Elucidating the mechanism of mRNA transport using novel imaging technologies



Danail Stoychev
Exeter College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy
Trinity 2023

Abstract

One of defining features of the cells of the nervous system is their strong polarity. They have long projections that extend to distances that could be thousands of times the length of the cell body. To support their function and to maintain homeostasis, the cells need to supply signals and resources like proteins to such remote sites away from the nucleus. This is particularly important for dynamic processes such as cytoskeletal rearrangement, membrane trafficking, and synaptic plasticity. There is considerable evidence that at least some of this peripheral activity is regulated by the transport and local translation of mRNA molecules at the distant cytoplasm of neurons. The molecular mechanisms which govern the post-transcriptional regulation of these processes in neurons is still not fully understood, and even less is known in glia. This is partly due to the technically challenging aspect of capturing live mRNA dynamics and partly because only a few examples have been studied in detail. I begin to fill this technological gap by assessing and applying state-of-the-art methods from the fields of microscopy and image processing. This includes the use of adaptive optics for aberration correction and for remote (optical) focusing, 3D structured illumination microscopy for optical sectioning and for super-resolution, and the development of an open-source software platform that can enable all this functionality in fast complex experiments. On the image processing side, it includes the application of deep learning to the denoising of fluorescence image data for fast live imaging with minimised phototoxicity and photobleaching. Empowered by all these technological advancements I establish an improved live motility assay for the study of mRNA dynamics, using the *Drosophila* peripheral nervous system and especially the neuromuscular junction, which is arguably the simplest and easiest to use model to study these processes live. The MS2-MCP system is used to insert into mRNAs of interest special structures to which RNA-binding proteins, fused to a fluorescent protein such as GFP, can then bind. I apply this assay to investigate a specific mRNA, *shot*, previously identified as encoding important protein regulators of synaptic development and linked to neurodevelopmental disorders.

Statement of Originality

I declare that, to the best of my knowledge, the contents of this thesis are original and product of my own work, unless otherwise stated. The thesis has not been submitted, in whole or in part, for any other degree or qualification in this or any in any other institution. The work described in this thesis was carried out in the Department of Biochemistry, Oxford University between October 2019 and October 2023.

I confirm that the work carried out in collaboration with others is clearly identified in the contents of the thesis. For convenience, it is also summarised here:

- Parts of the software developments described in chapter 3 were done in collaboration with Matthew Wincott. These include the interactive specification of ROI, the aberration correction results window, the DM viewer, and the generalisation of the correction routines.

- The latest iteration of the upright imaging system described in chapter 4 was built and maintained by Jingyu Wang. It had a long history and many people had contributed to its development over the years. The adaptive optics work described in the same chapter was also carried out jointly with Jingyu Wang. The diagram in Figure 4.1 is based on a drawing by Jingyu Wang. The sample in Figure 4.22 was prepared by Dalia Gala and imaged by Jingyu Wang. The sample in Figure 4.27 was prepared by Andreas Gerondopoulos.

- The *shot*-MS2 fly line used in chapter 6 was developed by Francesca Robertson. She also prepared the samples described in this chapter and she acquired the data presented in Figure 6.1.

Statement of Mitigation

The work described in this thesis was severely affected by multiple different circumstances, which created critical constraints and necessitated various adjustments along the way.

The first major circumstance was the COVID-19 pandemic, particularly the lockdowns and other restrictions throughout 2020 and 2021. They severely limited the ability to do lab work, especially microscopy development which required access to specific equipment and facilities. Moreover, the pandemic aligned with the beginning of the DPhil project, which was a crucial period for learning for someone who comes from an engineering background. This period was intended to build up the fundamental wet lab skills required to support the project, by receiving hands-on lessons and mentoring, including how to work with *Drosophila* and how to carry out various sample preparation protocols. The pandemic is one of the reasons why the project has a significant software development component to it, because this was something that could be done remotely and for the most part independently.

Another circumstance which had a significant impact on the DPhil project was the unavailability of the microscopy system described in chapter 4. Due to delays and other unforeseen events, largely caused by the COVID-19 pandemic, the funding for the instrument ended before it could be used to study live mRNA dynamics. Without funding, the instrument had to be disassembled and the respective laboratory spaces had to be vacated.

Finally, a few months before the end of the DPhil project, the primary supervisor made the decision to move to a different institution. This added more work and stress to the already difficult period of finalising the project. The move also made some of the laboratory equipment unavailable, because it was packed and prepared to be shipped to the new location.

