D data is significantly larger than in the 2D case. It is also highly sensitive to the less accurate z axis measurement. In contrast, when the spread estimate is based on the distance to object centre, the difference between 2D and 3D estimates is very small, with the 3D estimate often being slightly smaller than that obtained in 2D. The volume estimates based on the convex hull calculation are also reasonably similar in both cases.

14.2.5   Characteristics of the Ten Point Model

As the ten point model is both a reasonably complicated system, and still relatively new, a rigorous quantitative characterisation has not yet been possible. The application to both bead and biological specimens has however allowed several empirical observations to be made:

- For very small, essentially diffraction limited objects (e.g. beads) increasing the number of points used in the model increases the estimated size. For larger objects (>200nm), this increase becomes less noticeable.

- The sizes of very small objects tend to be overestimated, and no longer follow a completely linear relationship with respect to the real sizes. Nonetheless, it is possible to separate 200 and 400nm beads, and with a little luck 100 and 200nm in confocal images.

- When an object is too large to be approximated by the chosen number of points a separation into individual points is observed in the resulting fitted image.

- A poor signal to noise ratio can result in points being assigned to noise spikes away from the object, and thus an incorrect size estimate.

- The observed axial spread is typically larger than the lateral spread. It is hard to say whether this is due to an underestimation of the axial PSF size, or the poorer axial accuracy which one would expect through the anisotropic PSF.

- The error in size measurements based on this method would seem to be in the order of 20-30%, with a floor at around 70-100nm. When the PSF approximation is tuned, using, for example, beads it is possible to further reduce this error.

Drawing a few guidelines from these observations;

- Objects for which no lateral structure are visible (diffraction limited) are better analysed with either the lateral Gaussian fit, or SMI microscopy, as these methods will give more accurate results.

- At the present stage, it is highly advisable to examine all the fit results by hand to verify that the object is neither too large, nor so noisy that points have been assigned to noise peaks. This is easily done by examining the z-projections of the original images and fit result.

- Bead measurements in the same optical conditions, on which to base the PSF approximation, should improve the accuracy. This is however less critical than for lateral Gaussian based measurements as, due to the larger object size, the tolerable error is also increased.

14.2.6   Application to Widefield and SMI

The same concept is in principle applicable to widefield images. On simulated images with small objects and a theoretical (NA 1.4) PSF, quite reasonable results are obtained. A 3D Gaussian is however a rather poor approximation for a widefield PSF, and when applied to real images the results are decidedly mediocre. Assuming a suitable, computationally inexpensive, approximation to the widefield PSF can be found, there is however no reason

why the method should not also be applied to widefield images. Somewhat better results than with the 3D Gaussian are obtained when assuming that the PSF is a Gaussian in the lateral direction with the same integrated intensity in each slice, but with a width which depends linearly on depth. This model however lacks the characteristic of retaining a (relatively) sharp edge in the defocussed slices along the aperture angle[11]. Incorporation of this characteristic into the PSF model is likely to be necessary in order to obtain a reasonable accuracy in the z direction.

Given a feasible widefield PSF approximation, one could easily incorporate the SMI axial modulation. This would allow sensible, and, given the additional axial information, potentially very accurate, size and form estimates on those objects for which lateral structure is visible (implying a size significantly larger than 200nm), and yet exhibit some modulation in the SMI microscope. Most gene regions fall into this category.

---

[11]The typical hourglass shape

Part IV

Outlook

Chapter 15

Outlook/Ideas

The aim of this chapter is to briefly outline a few of the ideas and experiments which came up during the course of my thesis but that I was unable to fully investigate. It is hoped that these might prove useful for future students.

15.1 SMI Structured Illumination

The possibility of combining the SMI with other forms of structured illumination to obtain either an improvement in resolution, in imaging accuracy, or in both has been previously raised [3]. Some form of lateral structured illumination also has the potential to remove one of the most serious current limitations, the lack of optical sectioning. This lack of sectioning currently restricts sensible measurements to cases where there are only a very small number of objects within the cell (e.g. gene loci) or, alternatively to very thin samples. For the replication foci measurements, for example, we had to resort to cryosections or alternatively hypotonic flattening to ensure optical separation of the points. The possibility of making such measurements in an intact cell would thus be highly desirable, warranting the introduction of further illumination structure, if only to provide optical sectioning.

15.1.1 Slit Scan SMI

As a first experiment with structured illumination on the SMI, I placed a variable slit in the illumination path at a conjugate image plane. This resulted in a single line of illumination, similar to what might be obtained from a slit scan confocal. It was possible to scan the object under this line by using the y-axis of the object piezo. A 2D image could then be reconstructed by taking the same line of pixels along the x-direction at several different shifts of the y piezo. Figure 15.1 shows the improvement in optical sectioning and hence contrast achieved with the slit-scan method.

Initial attempts were made to combine the slit-scan with SMI illumination. This required the use of a 4Pi like preparation with two cover slips and a smaller inter-objective separation. A second slit (improvised from a pair of razor blades) was introduced into the conjugate plane of the second objective. Whilst both modulation and sectioning could be
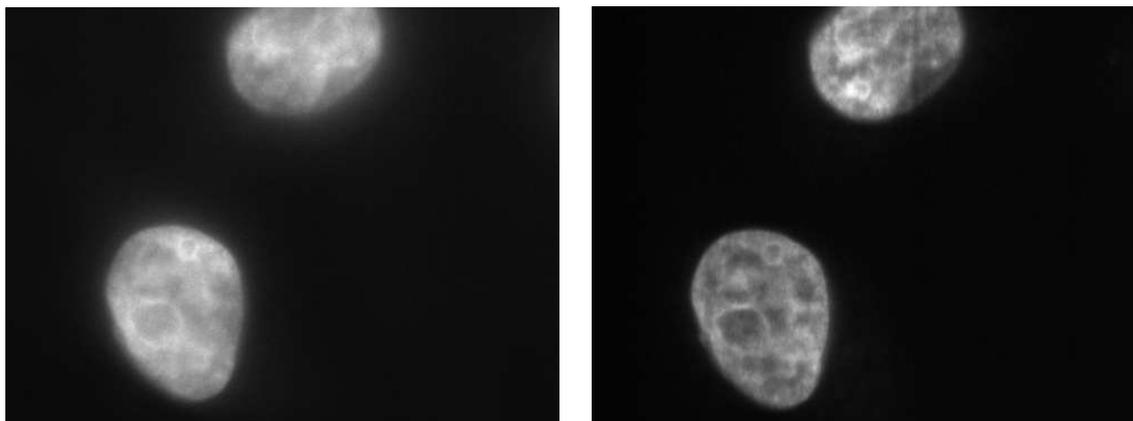
Figure 15.1: Widefield (left) and slit scan (right) images of a mid-section through GFP-H2A expressing HeLa cells (Specimen H. Albiez, LMU München). Both images were taken with the SMI microscope with one arm blocked. The improvement in contrast through use of the slit-scan technique is noticeable.

demonstrated on beads, it was anything but homogenous over the field of view. This can be attributed to the poor quality of the slits and alignment difficulties stemming from the rather makeshift nature of the setup for positioning them. When using well matched and properly adjustable slits there should be nothing to stop this method from functioning.

15.1.2 Outlook - Grating Based

A minor problem with the slit scan method is speed; one image must be obtained for each pixel in the y direction, multiplying the acquisition time by the size of the image in y. Grating based structured illumination also has the potential to introduce optical sectioning, as well as offering a significant increase in the lateral resolution. It should also be very much faster than a slit-scan approach, being only around 3-5 times slower than normal imaging. A grating based structured illumination could be implemented by introducing gratings in each of the arms at the location of the slits.

The reconstruction software which would be necessary for grating based SI is very similar to that required for SW-TIRF. I developed a simple reconstruction algorithm for use with the SW-TIRF microscope under development in the Cremer group, which might be able to serve as a basis for a full 3D reconstruction procedure for the combination of SMI and structured illumination. Figure 15.2 shows a simulation of the results which could be obtained using an adaptation of this algorithm and lateral structured illumination at three angles in a widefield microscope.
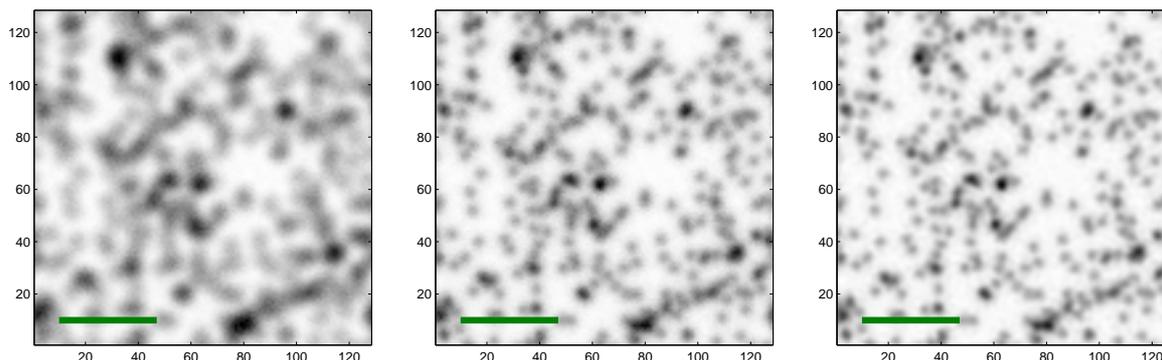
Figure 15.2: Simulated images showing the comparison between a widefield image of randomly distributed beads and a the images which could be obtained using lateral structured illumination (at angles of 0 and ± 60 degrees. From left to right: widefield image, reconstruction, reconstruction after Wiener filter. The scale bar is 1μm.