The main consequence of all these circumstances, which severely impacted the trajectory of the project, is that it was not possible to build a single unified imaging

system, capable of both the adaptive optics and the image denoising technologies described in chapters 4 and 5, respectively. The original idea was to use both of them in the improved motility assay described in chapter 6, but ultimately only the image denoising was used in the respective experiments, because the microscope with the adaptive optics did not exist at that point. To address this complication, I present in chapter 6 theoretical work, in the form of simulation, which demonstrates the need for the combined technologies for studying live mRNA dynamics.

Acknowledgements

I want to thank my supervisors, Prof Ilan Davis and Prof Martin Booth, for their continual support and kindness, for putting their trust in me, and for sharing their vast knowledge and experience with me. I also want to thank Richard Parton for teaching me the fundamentals of microscopy and fly work, amongst many other things, and for always being so reliable, patient, and a good company. Similarly, I had an amazing time working together with Jingyu Wang, whose passion and hard work were always an inspiration. I want to thank him for teaching me so much about optics, adaptive or not, and about engineering in general.

Despite many challenges, even a global pandemic, I had a really enjoyable time during my degree and this was in no small part due to the wonderful environment fostered by the Davis Lab and the Dynamic Optics and Photonics group. I want to give special thanks to Darragh Ennis for always being so helpful, knowledgeable, and reliable, and for bringing worryingly many sweets to the lab. Fran for working together with me, for teaching me about flies and genetics, and for being so cheerful and accommodating. Dalia, Jeff, and Maria who were around throughout my entire degree and whose company I enjoyed greatly. Also all the other members from the Davis Lab who made this such a brilliant experience: Aino, Ana, Arianna, David, Jana, Josh, Julia, MK, Martin, Sophia, Tai, and all the Part II students! I want to thank all the people from across the road, too, for being welcoming and friendly, and for always engaging my curiosity: Andrew, Asad, Chao, Jiahe, Karen, Matthew, Patrick, Qi, and everyone else. I want to use this opportunity to also thank everyone I interacted with from the Castello Lab: Alfredo, Cat, Enzo, Kate, Manuel, Marko, and Wael.

I am also really grateful for the opportunity to work with all the other people from the Micron research group who I have not mentioned already – David, Ian, Lothar, Mick, and Nick – as well as everyone else involved in the bioimaging facility: Carina, Deirdre, Nadia, and Niloufer.

I am thankful to the Biotechnology and Biological Sciences Research Council, the

Oxford Interdisciplinary Bioscience Doctoral Training Programme, and Aurox Ltd. for providing this exceptional opportunity and for awarding me the studentship. My industrial placements at Aurox were really enjoyable and I learned a lot from them. I am grateful to everyone in the company, especially Rimas Juškaitis and Tony Wilson.

I want to thank Exeter College as well, for all the support they have given me in the last four years, for being so reliable and well-ordered, and for providing me with excellent accommodation for such a long time.

Contents

Glossary **11**

1 Introduction **13**

1.1 Overview and motivation 13

1.2 Aims and objectives 16

1.3 Document structure 17

2 Theoretical background **18**

2.1 Structure and characteristics of the cells of the nervous system 18

2.1.1 Polarity 18

2.1.2 Tri-partite system 18

2.1.3 Cytoskeleton 19

2.1.4 Long-distance transport of mRNA 19

2.2 The genetic toolkit of *Drosophila* 20

2.2.1 Targeted gene expression 20

2.2.2 Gene knockdown 20

2.2.3 Labelling of mRNA 20

2.3 Light microscopy for live imaging 21

2.3.1 Producing contrast via fluorescence 21

2.3.2 The harmful effects of light on living matter 21

2.3.3 Detection of light and image formation 22

2.3.4 Types and sources of noise in an imaging system 22

2.4 Adaptive optics for fluorescence microscopy 23

2.4.1 Optical aberrations 23

2.4.2 Characterising the performance of an imaging system 24

2.4.3 Classification of aberrations 25

2.4.4 Wavefront sensing . . . 29
2.4.5 Aberration correction algorithms . . . 31
2.4.6 Types of adaptive elements . . . 35
2.4.7 Remote focusing . . . 36
2.5 Denoising of live fluorescence images . . . 38
2.5.1 Ill-posed nature of the denoising problem . . . 38
2.5.2 Traditional image denoising algorithms . . . 41
2.5.3 Neural networks for image denoising . . . 43

3 Developing software for the control of imaging systems 47
3.1 Introduction . . . 47
3.2 Aims and objectives . . . 50
3.3 Results . . . 51
3.3.1 Functional software changes . . . 51
3.3.2 Improvements to the GUI . . . 55
3.4 Discussion and conclusions . . . 57