15.2 4Pi Abuse

There are several ways in which one could use the 4Pi microscope to do things which were not necessarily anticipated in the original design. A few scientifically interesting possibilities are outlined below.

15.2.1 Double Excitation 4Pi-A

The use of 2-photon excitation in the 4Pi requires very stable labelling. This can be quite restrictive when it comes to what types of objects can be measured. The normal solution to weakly labelled objects is to use 1-photon 4Pi type C. Unfortunately it seems unlikely that a type C upgrade for the commercial 4Pi will be available any time soon[1].

One possibility to allow 1-photon imaging with the existing 4Pi-A setup would be to use a dye which can be excited at two (well separated) wavelengths, or alternatively a double labelling. Images could then be taken at both wavelengths, resulting in two images in which the sidelobes will be at different distances. When observed in the Fourier domain, the peaks of the image at the longer wavelength should fall in the dip from the shorter wavelength, partly filling this up. Simulations suggest that the 488nm and 633nm lines have sufficient separation to result in full OTF support when used with the glycerol objectives, although the margin would be fairly small. With the higher aperture angle oil objectives due to be introduced soon, this technique should certainly be practical. There are various possible ways to combine the information from both wavelengths, but the most promising would probably be some form of ML deconvolution. It is also likely that one would need to cope with an arbitrary phase shift between the PSFs at the two wavelengths.

---

[1]Priv. Comm. T. Szellas, Leica Microsystems Inc.

15.2.2   4Pi FT-Spectrometry

Fourier transform spectrometers usually use a Michelson interferometer to measure the coherence properties of a light source. The spectrum can then be obtained by taking the Fourier transform of the interference signal obtained whilst varying the path difference between the two arms. One can change the path difference in the 4Pi interferometer using the phase piezo, and a long phase scan should give very similar information to that obtained from a traditional Michelson based FT spectrometer. The most important difference is that both excitation and emission light are brought to interference resulting in a signal which is the multiplication of the two coherence patterns. In the Fourier domain this would be reflected in a convolution of the two spectra. If using the visible CW lasers, the width of the excitation spectrum will be very small, and negligible in comparison to that of the emission. One would thus effectively only measure the (shifted) emission spectrum.

15.2.3   Phase Scan

4Pi images are usually obtained directly by scanning with a small axial step size. Another possibility would be to make standard x-y-z stacks with an axial step size as for normal confocal images[2]. When accompanied by a phase scan it would be possible to use a reconstruction algorithm similar to that used for other structured illumination techniques such as PEM and OMX[3]. This method has the disadvantage of being indirect (i.e. more data processing is required), and that adjustment is likely to be a little harder to perform alignment than when making x-z slices[4].

Some advantages of the method would be that acquisitions could be accelerated by avoiding the oversampling in x which is usually required to give a sensible z voxel size, and that it might be easier for users of normal confocal microscopes to learn. Most significant, however, is the simplification such an acquisition mode would offer for imaging in cases of refractive index mismatch. Refractive index mismatch in normal 4Pi images can be corrected for (see Chapter 8), but this currently requires an iterative deconvolution algorithm. If x-y-z-*phase* stacks were acquired it would be possible to correct for phase changes in a linear, one step process. This would make the image reconstruction much faster than would be the case with an iterative deconvolution, and considerably more flexible.

15.3   Why do we see different spots in confocal and SMI?

A relatively common complaint from biologists using the SMI is that, whereas they could clearly identify two 'specific' spots in confocal images, the SMI images either do not show a specific labelling or, equivalently, show a large number of equally bright spots which cannot be readily separated into specific and non-specific. As the SMI is significantly more sensitive and has a better linearity than confocal microscopy, it is difficult to conceive of a

---

[2]around 150nm would probably be appropriate

[3]as for 1D PEM, 3 different phases would be required, although more are likely to give a better result

[4]although phase adjustment shouldn't be necessary

reason why one might be able to see things under the confocal but not with the SMI. As high quality labelling performed within the group using conventional FISH as well several experiments with external collaborators have shown a good correspondence between SMI images and those obtained with the confocal, the tendency has been to write the complaints off as being due to a poor labelling and a rather optimistic interpretation of the confocal data. The few opportunities I have had to view the samples in question using the confocal would not seem to refute this supposition.

Nonetheless, whilst certainly a factor in some cases[5], it would be premature to assign all differences to subjectivity without further investigation.

15.3.1 Acceptor Saturation

If fluorophores are too closely packed they can quench each other. One possible cause of quenching is energy transfer between fluorophores in the form of FRET or homo-FRET. With typical FISH labelling it should be quite possible to achieve a labelling density sufficient for homo-FRET to be observed. A fluorophore can, however, only act as an acceptor if it is in the ground state. As the intensity is increased the proportion of molecules in an excited state at any given time increases. One would thus expect FRET based quenching to decrease. Expressing this mathematically, the relationship between excitation (I_{ex}) and emission (I_{em}) intensities would usually be given as:

$$I_{em} = \sigma Q I_{ex} \left(1 - \frac{I_{ex}}{I_{ex} + I_{sat}} \right) \tag{15.1}$$

where σ is the absorption cross-section, Q the quantum yield and I_{sat} the intensity required for 50% of fluorophores to be in the excited state. In the case of quenching due to a single energy transfer, this can be modified to give:

$$I_{em} = \sigma Q I_{ex} \left(1 - \frac{I_{ex}}{I_{ex} + I_{sat}} \right) (1 - E) \tag{15.2}$$

where E is the efficiency of the energy transfer. Using the Förster formula;

$$E = \frac{r_0^6}{r_0^6 + r^6} \tag{15.3}$$

where r_0 is the Förster radius and r the distance between the two fluorophores. In a population of fluorophores, the average distance r between two fluorophores is inversely proportional to the concentration. If we start to saturate our population, the average distance between two fluorophores capable of acting as acceptors will increase, giving the following relationship:

$$r(I_{ex}) = \frac{r}{\left(1 - \frac{I_{ex}}{I_{ex} + I_{sat}} \right)} = r \frac{I_{ex} + I_{sat}}{I_{sat}} \tag{15.4}$$

---

[5]In the *DMD*-locus measurement attempt, for example, significant differences in the interpretation of widefield images of the cells were obtained from the different people involved.
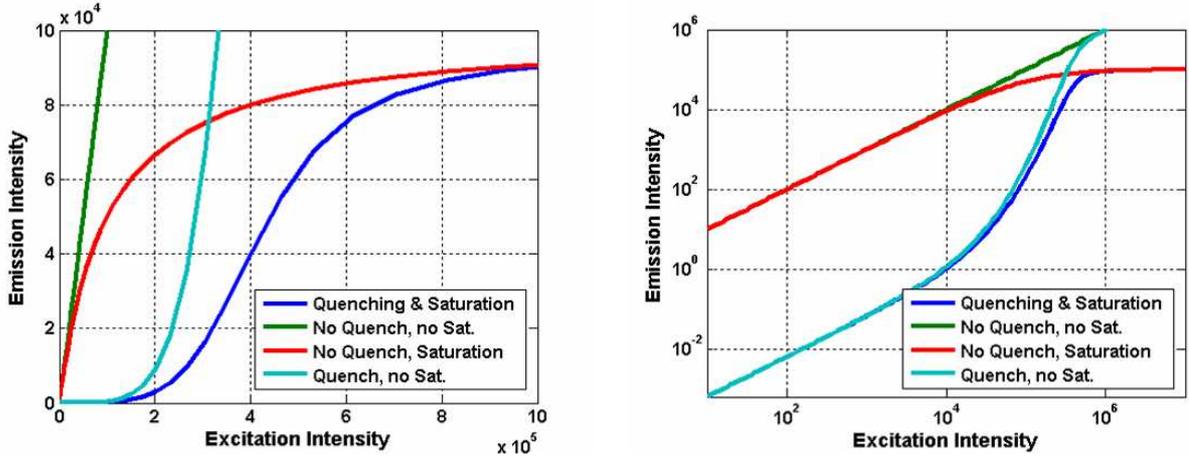
Figure 15.3: Left: Fluorescence response of a normal labelling, and one with saturable quenching. The responses without normal saturation are also shown in order to illustrate the source of the various parts of the curve. **Right:** The same with a dual logarithmic scale shoing that in the case of quenching the response remains approximately linear through several orders of magnitude, only showing a non-linearity as one approaches saturation. The intensity scale in both cases is arbitrary.