4 Adaptive optics for live fluorescence imaging 61
4.1 Introduction . . . 61
4.2 Aims and objectives . . . 63
4.3 Methods and materials . . . 64
4.3.1 Microscope design . . . 64
4.3.2 Control of Alpao DM69-15 deformable mirror . . . 66
4.3.3 Calibration of the deformable mirror (DM) . . . 70
4.3.4 Sensorless aberration correction . . . 73
4.3.5 Correction of system aberrations . . . 76
4.3.6 Structured illumination microscopy . . . 78
4.3.7 Fly stocks . . . 78
4.3.8 Sample preparations and fluorescent labelling . . . 79
4.4 Results . . . 80
4.4.1 Remote focusing for fast motionless volumetric imaging . . . 80
4.4.2 Sensorless aberration correction in the *Drosophila* PNS . . . 84
4.4.3 Sensorless aberration correction for deep imaging . . . 85

4.4.4 Sensorless adaptive optics for live imaging . . . 87
4.4.5 Multi-plane aberration correction . . . 87
4.4.6 Extended focus . . . 87
4.5 Discussion and conclusions . . . 89

5 Denoising of fluorescence images **93**
5.1 Introduction . . . 93
5.2 Aims and objectives . . . 96
5.3 Methods and materials . . . 97
5.3.1 Microscopy . . . 97
5.3.2 Computational environment . . . 97
5.3.3 Datasets and sample preparation . . . 97
5.3.4 Selection of deep-learning denoising methods . . . 98
5.3.5 Baseline denoising method . . . 99
5.3.6 Neural network training . . . 99
5.3.7 Image quality metrics . . . 101
5.4 Results . . . 102
5.4.1 Acquisition of training and evaluation data . . . 102
5.4.2 Denoising performance . . . 106
5.5 Discussion and conclusions . . . 113

6 Developing a motility assay for mRNA dynamics **118**
6.1 Introduction . . . 118
6.2 Aims and objectives . . . 120
6.3 Methods and materials . . . 121
6.3.1 Fly strains and sample preparation . . . 121
6.3.2 Microscopy . . . 121
6.3.3 Image processing . . . 122
6.4 Results . . . 122
6.4.1 Labelling *in vivo* live single mRNA molecules in a tissue-specific manner . . . 122
6.4.2 Image processing pipeline . . . 124
6.4.3 Particle track analysis . . . 126

6.4.4 Simulation of optical aberration and noise 129
6.5 Discussion and conclusions 129

7 Conclusions and future work **133**
7.1 Software for microscopy 133
7.2 Adaptive optics for fluorescence microscopy 134
7.3 Denoising of fluorescence images 135
7.4 Live mRNA motility assay 137

Appendices **140**

A Details about the Alpao DM69-15 **141**
A.1 Firmware bug 141
A.2 Step response 141

B Mathematical proofs **146**
B.1 Minimising the MSE 146
B.2 Comparison between noisy and clean training targets 147

C Complete denoising metrics **148**

D smFISH probe sequences **152**

Glossary

CMOS complementary metal–oxide–semiconductor. 22, 23, 29, 70

CNN convolutional neural network. 44–46, 93, 95, 96, 135, 137

CNS central nervous system. 15, 119

DM deformable mirror. 8, 35, 36, 52, 60–62, 64, 65, 67, 68, 70–72, 76, 77, 80, 82, 89, 92, 134, 143, 145