Substituting in equations 15.3 and 15.2 gives us:

$$I_{em} = \sigma Q I_{ex} \left(1 - \frac{I_{ex}}{I_{ex} + I_{sat}} \right) \left[1 - \frac{r_0^6}{r_0^6 + \left(r \frac{I_{ex} + I_{sat}}{I_{sat}} \right)^6} \right] \quad (15.5)$$

When plotted (with arbitrary values: $I_{sat} = 10^5$, $r_0 = 10$nm, $r = 2$nm and $\sigma Q = 1$) the curves in Fig. 15.3 are obtained. As the intensity in a confocal microscope is usually within an order of magnitude or two of that required for saturation, a nonlinearity of this nature has the potential to be relevant for confocal imaging. Equation 15.2 assumes that the quenching process is complete after one energy transfer. In the case of homo-FRET, however, fluorescence from the acceptor is also likely and transfers between several fluorophores may need to occur for effective quenching. It is not entirely clear exactly what form the equation would need to have to take this into account, but one would expect the process to become even less linear. This assumption of increasing non-linearity is based on the form obtained for two consecutive interactions:

$$I_{em} = \sigma Q I_{ex} \left(1 - \frac{I_{ex}}{I_{ex} + I_{sat}} \right) \left[1 - E_1 \left(1 - E_2 \right) \right] \quad (15.6)$$

Essentially the same non-linearity was proposed by Schönle *et Al* [85] as being exploitable for sub-resolution imaging. In this paper a standard FRET pair was assumed, and the possibility of homo-FRET or self quenching is not explicitly mentioned (although the principle remains the same). A more significant difference lies in the assumptions inherent
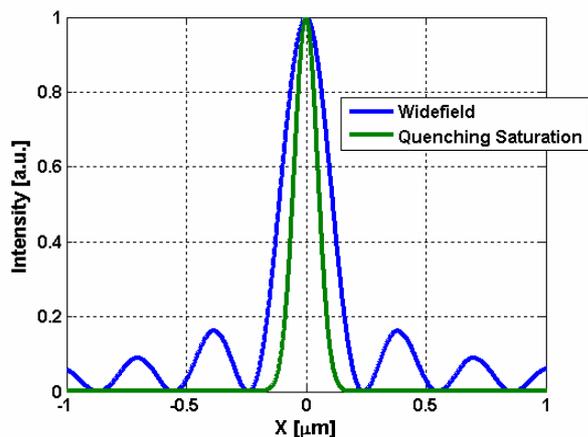
Figure 15.4: An example of what might be obtained when using the non-linearity as described above for resolution improvement.

in the mathematics. The considerations here are based on a saturation dependant effective distance between donor and acceptor, whereas Schönle assumes a chain consisting of a donor and one or more acceptors at a fixed distance. The most important difference in the predictions of the two models is that the chain model predicts non-linearity at all intensities, whilst with the concentration based maths significant non-linearity is only experienced as one nears saturation. It is unclear which is likely to be more realistic in practice - the chain model is certainly the best for fluorophore chains synthesised to reproducibly produce the effect. For relatively dense stochastic labelling of a structure, it is however possible that the population model would give better results.

In any case a non-linearity of at least quadratic, if not significantly higher, order should be achievable. This obviously lends itself to resolution improvement. Figure 15.4 shows the fluorescence intensity one might obtain in the focus, were Eqn. 15.2 to hold. As one needs to operate so near saturation it likely that the non-linearity could only be used with relatively photo-stable dyes.

Experimental Verification

Whilst the theoretical predictions are relatively clear, I am unaware of an experimental verification of the concept (and neither are the authors of the original paper[6]). If such non-linearities were to crop up when using normal labelling, they might well be detrimental to the correct interpretation of experimental data. This, as well as the promise of an additional technique for resolution improvement, should warrant further investigation of the phenomenon.

---

[6]Priv. Comm. A. Schönle

Chapter 16

Conclusion

This thesis investigates the use of advanced light microscopic techniques to measure structural parameters of fluorescently labelled objects with an accuracy which is significantly better than the optical resolution. Particular focus is placed on the techniques of Spatially Modulated Illumination and of 4Pi microscopy, although some attention was also paid to improving data analysis in confocal microscopy. Simulations were also performed to investigate some of the limitations inherent in SMI size estimations involving arbitrarily shaped objects.

The existing SMI microscope setup was modified to provide for excitation at 568nm. This, along with several simplifications of the setup, the introduction of a more sensitive black and white camera, and a rewrite of the control software have brought the SMI microscope a good deal closer to uncomplicated routine use. A refinement and automatisation of the data analysis procedures was also undertaken, and it is now possible to perform the data evaluation with a minimum of user intervention. The SMI analysis software now includes an automated object finding procedure, a fit algorithm capable of dealing with asymmetric axial profiles, and for which all starting parameters are automatically determined, and a lateral Gaussian fit for position and approximate size determination in the x-y plane.

The SMI microscope was applied to several biological problems during the course of the thesis, two of which are presented in detail this thesis. The first of these was the size analysis of replication foci (Chapter 11), in which it was shown that, despite a clustering of the foci into higher order structures, no significant change in the size of the underlying foci takes place as the cells pass through S-Phase. The unit size was estimated to be approximately 125nm. This corroborates results obtained using electron microscopy. In contrast to the EM study, our results are supported by a substantially larger sample size and include the comparison of various labelling strategies (protein/DNA and XFP/antibody). They additionally reveal a large variation in the size spectra between individual cells, which could either be the result of the high variability inherent in most biological systems, or indicate the presence of some 'hidden variable', potentially unrelated to S-phase position which has an effect on loci size.

The second biological problem (Chapter 13) involved the structure of centromere, or more specifically the distribution of a couple of inner kinetochore proteins. This represents

the first SMI measurements using a double labelling with GFP and mRFP. EM studies of the kinetochore during metaphase show a disk like structure with a diameter of approximately 350nm and a thickness of ∼30nm, the exact structure of the kinetochore in interphase remains however unclear. Due to the somewhat larger size of the kinetochore complex, completely unambiguous results could not be obtained. It was however possible to say that the structure is relatively compact, and that a completely flat disc is unlikely. A spherical form would also be unlikely. The results would be compatible with a somewhat crumpled disc, a configuration which is certainly biologically feasible. They are, however, many other structures which would give a similar result.

The reliance on an object model and the problems this poses for SMI size measurements on arbitrarily shaped objects have been discussed previously ([104]). In this thesis consideration of SMI measurements on non spherical objects has been extended further by analysing simulated gene structures (Chapter 12). These studies raise some serious questions about the validity of size estimation with non-spherical objects, particularly for sparse objects such as polymer random walks (as one would expect from chromatin). Nonetheless if the structures are small enough, and can be assumed to be relatively compact, a scaling relationship can be shown to exist. Generally one must be extremely careful when attempting to make quantitative size estimates for arbitrary objects. For structures which are larger, or about which relatively little is known, a better strategy is to simulate possible structures and compare the modulation depth obtained from these structures with those measured experimentally. Unfortunately one is currently limited to qualitative statements in this case.

The same algorithms written for SMI data evaluation were adapted to work with the 4Pi microscope. Size estimation did not work nearly as well with the 4Pi as with the SMI, a fact which may be due to the poor stability of the commercial 4Pi. In addition to the model based algorithms, a deconvolution procedure was developed for 4Pi data in which a refractive index mismatch induced phase shift in the PSF can be corrected (Chapter 8). This algorithm was successfully applied to several experimental data sets.

Motivated by the production of several gene-region labellings within the group which turned out to be too large for reliable size measurements using the SMI microscope, model based fit algorithms for use with confocal microscopy were investigated (Chapter 14). Both the Gaussian fit and the *multi-point model* allow measurements with an accuracy which is significantly better than the resolution and thus what can be achieved using threshold based methods. They also lack the subjectivity associated with setting a threshold. The multipoint model uses a somewhat more sophisticated object model, allowing measurements of objects which are not approximately Gaussian/spherical in shape, it is thus particularly suitable for the measurements of gene structures in the ∼200kbp range.

The algorithms and SMI microscope hardware have now reached a level of maturity where they can easily used by a wide user base including biologists. This is evidenced by the broad acceptance within the group and their use in in several other projects in addition to the examples presented in this thesis. Minor modifications also allow the data analysis tools to be used for the evaluation of PALM data. The combination of structured illumination approaches such as SMI and PALM is likely to improve the localisation accuracy and

hence the effective resolution. Before the introduction of structured illumination can be considered, however, several basic problems pertaining to 3D localisation such as blinking and bleaching must be addressed.

Due to the fact that most of the techniques presented attempt to estimate quantities with an accuracy better than the imaging resolution, some form of assumption about the object, usually in the form of a mathematical model, is always required. When using such approaches it is vitally important that the assumptions are reasonable. This requires a good mathematical description of the object and imaging process in addition to good labelling practices and a sensitive and properly calibrated microscope. A failure in any of these aspects is likely to prevent useful information from being obtained.

Part V

Appendices

Appendix A

pySMI – SMI Control Software

The incorporation of shutters required significant changes to the SMI control structure. While it would have been possible to make these changes in the original (c++, MSFC) program, a rewrite in Python had several advantages. Python[1] is a high level interpreted language which is fully object oriented and, due to the structure of the language, characterised by concise, easily readable code. Because Python is an interpreted language, there is a small performance hit over a compiled language such as c. This is largely offset by a faster development time, and the ability to easily integrate c/c++ code for performance critical sections.

A.1 Program Structure

A.1.1 Low Level c++ Layer

The program sections which need access the hardware, or are performance critical were implemented in c++ and bound into Python using SWIG[2]. The hardware control classes for the stepper motors and piezos were taken from the original SMI project with minor modifications; the camera class required extensive modifications, and the shutter control was written from scratch. In order to ensure adequate performance, image storage, the transfer of images from the camera, and the translation of image data into a picture for display are also implemented in c++.

A brief description of each class below in the c++ layer is given below:

CCamera Hardware interface to camera - includes code to transfer images into data stack.

CStepOp Controls stepper motors.

CPiezoOp Controls piezos.

CShutterControl Controls shutters by toggling pins on parallel port.

---

[1]http://www.python.org

[2]Simple Wrapper Interface Generator [9] - a program which automatically generates the wrapper code required to call c/c++ code from several high level languages including Python

CDataStack In memory implementation of a data stack, designed to be easily loaded from and written to the .kdf format used for SMI data storage. Each separate channel is stored in a contiguous block of memory.

LineProfile Code for extracting a profile from a data stack.

CBaseRenderer Virtual base class (cf. Java `interface`) defining what functions a renderer (code for generating the displayed picture) should offer. Basis of an extensible system of renderers - eg pseudocolour. A renderer draws on a bitmap provided by the calling function; in the Python implementation a c pointer to a bitmap created by the windowing library (wxWindows)[3].

CLUTRGBRenderer The only implemented renderer - three channels as RGB using lookup tables.

CDisplayOpts Options for renderer, stating gains, offsets etc for the different colour channels.

A.1.2 Python Implementation layer

Sitting on top of the low level c++ layer are several Python classes providing the 'glue' between the various hardware components and implementing the image acquisition process. The implementation of the procedures responsible for defining the sequence of piezo movements, camera exposures and shutter states in Python allows the acquisition process to be easily customised and extended. It is however important to maintain a reasonable level of performance in the acquisition code. To this end one important trick is used; whilst the transfer of data from the camera into the in-memory storage is initiated from the Python code, the process itself takes place in the complied c++ code. To facilitate this, the `CCamera` object knows about the `CDataStack` object, and how it should write data into it.

The two most important classes in the Python implementation layer are `PreviewAcquisator` and `SimpleSequenceAcquisator`. `PreviewAcquisator` provides the functionality required to obtain a sequence of pictures from the camera whilst operating the shutters and is used for the live display. Stubs for piezo movements between images are also included but do nothing. `SimpleSequenceAcquisator` extends `PreviewAcquisator` adding code to the piezo movement stubs and is used for the acquisition of image stacks. Both classes obtain camera and shutter parameters from the `chaninfo` class.

A.1.3 wxPython GUI

The pySMI GUI is written entirely in Python, using the wxPython windowing library. It reproduces the functionality of the previous c++ based software with the following significant differences; it supports an arbitrary number of colour channels (necessitated

---

[3]Note this a somewhat obscure workaround to avoid drawing to anything in Python. I have explicitly mentioned it here because it is important for performance and could be a cause for confusion to anyone wanting to work with the code.

by the shutters - the previous version supported either 1 or 4 corresponding to black and white or colour cameras), it can show xz and yz slices through data stacks, and it lets you have more than one stack open at any one time. Also new is the console window which is discussed below. A detailed description of the implementation is unrealistic here, hopefully the few screen shots shown in Figs. A.1, A.2 and A.3 will suffice to give a reasonable idea.
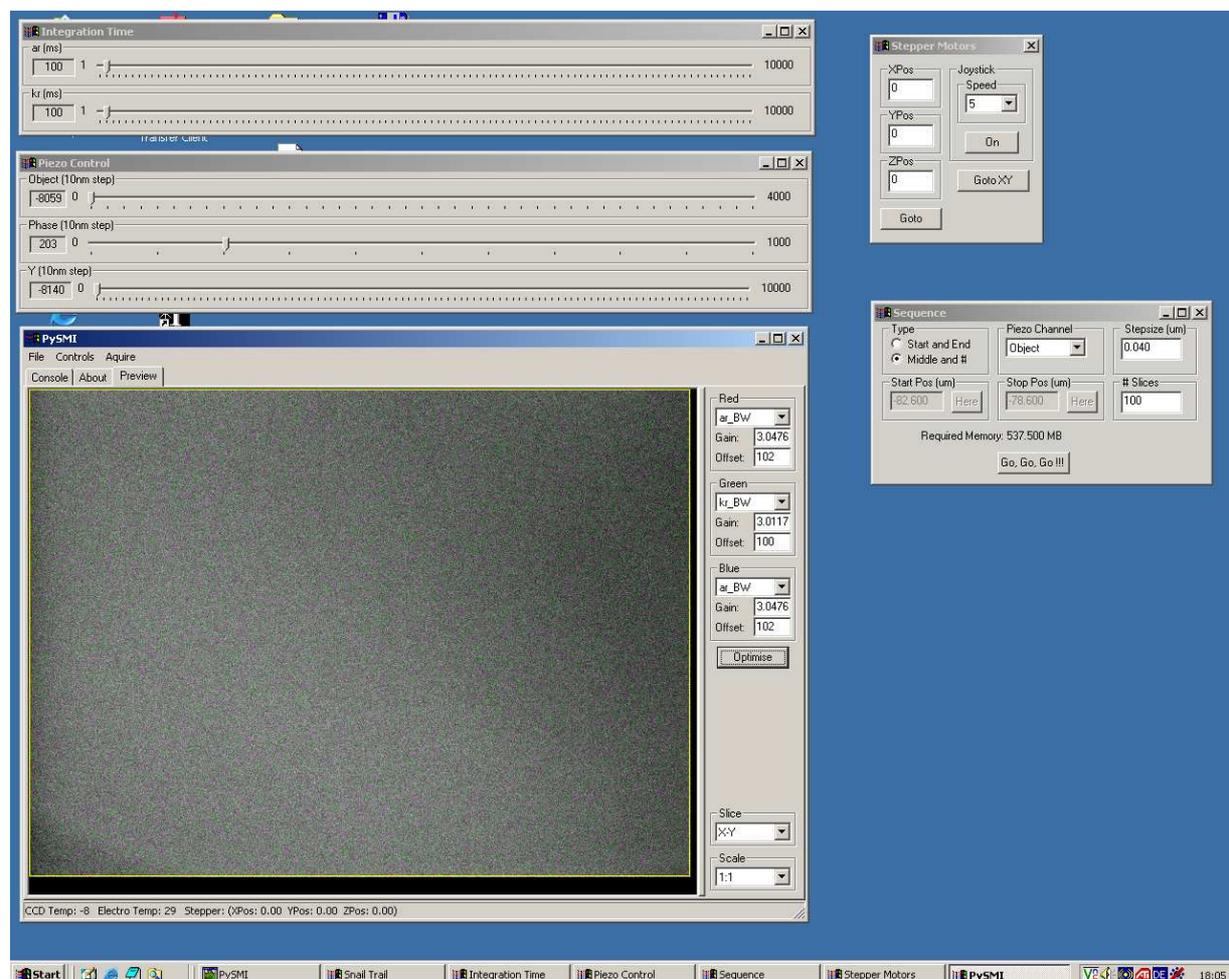
Figure A.1: pySMI control software showing main window, piezo, integration time and stepper motor controls, and the sequence acquisition dialog.

A.2 Expanding pySMI

One of the original goals motivating the move to Python was to make the maintenance and adaptation of the SMI software easier. Due to the fact that Python is a high level language, and due to the modular design of the software it is very easy to write small scripts to implement a certain functionality. Indeed, due to Python's weak typing, using an alternative piece of hardware is as easy as substituting another class which provides the