GFP green fluorescent protein. 21, 79, 85, 87, 98, 122

GUI graphical user interface. 49, 50, 52, 55, 57, 66, 133

HRP horseradish peroxidase. 79, 85

MCP MS2-coat protein. 21, 118, 119, 121, 122, 124, 131, 138

mRNA messenger RNA. 14–17, 20, 96, 97, 118–122, 124–126, 129, 131, 132, 137–139

MSE mean squared error. 46, 101, 108, 122, 146, 147

NMJ neuromuscular junction. 15, 18, 79, 84, 85, 119, 120, 138

NSS non-local self-similarity. 42, 135

OTF optical transfer function. 25, 34

PCA principal component analysis. 126, 138

PNS peripheral nervous system. 15, 18, 119–121, 124, 139

PSF point-spread function. 24, 25, 34, 77, 84, 129, 131

PSNR peak signal-to-noise ratio. 101, 102, 106, 108, 109

RFP red fluorescent protein. 79, 87

RI refractive index. 27, 29, 36, 85

RNA ribonucleic acid. 20, 21

RNAi RNA interference. 20, 120

ROI region of interest. 55, 143

SIM structured illumination microscopy. 52, 63, 64, 78, 84, 85, 87, 92

SLM spatial light modulator. 35, 52, 61, 62, 65, 78

smFISH single-molecule fluorescence in situ hybridisation. 20, 138

SNR signal-to-noise ratio. 93, 103, 138

SSIM structural similarity index measure. 101, 102, 106, 109

UAS upstream activation sequence. 20, 120, 121

YFP yellow fluorescent protein. 79, 85

Chapter 1

Introduction

1.1 Overview and motivation

One of the most fascinating, and certainly characteristic, properties of the cells of the nervous system is their polarity. They exhibit strong asymmetry and there is directionality to their function. As a consequence of this, they feature cytoplasmic projections which can reach lengths that are thousands of times longer than the diameter of the cell body and which make connections with thousands of other cells, with specialised structures called synapses. Despite the commonality between these projections and their respective synapses, they have individual functionality, transmitting and receiving information independently, and thus they also have their own unique requirements for molecular signals and resources, such as proteins.

According to flow of genetic information, protein is an end product which traces its origin to the genome, almost all of which is located in the nucleus in the cell body. However, even fast protein transport can take more than a week to arrive from the nucleus to the far ends of some of the longest human axonal projections [1]. But the nervous system is highly dynamic, constantly integrating information from the environment, and the peripheral activity of its cells requires reactions on much shorter time scales. As a side note, the matters of efficiency and energy expenditure in the delivery of molecular signals and resources are important, too, but not strictly necessary for the success of the delivery.

These contrasting observations lead to the important question how can the cells of nervous system support a diverse functionality in remote sites and at reasonable time

scales. There is emerging evidence that the answer to this riddle is decentralisation, where some messenger RNAs (mRNAs) are transported and locally translated, on demand, in a very flexible, efficient, and responsive manner [2]. However, the underlying molecular mechanisms are not yet completely understood in neurons and even less is known in glia. The need to elucidate the role of mRNA transport is further reinforced by the fact that its dysregulation in the nervous system is linked to various states of disease, including spinal muscular atrophy, amyotrophic lateral sclerosis, frontotemporal dementia, and Fragile X syndrome [3, 4]. Together with being a vital part of the machinery that makes up life, mRNA is perhaps the most important biomolecule to examine in terms of transport. So important, in fact, that it was recently discovered that it even hitchhikes when other cargo is being transported, such as lysosomes [5].

There are different ways to study the transport of mRNA, but fluorescent microscopy is the only method that can provide the necessary spatial and temporal resolution, as well as contrast, to disentangle the dynamic behaviour [6]. The fluorescent labelling of molecules in preparation for this method is already a challenge of its own, where it is not always possible to generate sufficient contrast, but there are many technical obstacles to the microscopy, too. Some of them include:

- Photobleaching – fluorescence-emitting molecules are inherently unstable and causing them to emit light leads eventually to an irreversible permanent depletion
- Phototoxicity – light is harmful to living matter, leading to stress and thus non-physiological condition, and in large amounts leads to death
- Light scattering – biological tissue usually has high optical heterogeneity which causes light to scatter and limits its penetration depth
- Optical aberrations – refractive index mismatch and deficiencies in the optical system cause the light to deviate from its ideal course, leading to reduced contrast and resolution

Live fluorescence microscopy can be used with cultured human cells but this is a rather artificial and simplistic system. Various model organisms have emerged as a way to overcome these limitations and to study complex phenomena. Amongst them, *Drosophila* stands out as a particularly suitable candidate for studies of the nervous

system. The ease of use and the powerful genetic toolkit, coupled with modern mRNA labelling techniques, presents a unique opportunity to investigate mRNA transport live and in vivo. Especially convenient is the peripheral nervous system (PNS) and the neuromuscular junctions (NMJs) which provide the means to examine single motor neurons, their interaction with glial cells, and the synaptic connections they make with muscles. Depending on the developmental stage and the sample preparation there are ways to easily access axonal projections and NMJs for live imaging with standard light microscopes. Examining the central nervous system (CNS) is an option, too.

I make use of this unique opportunity to develop an improved motility assay for studying mRNA transport. I address some of the challenges associated with live fluorescence imaging by assessing and applying state-of-the-art technologies related to microscopy and image processing. These technologies are not arbitrary. They are specifically chosen for their relevancy to the task, and evidence hints at their strong potential.

One of these technologies is adaptive optics, where light is reshaped using special type of devices. The main applications are aberration correction and remote focusing, which is the ability to shift the focus in a purely optical way. The advantages of the former are obvious, as reduced aberrations lead to increased resolution and contrast, and to more accurate measurements, which make it particularly important in quantitative microscopy. Remote focusing, on the other hand, can bring various improvements depending on the way in which it is implemented, but a common advantage is the motionless acquisition of volumetric data. This minimises the disturbances to the sample, caused by traditional mechanical actuation, and it creates new experimental opportunities, such as fast volumetric imaging combined with electrophysiology, where micropipettes can remain stationary without the danger of breaking them or perturbing the sample.

The use of adaptive optics requires sophisticated software control. I implement this on top of an existing free open-source microscopy software platform. It benefits from the available infrastructure for fast live imaging and it pushes the system to its limits, taking full advantage of the underlying hardware. For convenience, the software and the adaptive optics are part of an upright structured-illumination super-resolution microscope. Although there is a trade-off between resolution, light dose, and speed, this kind of super-resolution methods are one of the most gentle and most suitable for

live imaging.

In terms of image processing, I evaluate the performance of several state-of-the-art deep-learning methods for image denoising. The advantage of this ability to lower the excitation light dose, which is the only way to reduce the harmful effects of phototoxicity and photobleaching, but it results in noisy data. Deep-learning methods have recently revolutionised certain areas of computer vision, such as object recognition. It is speculated that they can recapture some of this success in other domains, such as image denoising, but this remains to be proven. Unfortunately, many deep-learning denoising methods are developed on synthetic or non-representative datasets, and therefore gauging their performance is difficult without practical assessment.

Finally, I demonstrate the potential of the motility assay by investigating the dynamics of a specific mRNA which is known to be present in the distant parts of neurons and glia, but not much is known about its transportation. That being said, the focus is the generic assay and the particular gene is merely an interesting test subject.

1.2 Aims and objectives

The aim is to use novel imaging technologies to develop an improved motility assay for studying live mRNA dynamics in the *Drosophila* nervous system. The objectives are the following:

- Build the software foundation necessary for adaptive optics and establish a unified platform for fast live imaging
- Implement adaptive optics capable of sensorless aberration correction and remote focusing, and assess their performance for live fluorescence imaging
- Evaluate state-of-the-art deep-learning denoising methods and develop a pipeline suitable for the restoration of live low-light fluorescence image data
- Demonstrate the effectiveness of the motility assay by investigating the dynamics of a specific mRNA

1.3 Document structure

Chapter 1 introduced the topic of the thesis and the motivation behind it. Chapter 2 presents the theory behind all the different topics covered in the thesis. Chapter 3 describes all software developments, most of which are later used as part of the work in subsequent chapters. Chapter 4 describes the implementation and assessment of adaptive optics for aberration correction and for remote focusing. Chapter 5 presents the assessment of several deep-learning denoising methods and compares them to a baseline traditional method. Chapter 6 describes the improved motility assay and demonstrates its practical use with a concrete mRNA. Finally, chapter 7 summarises the main results, draws conclusions, and outlines future research directions.

Chapter 2

Theoretical background

2.1 Structure and characteristics of the cells of the nervous system

2.1.1 Polarity

The nervous system consists of two types of cells: neurons and glia. Both of them exhibit strong polarity, which leads to long cytoplasmic projections which extend from the cell body, branch off, and make connections with thousands of other cells on average. Projections usually represent more than 99% of the cytoplasmic content of the cell and their length is thousands of times that of the cell body [7].

2.1.2 Tri-partite system

Neurons are not the only participant in the structure of the synapse. They are accompanied by glial cells, and the interaction between the two is important for the regulation of many processes. There are many different types of glial cells. Three of them are present in the *Drosophila* PNS. Perineurial and subperineurial (sometimes grouped together as perisynaptic) glia cover the surface of the brain and its axonal projections, while wrapping glia surrounds the neurons inside axonal projections. Together, all three form the fly blood-brain barrier. A schematic of this arrangement, both for a cross section of an axonal projection and for the NMJ, is shown in Figure 2.1.

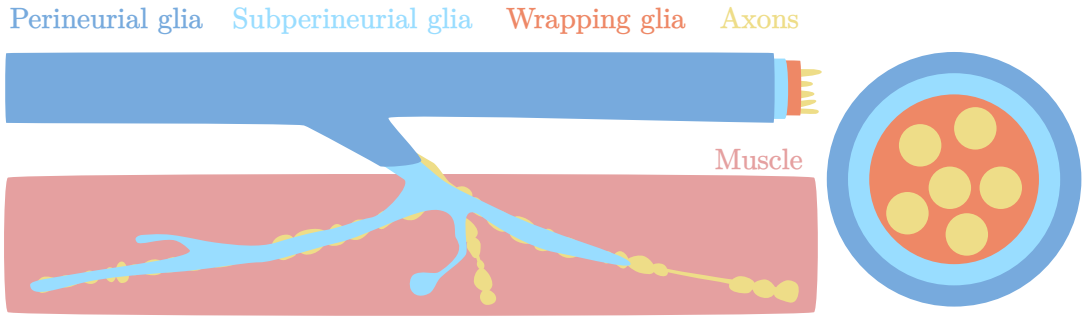


Figure 2.1: Neurons and glial cells at the neuromuscular junction, and a cross-section of the axonal projection.