Figure A.2: Channel control: configured shutter controlled channels are shown in the channel dialog (**top**). Editing or adding a channel brings up the channel edit dialog (**bottom**) where the pins on the parallel port to be toggled and the colour elements of the CCD to be read out can be selected. Note that the *Colours* options are for a colour CCD with Bayer filter - for a B/W CCD *B/W* is the only sensible choice. B/W can also be used with a colour CCD in which case all channels are averaged.
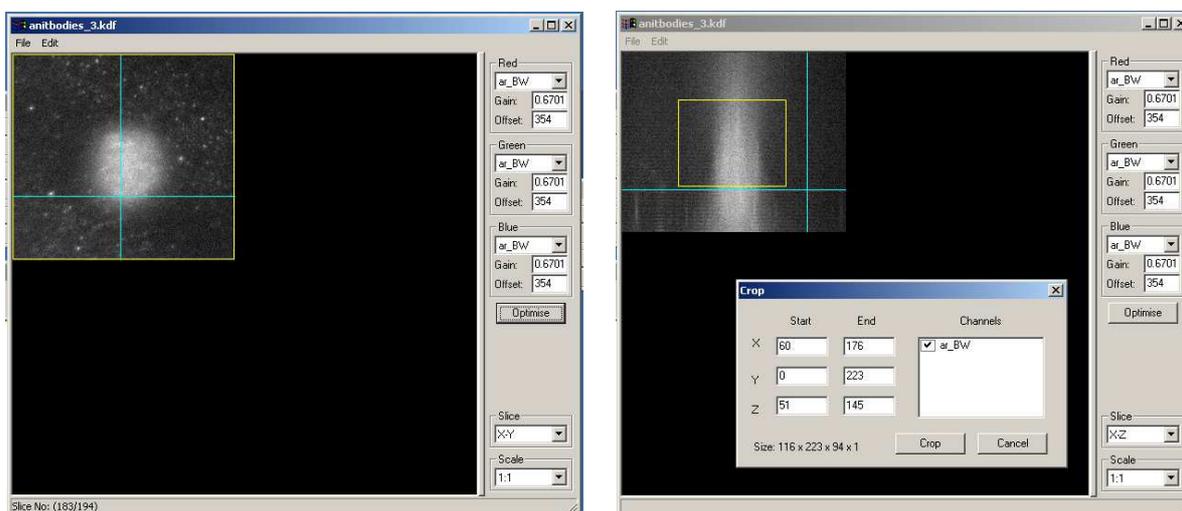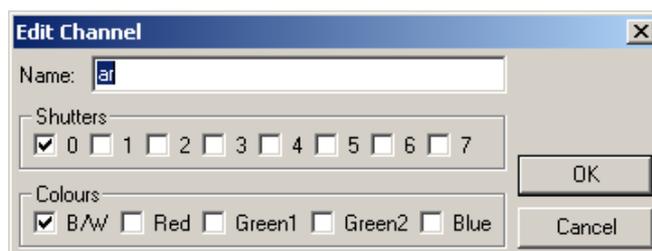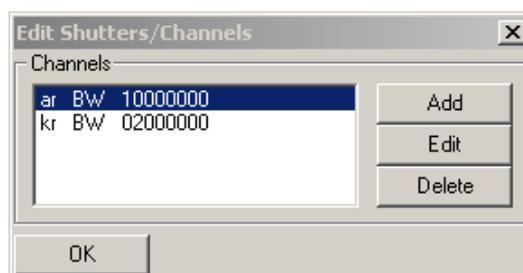
Figure A.3: pySMI includes a simple viewer for 3D image stacks, supporting slicing in xy, xz, and yz planes, simple colour map manipulations and the cropping of images after acquisition. The viewer portion of pySMI, being hardware independent, also runs on Linux and should be able to be easily ported to any platform for which gcc and wxPython are available.

same interface as the original one. This can even be done while the program is running! An indication of the success of this approach is that pySMI is currently running on five different microscopes within the group. These microscopes have vastly different hardware configurations with the only common factor being the use of a PCO camera on each.

A.2.1   init.py

Rather than reading static parameters from a classical configuration file (eg .ini files, windows registry etc...) pySMI obtains its configuration from the `init.py` file which is a valid

Python script and is executed during program startup. Within this script it is possible to perform significant customisation; for example the class to use for the piezo controller is initialised here, and the 'Snail Trail' display currently offered on SMI1 is implemented entirely within `init.py`.

A.2.2   The pySMI Console

One of the most useful features of pySMI when developing or testing new code is the console. This exposes the Python interpreter along with the global namespace of the program to the user. It is thus possible to interactively execute Python code and to interact at a low level with the various program components. This allows existing components to be debugged and new components tested without the need to incorporate them into the gui. It is also possible to alter the running program.

A.2.3   Different Hardware

PIFoc

Control for PIFocs and other PI piezo devices using either the E662 or E816 control interfaces has been implemented by redefining the piezo class. The code for the E662 controller is shown below as an example of how simple it is to implement such an extension:

```python
import serial;

class piezo_e662:
    def __init__(self, portname='COM2', maxtravel = 100.00):
        self.max_travel = maxtravel
        self.ser_port = serial.Serial(portname, 9600, rtscts=1, timeout=30)
        self.ser_port.write('DEV:CONT REM\n')

    def MoveTo(self, iChannel, fPos, bTimeOut=True):
        if (fPos >= 0):
            if (fPos <= self.max_travel):
                self.ser_port.write('POS %3.4f\n' % fPos)
            else:
                self.ser_port.write('POS %3.4f\n' % self.max_travel)
        else:
            self.ser_port.write('POS %3.4f\n' % 0.0)

    def GetPos(self, iChannel=1):
        self.ser_port.write('POS?\n')
        res = self.ser_port.readline()
        return float(res)
```

```
def GetControlReady(self):
    return True

def GetChannelObject(self):
    return 1

def GetChannelPhase(self):
    return 1
```

As a comparison, the c++ code used to implement the piezo control on SMI1 using an earlier PI piezo controller (having a different but relatively similar command set) weighs in at 1082 lines[4].

Tomograph

Likewise the changes required to incorporate tomograph rotation, including menu items, keyboard shortcuts, and automatic correction of z-position to ensure that the object remains in focus after the rotation required only a few lines of code.

A.3 Room for Improvement

There are several features which, when implemented, could significantly improve the control software. Some of these are as follows:

Multiple piezo acquisitions Initial code for this exists - storing the second motion direction in the channel direction. A GUI is has however not been implemented. Multiple piezos acquisitions are important in conjunction with lateral structured illumination (eg: slit scan) or when trying to minimise acquisition time for a 3D stack (3 phase steps at a z step of 200-250nm would still satisfy Nyquist, whilst requiring a somewhat smaller total number of images).

Support for .tiff images Whilst .kdf is a very easy format to deal with, it is not nearly as universal as .tiff[5]

Tracking of slide movements Test code for recording the movements of the slide and hence providing a record of the areas which have been examined is implemented on SMI1. This is however lacks many features eg: erasing the track after changing slides.

---

[4]This comparison is not entirely fair - the original c++ code could have been made significantly smaller if a pre-existing serial class library had been used. The example code also fails to implement multiple channels (the controllers we have are principally capable of operating in multi-channel mode through the use of daisy chaining, although this is not currently used). A factor of 5-10 in code length and a significant increase in readability are however quite normal.

[5]Although seeing as .tiff is such a broad specification, saving as .tiff does not guarantee that anyone else is going to be able to read it. I am not currently aware of any (widely available) program which supports the 4D, 16bit tiffs that we would need.

Tighter integration of 'cell finder' It would be a relatively trivial task to read the data from the 'cell finder' coordinate transform software, and, after manually finding the reference points, to automatically relocate the cells. This could be integrated with the above item to provide a graphical picture of the slide with both the locations of the cells and the path taken shown.

Interfacing the data as a NumArray type Data storage is currently implemented in a c++ class to which there is only indirect (slow) access from the Python code. When a little wrapper code was written such that the NumArray interface could be used, then it would be very easy to implement a significant amount of data processing within pySMI through the use of SciPy[6].

Support for other camera types pySMI currently only has support for PCO sensicam cameras. The addition of other camera interfaces would obviously be desirable.

Better handling of image metadata There are several aspects here which could be improved including a couple which fall into the category of bugs:

- BUG: Data in the .log file is only correct if the image is saved before any other images are acquired. Otherwise the .log file refers to the most recent stack.

- Metadata should include voxelsize, as this varies between microscopes and will change if objectives / tube lens, binning etc. are changed. Currently the evaluation software has the values for SMI1 hard coded.

- The metadata could also be directly entered into the database used to store the settings for the analysis software - this could allow individual images to be traced through all stages of analysis. Whilst highly desirable in a clinical context, it is probably overkill in our case.

---

[6]SciPy (www.scipy.org) provides a fast, vectorised, implementation of a large number of basic mathematical operations, not dissimilar to Matlab. It is relatively easy to translate existing Matlab code into Python + SciPy with the advantages of increased stability and a clearer coding style.

Appendix B

Biological Methods

B.1 KIR2.1 Pottassium Channels

B.1.1 Cell biology

N-terminal fusion constructs of Kir2.1 channel subunits with enhanced Yellow Fluorescent Protein (eYFP) were designed by inserting the respective cDNA in-frame into the pEYFP-C1 eukaryotic expression vector (Clontech) using the EcoRI and BamHI restriction sites as described in Stockklausner et al [95]. Human embryonic kidney (HEK) 293 cells were grown in MEM (Minimum Essential Medium, Invitrogen), supplemented with 10% fetal calf serum and 1% penicillin/streptomycin at 37 °C and 5% CO_2. The cells were grown on coverslides provided by Leica Microsystems. At 80% confluence, the cells were transfected with the respective cDNA using Fugene 6 (Roche) following the supplier's instructions.

B.1.2 Sample preparation

The two coverslips specifically designed for use with the Leica 4Pi Microscope were provided by Leica Microsystems (Mannheim, Germany). Transfected cells were fixed in 4% para-formaldehyde (PFA) in phosphate-buffered saline (PBS) for 20 min at 4 °C on one of the coverslips. The fixed cells were embedded in 87%(volume) Glycerol, 50 mM Tris-HCL, pH= 7.4. On the other coverslip we immobilised green fluorescence beads (diameter $100nm$, G100, Duke Scientific) to implement reference objects for evaluation of the surface structure. For sealing the coverslips we used nail-polish and dental glue (Twin-sil A+B).