2.1.3 Cytoskeleton

One of the defining components of all cells, including those of the nervous system, is the cytoskeleton. This a complex network of filaments, made of protein, which provides cells with mechanical and structural support, as well as helping with the transport of biomolecules. There are three types of filaments, with diameter sizes from around 6 nm to 25 nm. The finest are microfilaments (actin), followed by intermediate filaments, and the thickest are microtubules.

Actin and microtubules are polarised and they have directionality, with the two ends of a filaments called plus and minus ends. In neurons, actin and microtubles can have both uniform and mixed polarity, depending on the compartment of the cell. For example, in axonal projections the minus ends point towards the nucleus, whereas in dendrites the polarity is mixed.

The cytoskeleton, much like the cell it supports, is a dynamic structure which is constantly being reorganised to accommodate the needs of the cell.

2.1.4 Long-distance transport of mRNA

There are two main ways of moving molecules around a cell: passively via diffusion and actively via molecular motors. These are protein complexes that convert the energy harnessed from ATP hydrolysis into mechanical motion. They are most often associated with the group of cytoskeletal motors which attach to the filaments of the cytoskeleton and transport molecular cargo along them: myosins in the case of actin and kinesins and dyneins in the case of microtubules. Kinesins move towards the plus ends, while dyneins towards the minus ends. The binding between cargos and motors is accomplished by

special class of proteins called adaptors.

2.2 The genetic toolkit of *Drosophila*

2.2.1 Targeted gene expression

One of the most powerful genetic tools in *Drosophila* is the ability to selectively activate a gene. This is most commonly achieved with Gal4, a transcription factor which comes from yeast [8]. It binds to upstream activation sequence (UAS) which then activates transcription of the neighbouring gene. On its own, Gal4 does not have a strong effect in flies, which also do not naturally have many UAS regions. It is only when the two come together that the targeted gene is turned on.

In practice, the *GAL4* gene is inserted into the *Drosophila* genome in such a way that it is under the control of a promoter or an enhancer which is expressed only in specific cells or tissues. These fly lines are called drivers. An example is the pan-glial driver Repo-Gal4 which would make all glial cells emit fluorescence under excitation with light when combined with a fluorescent protein under the control of UAS.

2.2.2 Gene knockdown

The translation of mRNA to protein can be suppressed by small non-coding ribonucleic acid (RNA) in a process called RNA interference (RNAi). This effectively inhibits the expression of the gene and significantly lowers the protein concentration, but it does not eliminate it entirely, as in the case of a knockout where the gene is completely inactivated or removed.

Drosophila has large libraries of RNAi lines that cover a significant proportion of the genome. They are usually fused to UAS for targeted gene silencing by Gal4. This is particularly useful in functional studies where knockdown experiments can elucidate some of the roles of a gene.

2.2.3 Labelling of mRNA

There are various ways to label mRNA molecules, with some of the most common ones being single-molecule fluorescence in situ hybridisation (smFISH) and the insertion of MS2 stem loops. The former is suitable for fixed material and it works by hybridising

a set of complementary oligonucleotide probe sequences to a target RNA. More probes correspond to stronger signal, and they can also be labels with many different types of fluorophores.

The MS2 stem loops come from a bacteriophage of the same name and they represent special types of secondary RNA structure [9]. The loops are recognised by the MS2-coat protein (MCP) which can be fused to a fluorescent protein or other reporters. The insertion of the MS2 sequence is usually in the 3' or 5' untranslated region.

2.3 Light microscopy for live imaging

2.3.1 Producing contrast via fluorescence

Fluorescence is a process of emitting light after the absorption of energy, usually in the form of light. Molecules possessing this property are called fluorophores. They have delocalised electrons which can jump from one stable "ground" state to an unstable "excited" state. Eventually, an electron returns back to the ground state by emitting the extra energy in the form of light. In other words, when light of specific wavelength is shone on a fluorophore, the fluorophore will absorb photons, become excited for a brief moment (called lifetime), and then return to its ground state by emitting photon of lower energy than the one that excited it.