B.1.3 4Pi Microscopy

The measurements presented here were performed on a commercial prototype 4Pi Microscope at Leica Microsystems Inc, Mannheim, Germany. This is a 2-Photon 4Pi microscope of type A (the excitation foci are brought to interference, but interference of the detected light is not used). 2-Photon Excitation was provided using a Ti-Saphire laser producing ps pulses and tunable between 740 and 990nm. A matched pair of glycerol immersion objectives (100x, NA=1.35) were used with quartz coverslips and 87% glycerol as a mounting and

immersion medium. A voxel size of 26.2nm in x and z and 30.5nm in y was chosen and an xzy scan performed resulting in a final image size of 512x25x512 voxels. The pinhole and beam expander were set to 1 Airy and 3 respectively, a line average of 16 and accumulation of 2 were used. The excitation wavelength was approximately 970nm.

B.2   Replication Foci

B.2.1   Cell culture

C2C12 mouse myoblast cells (Yaffe and Saxel, 1977) and C2C12 cells stably expressing GFP-tagged PCNA (Leonhardt et al., 2000) were cultured using DMEM supplemented with 20% fetal calf serum at 37 °C in a humidified atmosphere of 5% CO2.

B.2.2   In situ replication labelling

DNA synthesis foci were labelled by adding the thymidine analogue BrdU to the cell culture medium at a final concentration of 10-4 M for 10 minutes.

B.2.3   Hypotonic treatments

Nucleotide-labelled C2C12 cells were trypsinised, the cell suspension was washed twice with cold PBS and resuspended in a small volume (50 μl) of PBS. Excess of prewarmed 50mM KCl solution was added to the suspension and the cells were let to swell at 37 °C for another 10 min. The cells were pelleted again and resuspended in 30 l of the hypotonic solution. Freshly made cold methanol:acetic acid mixture (3:1) was added [drop by drop] and the cells were kept for at least 30 min at - 20 °C . The resulting nuclei were pelleted and excess of the fixative was removed. Nuclei were gently pipetted in the remaining volume, transferred to the microscope slides and air-dried. The slides were kept at - 20 °C until immunostaining.

B.2.4   Cryosections

Cells were fixed in 4% freshly-depolymerised paraformaldehyde in 250 mM HEPES (pH 7.6, 10 min, 4 °C), refixed in 8% paraformaldehyde in 250 mM HEPES (pH 7.6, 2 h, 4 °C), scrapped and pelleted during the latter fixation (Branco et al., 2005). Cell pellets were embedded (2 h) in 2.1 M sucrose in PBS, transferred onto a copper block, frozen and stored in liquid nitrogen. Cryosections (100-150 nm thick) were cut with a glass knife using an ultracryomicrotome (UltraCut UCT52, Leica, UK), captured on drops of 2.1 M sucrose in PBS, transferred to 0.17 mm thick glass coverslips or Cellocate coverslips (Eppendorf), and stored at -20 °C .

B.2.5 Anti-BrdU immunostaining

All incubations were at room temperature unless otherwise stated. Briefly, formaldehyde fixed cells were permeabilised by 10 min incubation in PBS containing 0.5% Triton X-100, followed by blocking for 30 min in PBS containing 0.2% fish skin gelatin. Hypotonically treated methanol:acetic acid fixed cells were not further permeabilised and directly blocked in PBS with 0.2% gelatin. Incorporated BrdU was detected by incubating for one hour with anti-BrdU mouse monoclonal antibody (6 μg/ml; clone B44, Becton Dickinson) at 37 °C in a solution (33 mM Tris-Cl pH 8.1, 0.33 mM $MgCl_2$, 0.5 mM beta-mercaptoethanol, 1% BSA) containing DNase I (20 U/ml) to expose the epitope. Detection was performed with an anti-mouse IgG antibody conjugated with FITC (Jackson Immuno Research). Samples were mounted in Mowiol or VectaShield for confocal and conventional microscopy and in VectaShield (Vector Laboratories) for high-resolution microscopy.

B.2.6 Anti-GFP immunostaining

Cryosections were rinsed (3x) in PBS, incubated (15 min) in 20 mM glycine in PBS, rinsed (3x) in PBS, treated (10 min) with 0.1% Triton X-100 in PBS, blocked (30 min) with PBS+ (PBS supplemented with 1% BSA, 0.2% fish skin gelatin, 0.1% casein; pH 7.6), incubated (2 h) with rabbit antibodies anti-GFP (diluted 1:50 in PBS+; affinity-purified rabbit polyclonal antibodies raised against full length GFPS65T), washed (5x, over 1 h) in PBS+, incubated (1 h) with AlexaFluor488-conjugated antibodies against rabbit IgGs (in PBS+; raised in donkey; Molecular Probes, Eugene, OR), rinsed (3x, 30 min) in PBS+, rinsed (3x) in PBS, counterstained (45 min) with TOTO-3 (2 μM, Molecular Probes) in 0.05% Tween-20 in PBS, rinsed successively in 0.05% Tween-20 in PBS and then PBS, before coverslips were mounted in VectaShield (Vector Laboratories, Peterborough, UK) on glass slides coated with 0.1 μm FluoresbriteTM Carboxylate YG microspheres (Polysciences, Inc., Warrington, PA). The long incubation times used allow for antibody accessibility throughout the thickness of cryosections, providing the highest sensitivity (Branco et al., 2005).

Acknowledgements

Bibliography

[1] E. Abbe, *Beiträge zur theorie des mikroskops und der mikroskopischen wahrnehmung*, Arch. f. mikr. Anat. **8** (1873), no. 9, 413–468.

[2] B. Alberts, A. Johnson, J. Lewis, R. Martin, K. Roberts, and P. Walter, *Molecular biology of the cell*, 4 ed., Garland Science, 2002.

[3] B. Albrecht, *Hochpräzisions-fluoreszenzmessungen mit räumlich modulierter anregung*, Ph.D. thesis, Universität Heidelberg, 2002.

[4] D. Baddeley, C. Carl, and C. Cremer, *4Pi microscopy deconvolution with a variable point-spread function*, Applied Optics **45** (2006), no. 27, 7056–7064.

[5] B. Bailey, D.L. Farkas, D.L. Taylor, and F. Lanni, *Enhancement of axial resolution in fluorescence microscopy by standing-wave excitation*, Nature **366** (1993), no. 6450, 44–48.

[6] R. Barrett, M. Berry, T. F. Chan, J. Demmel, J. Donato, J. Dongarra, V. Eijkhout, R. Pozo, C. Romine, and H. Van der Vorst, *Templates for the solution of linear systems: Building blocks for iterative methods, 2nd edition*, SIAM, Philadelphia, PA, 1994, `http://www.netlib.org/linalg/html_templates/Templates.html`.

[7] J. Bastert, *Präzisionsgrräßen- und Distanzmessungen an multispektral fluoreszierenden Probekörpern mit der Spatially Modulated Illumination Microscopy*, Master's thesis, Universität Heidelberg, 2006.

[8] C. Batram, *Hochpräzisionsgrößenmessung und Bestimmung der Topologie subchromosomaler Domänen nach molekularer Markierung unter Berücksichtigung des Einflusses der Zellfixierung*, Ph.D. thesis, Universität Heidelberg, 2007.

[9] D.M. Beazley, *SWIG: An easy to use tool for integrating scripting languages with C and C++*, Proceedings of the 4th USENIX Tcl/Tk Workshop (1996), 129–139.

[10] AS Belmont and AF Straight, *In vivo visualization of chromosomes using lac operator-repressor binding.*, Trends Cell Biol **8** (1998), no. 3, 121–4.

[11] E. Betzig, G.H. Patterson, R. Sougrat, O.W. Lindwasser, S. Olenych, J.S. Bonifacino, M.W. Davidson, J. Lippincott-Schwartz, and H.F. Hess, *Imaging Intracellular Fluorescent Proteins at Nanometer Resolution*, Science **313** (2006), no. 5793, 1642.

[12] J. Bewersdorf, R. Schmidt, and SW Hell, *Comparison of I 5 M and 4Pi-microscopy*, J Microsc **222** (2006), 105–17.

[13] C. Mar Blanca, J. Bewersdorf, and S. W. Hell, *Determination of the unknown phase difference in 4pi-confocal microscopy through the image intensity*, Optics Communications **206** (2002), 281–285.

[14] A. Bloess, Y. Durand, M. Matsushita, H. van Dermeer, GJ Brakenhoff, and J. Schmidt, *Optical far-field microscopy of single molecules with 3. 4 nm lateral resolution*, Journal of Microscopy **205** (2002), no. 1, 76–85.

[15] M. Born and E. Wolf, *Principles in optics*, 6. ed. ed., A. Wheaton, Co., Oxford, 1983.

[16] H. Bornfleth, *Precision distance microscopy and model-based segmentation for the analysis of the functional topology of the cell nucleus*, Ph.D. thesis, Universität Heidelberg, 1998.

[17] M.R. Branco and A. Pombo, *Intermingling of chromosome territories in interphase suggests role in translocations and transcription-dependent associations*, PLoS Biol **4** (2006), no. 5, e138.

[18] N.P. Carter, *Fluorescence in situ hybridizationstate of the art*, Bioimaging **4** (1996), no. 2, 41–51.

[19] E. Chung, D. Kim, Y. Cui, YH Kim, and PT So, *2D standing wave total internal reflection fluorescence microscopy: super-resolution imaging of single molecular and biological specimens.*, Biophys J (2007).

[20] BG Clark, *An efficient implementation of the algorithm 'CLEAN'*, Astronomy and Astrophysics **89** (1980), no. 3, 377–378.

[21] J.A. Cooper, *Effects of Cytochalasin and Phalloidin on Actin*, The Journal of Cell Biology **105** (1987), no. 4, 1473–1478.

[22] C. Cremer, P. Edelmann, H. Bornfleth, H. Luz, G. Kreth, H. Münch, and M. Hausmann, *Principles of spectral precision distance confocal microscopy of molecular nuclear structure*, Handbook of Computer Vision and Applications (B. Jähne, H. Haußecker, and P. Geißler, eds.), vol. 3, Academic Press, San Diego, New York, 1999, pp. 839–857.

[23] T. Cremer and C. Cremer, *Chromosome territories, nuclear architecture and gene regulation in mammalian cells*, Nat. Rev. Genet **2** (2001), no. 4, 292–301.

[24] G. Dellaire, H. Dehghani, R. Ching, C.H. Eskiw, and D.P. Bazett-Jones, *Interactions between chromatin and pml nuclear bodies: insights from fluorescence microscopy and electron spectroscopic imaging*, International Symposium: Optical Analysis of Biomolecular Machines, Berlin, July 2006.

[25] W. Denk, J. H. Strickler, and W. W. Webb, *2-photon laser scanning fluorescence microscopy*, Science **248** (1990), 73–76.

[26] M. Dyba and S.W. Hell, *Focal Spots of Size $\lambda/23$ Open Up Far-Field Florescence Microscopy at 33 nm Axial Resolution*, Physical Review Letters **88** (2002), no. 16, 163901.

[27] P. Edelmann and C. Cremer, *Improvement of confocal spectral precision distance microscopy (spdm)*, Optical Diagnostics of Living Cells III, Proc. SPIE, vol. 3921, 2000, pp. 313–320.

[28] A. Egner, M. Schrader, and S. W. Hell, *Refractive index mismatch induced intensity and phase variations in fluorescence confocal, multiphoton and 4pi-microscopy*, Optics Communications **153** (1998), 211–217.

[29] A. Esa, P. Edelmann, G. Kreth, L. Trakhtenbrot, N. Amariglio, G. Rechavi, M. Hausmann, and C. Cremer, *Three-dimensional spectral precision distance microscopy of chromatin nanostructures after triple-colour dna labelling: a study of the bcr region on chromosome 22 and the philadelphia chromosome*, J. Microsc. **199** (2000), 96–105.

[30] J. Essers, A.F. Theil, C. Baldeyron, W.A. van Cappellen, A.B. Houtsmuller, R. Kanaar, and W. Vermeulen, *Nuclear Dynamics of PCNA in DNA Replication and Repair*, Molecular and Cellular Biology **25** (2005), no. 21, 9350–9359.

[31] Antonio Virgilio Failla, Benno Albrecht, Udo Spöri, Andreas Schweitzer, Alexander Kroll, Georg Hildenbrand, Margund Bach, and Christoph Cremer, *Nanostructure analysis using spatially modulated illumination microscopy*, Com-PlexUs **1** (2003), 77–88.

[32] S. Fenz, *Zwei-Farbenhochpräzisonsdistanzmesungen mit dem SMI-Mikroskop*, Master's thesis, Universität Heidelberg, 2005.

[33] R. Freimann, S. Pentz, and H. Horler, *Development of a standing-wave fluorescence microscope with high nodal plane flatness*, Journal of Microscopy **187** (1997), no. 3, 193–200.

[34] JT Frohn, HF Knapp, and A. Stemmer, *Three-dimensional resolution enhancement in fluorescence microscopy by harmonic excitation*, Opt. Lett **26** (2001), 828–830.

[35] C. Gromann, *Vergleichende fuoreszenzoptische Untersuchungen zur Genom- Mikround Nanoarchitektur der BCR- und ABL- Region an Kontroll-, Remissions und CML-Präparaten*, Master's thesis, Universität Heidelberg, 2005.

[36] M. Gu, *Principles of three-dimensional imaging in confocal microscopes*, 1. aufl ed., World Scientific Publishing, Singapore, 1996.

[37] M. G. L. Gustafsson, *Surpassing the lateral resolution limit by a factor of two using structured illumination microscopy*, J. Microsc. **198** (2000), no. 2, 82–87.

[38] M.G.L. Gustafsson, *Extended resolution fluorescence microscopy*, Curr. Opin. Struct. Biol **9** (1999), no. 5, 627.

[39] ———, *Nonlinear structured-illumination microscopy: Wide-field fluorescence imaging with theoretically unlimited resolution*, Proceedings of the National Academy of Sciences **102** (2005), no. 37, 13081–13086.

[40] M.G.L Gustafsson, D.A. Agard, and J.W. Sedat, *I5M: 3D widefield light microscopy with better than 100nm axial resolution*, Journal of Microscopy **195** (1999), no. Pt 1, 10–16.

[41] R.P. Haugland, M.T.Z. Spence, I.D. Johnson, and A. Basey, *The Handbook: A Guide to Fluorescent Probes and Labeling Technologies*, Molecular Probes, 2005.

[42] Michael Hausmann, Ralph Winkler, Georg Hildenbrand, Jutta Finsterle, Andrea Weisel, Alexander Rapp, Eberhard Schmitt, Siegfried Janz, and Christoph Cremer, *COMBO-FISH: specific labeling of nondenatured chromatin targets by computer-selected DNA oligonucleotide probe combinations.*, Biotechniques **35** (2003), no. 3, 564–70, 572–7.

[43] E. Hecht and A. Zajac, *Optics*, Addison-Wesley, 1974.

[44] R. Heintzmann, *Resolution enhancement of biological light microscopic data*, Ph.D. thesis, Universität Heidelberg, 1999.

[45] R. Heintzmann and P.A. Benedetti, *High-resolution image reconstruction in fluorescence microscopy with patterned excitation*, Applied Optics **45** (2006), no. 20, 5037–5045.

[46] R. Heintzmann and C. Cremer, *Laterally modulated excitation microscopy: improvement of resolution by using a diffraction grating*, Proc. SPIE **3568** (1998), 185–196.

[47] R. Heintzmann, T.M. Jovin, and C. Cremer, *Saturated patterned excitation microscopy-a concept for optical resolution improvement*, Journal of the Optical Society of America A **19** (2002), no. 8, 1599–1609.

[48] S. Hell, G. Reiner, C. Cremer, and E.H.K. Stelzer, *Aberrations in confocal fluorescence microscopy induced by mismatches in refractive index*, Journal of Microscopy **169** (1993), 391–405.

[49] S. W. Hell, C. Mar Blanca, and J. Bewersdorf, *Phase determination in interference-based superresolving microscopes through critical frequency analysis*, Optics Letters **27** (2002), no. 11, 888–890.

[50] S. W. Hell, S. Lindek, C. Cremer, and E. H. K. Stelzer, *Measurement of 4pi-confocal point spread function proves 75 nm axial resolution*, Appl. Phys. Lett. **64** (1994), no. 11, 1335–1337.

[51] S. W. Hell and E. H. K. Stelzer, *Fundamental improvement of resolution with a 4pi-confocal fluorescence microscope using two-photon excitation*, Opt. Commun. **93** (1992), 277–282.

[52] G. Henneke, S. Koundrioukoff, and U. Huebscher, *Multiple roles for kinases in DNA replication*, EMBO Reports **4** (2003), no. 3, 252–256.

[53] S.T. Hess, T.P.K. Girirajan, and M.D. Mason, *Ultra-High Resolution Imaging by Fluorescence Photoactivation Localization Microscopy*, Biophysical Journal **91** (2006), no. 11, 4258.

[54] M. Hofmann, C. Eggeling, S. Jakobs, and S.W. Hell, *Breaking the diffraction barrier in fluorescence microscopy at low light intensities by using reversibly photoswitchable proteins*, Proceedings of the National Academy of Sciences **102** (2005), no. 49, 17565–17569.

[55] JA Högbom, *Aperture Synthesis with a Non-Regular Distribution of Interferometer Baselines*, Astronomy and Astrophysics **15** (1974), 417.

[56] B. Jähne, *Digital Image Processing*, Springer, 2002.

[57] C. Janeway, P. Travers, M. Walport, and M. Shlomchik, *Immunobiology: The immune system in health and disease.*, 5 ed., ch. 3, Garland, 2001, http://www.ncbi.nlm.nih.gov/books/bv.fcgi?rid=imm.section.323.

[58] AW Jones, J. Bland-Hawthorn, and PL Shopbell, *Towards a General Definition for Spectroscopic Resolution*, Astronomical Data Analysis Software and Systems IV **77** (1995).

[59] T.A. Klar, S. Jakobs, M. Dyba, A. Egner, and S.W. Hell, *Fluorescence microscopy with diffraction resolution barrier broken by stimulated emission*, Proceedings of the National Academy of Science **97** (2000), no. 15, 8206–8210.

[60] M.V. Klein and T.E. Furtak, *Optik*, Springer, 1988.