All living cells are fluorescent to a certain degree, because some of the molecules responsible for their correct function, such as flavins, are fluorophores. Fluorescence caused by such endogenous fluorophores is called autofluorescence, to distinguish it from exogenous sources, such as dyes, and artificial sources, such as genetically engineered fluorescent proteins. One of the most successful and widely used such proteins is green fluorescent protein (GFP) [10].

2.3.2 The harmful effects of light on living matter

While in its excited state, fluorophores are very volatile and they react with oxygen, because of its unusual triplet ground state. This permanently modifies the fluorophore and it enters a non-emitting state, called dark state. This process is known as photobleaching. During photobleaching, reactive oxygen species are also generated, which are toxic to living organisms and lead to deleterious physiological conditions when

accumulated.

2.3.3 Detection of light and image formation

Special light sensitive cameras are used to do capture the emission of light by a fluorophore. Such cameras have sensors that consist of arrays of pixels, each being a separate detector of its own. When photons land on the pixel they get absorbed and stored as electric charge. In this sense, the camera sensor is a device which converts light into electrical energy. This electric charge is eventually read out by the camera's electronics and converted into an intensity value, which is what makes up the final image. It is an array of intensities, corresponding to how many photons were converted into electric charge at each pixel over the exposure duration. There are different types of camera technologies, with the current state of the art being complementary metal–oxide–semiconductor (CMOS), and particularly the scientific grade. In these devices each pixel has its own set of electronics used for charge conversion. Pixel charges are eventually converted into digital signals, where each column of the array has a dedicated converter.

An alternative detector is the photomultiplier tube, which is used in applications where the emission light is scanned over the detector, such as the two-photon and the laser-scanning confocal microscopes. The reason is that such devices consist of a single detector, instead of an array of detectors like cameras, which has high sensitivity at the expense of loss of spatial information.

2.3.4 Types and sources of noise in an imaging system

The conversion of light into electrons by the camera is an imperfect process and besides losses it also suffers from stochasticity. The read out values can be described by a normal (Gaussian) distribution centred at the true value. The situation is further exacerbated by the quantum nature of light and the fact that photons can only arrive in batches, which results in a separate source of fluctuation called shot or photon noise. The amount of actual photon counts recorded by the detector will vary according to a Poisson distribution.

A pixel will thermally generate electrons, even in the absence of incident photons, which leads to dark noise that contributes to the same signal fluctuations that follow a Poisson distribution. Modern fluorescence imaging involves relatively short exposure

times and coupled with advances in cooling technology there is little heat generated in the sensor, and therefore dark noise is usually ignored. Finally, CMOS-based detectors are also subject to fixed pattern noise, which is caused by slight differences in the electronics of each pixel or the converters at each column, and it is not of stochastic nature in the same way that the other types of noise are.

2.4 Adaptive optics for fluorescence microscopy

2.4.1 Optical aberrations

Light was described mostly as a stream of photons in section 2.3. An equally valid interpretation is that of an electromagnetic wave, whose value at point x in space and at time t, is given by:

$$\psi(x,t) = \psi_0 \cdot e^{i(kx-\omega t)} \tag{2.1}$$

In equation 2.1 ψ_0 is the amplitude, $k = 2\pi/\lambda$ is the wave number with wavelength λ, and $\omega = 2\pi f$ is the angular frequency with ordinary frequency f. The phase of such a wave is given by $\arg(\psi) = kx - \omega t$.

When light waves travel, the surfaces of equal phase are called wavefronts. There are two main types of electromagnetic waves – plane and spherical. When a wavefront deviates from its ideal shape it is said to be optically aberrated. These aberrations lead to errors in measurements when light is used for probing, as in the fields of astronomy and microscopy. One of the main goals of adaptive optics is to correct the optical aberrations and therefore to restore the ideal wavefront.

When light is used to create an image, the light field at points in the object is mapped to points on a detector. Typically the detector measures the intensity of the light, which is effectively the photon density given by $|\psi|^2 = \psi_0^2$ and it is proportional to the number of photons. This requirement to collect or focus light, from or to a point, is often achieved with lenses. A plane wave passing through a positive lens becomes spherical. Likewise, a spherical wave turns planar. In this sense, an aberrated wavefront leads to an aberrated focal spot, which is no longer a simple point.