[61] Karel Koberna, Anna Ligasova, Jan Malinsky, Artem Pliss, Alan J Siegel, Zuzana Cvackova, Helena Fidlerova, Martin Masata, Marketa Fialova, Ivan Raska, and Ronald Berezney, *Electron microscopy of DNA replication in 3-D: evidence for similar-sized replication foci throughout S-phase*, J Cell Biochem **94** (2005), no. 1, 126–138.

[62] O. Kratky and G. Porod, *Rontgenuntersuchung geloster fadenmolekule*, Rec. Trav. Chim **68** (1949), 1105.

[63] V. Krishnamurthi, B. Bailey, and F. Lanni, *Image processing in 3d standing-wave fluorescence microscopy*, Proc. SPIE **2655** (1996), 18–25.

[64] C. Lanctôt, T. Cheutin, M. Cremer, G. Cavalli, T. Cremer, et al., *Dynamic genome architecture in the nuclear space: regulation of gene expression in three dimensions*, Nature Reviews Genetics **8** (2007), 104–115.

[65] F. Lanni, A. S. Waggoner, and D. L. Taylor, *Standing-wave luminescence microscopy*, US-Patent No. 4 621 911, 1986.

[66] K. Levenberg, *A method for the solution of certain problems in least squares*, Quart. Appl. Math **2** (1944), no. 2, 164–168.

[67] K. Lidke, B. Rieger, T. Jovin, and R. Heintzmann, *Superresolution by localization of quantum dots using blinking statistics*, Optics Express **13** (2005), no. 18, 7052–7062.

[68] D. Marenduzzo, K. Finan, and P.R. Cook, *The depletion attraction: an underappreciated force driving cellular organization*, The Journal of Cell Biology **175** (2006), no. 5, 681.

[69] D. Marenduzzo, C. Micheletti, and P.R. Cook, *Entropy-Driven Genome Organization*, Biophysical Journal **90** (2006), no. 10, 3712–3721.

[70] D.W. Marquardt, *An Algorithm for Least-Squares Estimation of Nonlinear Parameters*, Journal of the Society for Industrial and Applied Mathematics **11** (1963), no. 2, 431–441.

[71] H. Mathée, *Hochpräzisionsmikroskopie mit räumlich strukturierter Beleuchtung und Zweiphotonenanregung*, Ph.D. thesis, Universität Heidelberg, 2006.

[72] H. Mathée, D. Baddeley, and Cremer. C., *Spatially modulated illumination microscopy using one objective lens*, Optical Engineering(in press) (2007).

[73] T. Misteli, *Concepts in nuclear architecture*, BioEssays **27** (2005), no. 5, 477–487.

[74] _____, *Beyond the sequence: cellular organization of genome function*, Cell **128** (2007), no. 4, 787–800.

[75] G.L. Moldovan, B. Pfander, and S. Jentsch, *PCNA, the Maestro of the Replication Fork*, Cell **129** (2007), no. 4, 665–679.

[76] C. Münkel and J. Langowski, *Chromosome structure predicted by a polymer model*, Physical Review E **57** (1998), no. 5, 5888–5896.

[77] J. Odenheimer, G. Kreth, and D.W. Heermann, *Dynamic Simulation of Active/Inactive Chromatin Domains*, Journal of Biological Physics **31** (2005), no. 3, 351–363.

[78] L. Rayleigh, *Investigations in optics with special reference to the spectroscope, 1: resolving, or separating power of optical instruments*, Philosophical Magazine. Reprinted in: Scientific Papers of Lord Rayleigh **1** (1879), 8.

[79] W. H. Richardson, *Bayesian-based iterative method of image restoration*, J. Opt. Soc. of Am. **62** (1972), 55–59.

[80] C.L. Rieder and ED Salmon, *The vertebrate cell kinetochore and its roles during mitosis*, trends in CELL BIOLOGY (Vol. 8) (1998), 311.

[81] U.P. Roos, *Light and electron microscopy of rat kangaroo cells in mitosis*, Chromosoma **41** (1973), no. 2, 195–220.

[82] M.J. Rust, M. Bates, and X. Zhuang, *Sub-diffraction-limit imaging by stochastic optical reconstruction microscopy (STORM)*, Nature Methods **3** (2006), 793–796.

[83] A. San Paulo and R. Garcia, *High-Resolution Imaging of Antibodies by Tapping-Mode Atomic Force Microscopy: Attractive and Repulsive Tip-Sample Interaction Regimes*, Biophysical Journal **78** (2000), no. 3, 1599–1605.

[84] B. Schneider, *Spatially modulated illumination microscopy*, Ph.D. thesis, Universität Heidelberg, 1999.

[85] A. Schönle, P.E. Hänninen, and S.W. Hell, *Nonlinear fluorescence through intermolecular energy transfer and resolution increase in fluorescence microscopy*, Ann. Phys.(Leipzig) **8** (1999), no. 2, 115–133.

[86] M. Schrader, K. Bahlmann, G. Giese, and S. W. Hell, *4pi-confocal imaging in fixed biological specimens*, Biophysical Journal **75** (1998), 1659–1668.

[87] M. Schrader, S. W. Hell, and H. T. M. van der Voort, *Three-dimensional superresolution with a 4pi-confocal microscope using image restoration*, Journal Of Applied Physics **84** (1998), no. 8, 4033–4042.

[88] N.C. Shaner, P.A. Steinbach, and R.Y. Tsien, *A guide to choosing fluorescent proteins*, Nat Methods **2** (2005), no. 12, 905–909.

[89] CE Shannon, *Communication in the presence of noise*, Proceedings of the IRE **37** (1949), no. 1, 10–21.

[90] C.M. Sparrow, *On Spectroscopic Resolving Power*, The Astrophysical Journal **44** (1916), 76.

[91] U. Spöri, *Messungen mit dem SMI-Mikroskop*, Ph.D. thesis, Universität Heidelberg, 2004.

[92] U. Spöri, A.V. Failla, and C. Cremer, *Superresolution size determination in fluorescence microscopy: A comparison between spatially modulated illumination and confocal laser scanning microscopy*, Journal of Applied Physics **95** (2004), 8436.

[93] T. Staudt, MC Lang, R. Medda, J. Engelhardt, and SW Hell, *2, 2'-Thiodiethanol: A new water soluble mounting medium for high resolution optical microscopy.*, Microsc Res Tech **70** (2007), no. 1, 1–9.

[94] E. Steinwand, *4Pi Mikroskopie - Anwendung und Analyse bildverarbeitender Methoden*, Master's thesis, Universität Heidelberg, 2006.

[95] C. Stockklausner and N. Klöcker, *Surface expression of inward rectifier potassium channels is controlled by selective golgi export*, J. Biol. Chem. **278** (2003), 17000–17005.

[96] S.M. Tan, *Aperture synthesis mapping and parameter estimation*, Ph.D. thesis, Mullard Radio Astronomy Observatory, Cavendish Laboratory, Cambridge, 1987.

[97] S.M. Tan and C. Fox, *Physics 707 inverse problems (course notes)*, `http://www.math.auckland.ac.nz/~phy707/`.

[98] F.A. Tanious, J.M. Veal, H. Buczak, L.S. Ratmeyer, and W.D. Wilson, *DAPI (4', 6-diamidino-2-phenylindole) binds differently to DNA and RNA: minor-groove binding at AT sites and intercalation at AU sites*, Biochemistry **31** (1992), no. 12, 3103–3112.

[99] R.Y. Tsien et al., *THE GREEN FLUORESCENT PROTEIN*, Annual Review of Biochemistry **67** (1998), no. 1, 509–544.

[100] T. Tsukamoto, N. Hashiguchi, S.M. Janicki, T. Tumbar, A.S. Belmont, and D.L. Spector, *Visualization of gene activity in living cells*, Nature Cell Biology **2** (2000), 871–878.

[101] I. Upmann, *Aufbau und Anwendung mikroskopischer Systeme zur Messung an Einzelmolekülen*, Ph.D. thesis, Universität Heidelberg, 2002.

[102] O. Vafa and K.F. Sullivan, *Chromatin containing CENP-A and alpha-satellite DNA is a major component of the inner kinetochore plate*, Curr. Biol **7** (1997), 897–900.

[103] AM van Oijen, J. Köhler, J. Schmidt, M. Müller, and GJ Brakenhoff, *3-Dimensional super-resolution by spectrally selective imaging*, Chem. Phys. Lett **292** (1998), 183–187.

[104] C. Wagner, *Neue Methoden zur Größen- und Formbestimung von fluoreszierenden Nanostrukturen unter Verwendung der SMI-Mikroskopie*, Ph.D. thesis, Universität Heidelberg, 2004.

[105] JY WANG and DE SILVA, *Wave-front interpretation with Zernike polynomials*, Applied Optics **19** (1980), no. 9, 1510–1518.

[106] V. Westphal and S.W. Hell, *Nanoscale Resolution in the Focal Plane of an Optical Microscope*, Physical Review Letters **94** (2005), no. 14, 143903.

[107] S. Wilhelm, B. Gröbler, M. Gluch, and H. Heinz, *Die konfokale laser scanning mikroskopie*, Brochure, Carl Zeiss Mikroskopsysteme, D-07740 Jena, 2003.

[108] F. Zernike, *Diffraction and Optical Image Formation*, Proceedings of the Physical Society **61** (1948), no. 2, 158–164.