An optical system is a finite physical system and therefore the admitted light is always limited by what is called an aperture stop. Typically it corresponds to a circular opening, but this is not a requirement and it could equally be rectangular or some other

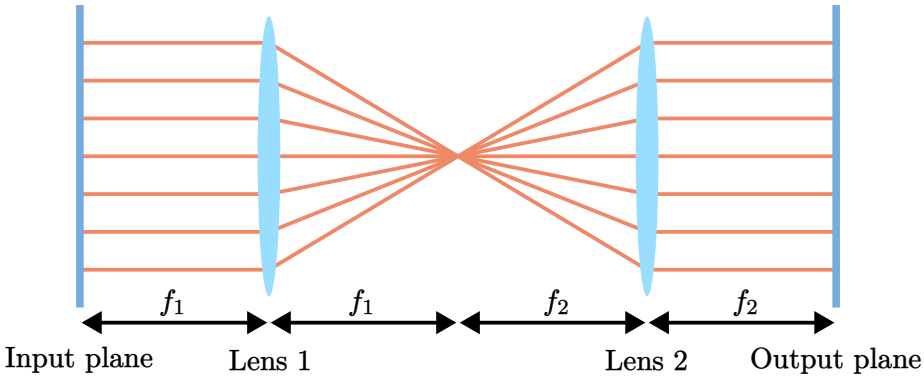


Figure 2.2: Reimaging of a plane with the use of a 4f relay system. Lenses 1 and 2 have focal lengths f_1 and f_2, respectively. The magnification factor is $-f_2/f_1$.

shape. Image of the aperture stop is called pupil and it is formed at a pupil plane. The light travelling through an optical imaging system can be described by the system's pupil function which corresponds to the wavefront (equation 2.1) at the pupil plane.

Accessing the aperture stop of an imaging system is not always convenient. For example, microscope objective lenses with high magnification have their pupil plane inside the objective lens' barrel. For this reason, pupil planes are often re-imaged with a 4f imaging system. This is an optical relay with two lenses and four focal lengths between the two conjugated planes, as shown in Figure 2.2, which is where the name comes from. The magnification of the system is determined by the ratio of the focal lengths of the two lenses: $M = -f_2/f_1$.

2.4.2 Characterising the performance of an imaging system

Optical imaging systems are diffraction limited. They act as low-pass filters of spatial frequency and thus there is a limit to the smallest resolvable feature (highest spatial frequency). Therefore, a point of the object plane results in a spread-out blurry spot in the image, with the shape of the spot determined by the parameters of the system. The function that describes this mapping is called point-spread function (PSF) and for an incoherent system it is equivalent to an impulse response, which provides a complete characterisation of the system's behaviour, including aberrations. Conceptually, the PSF can be thought of as the paint brush used to trace the object onto the image. Mathematically, this takes the form of convolution with the object plane $O(x, y)$ to produce the intensity image $I(x, y)$: $I(x, y) = \text{PSF} * O(x, y)$. In practice, there is no need for an infinitesimally small point to determine the PSF, any object below the

diffraction limit of the system will suffice. The Fourier transform of the PSF is called the optical transfer function (OTF) of the system, and it carries the same information but in frequency space. The PSF and the OTF can be directly derived from the pupil function by taking the squared magnitude of its Fourier transform or by autocorrelation, respectively.

2.4.3 Classification of aberrations

Optical aberrations can be described by the shape of the wavefront. Notable early attempt to provide a general mathematical treatment of the subject is the work of Gauss around the middle of the 19th century [11], building upon previous work from "Cotes, Euler, Lagrange, and Möbius" to derive what is known today as Gaussian optics. Gauss gave precise definitions of focal length and the cardinal points in geometrical optics, on systems with rotational symmetry, leading to a linear description of image formation. However, the elegance of the relations described by Gauss holds true only for rays that make small angles with the common axis of the lenses through which they pass, known as paraxial approximation. In general, the description of rays in geometrical optics involves trigonometric functions. For example, the rules of refraction are governed by a ratio of sine functions, which can be expanded in the following power series:

$$\sin(\theta) = \theta - \frac{\theta^3}{3!} + \frac{\theta^5}{5!} - \frac{\theta^7}{7!} + \cdots = \sum_{n=0}^{\infty} \frac{(-1)^n}{(2n+1)!}\theta^{2n+1} \tag{2.2}$$

Effectively, Gauss assumed a first-order approximation where $\sin(\theta) \approx \theta$. Such an approximation describes the general principles behind image formation but it ignores any deviations from this ideal condition and thus it ignores aberrations.

One year after the death of Gauss, Philipp Ludwig von Seidel published further work on optical systems with rotational symmetry, this time with the addition of the second term (third-order) of the power series, which was the first mathematical treatment of image errors and it led to the formal definition of five classic aberrations with the strongest influence on the quality of an image[12]:

- spherical aberration – focus varies with the distance from the optical axis
- coma – magnification varies with the distance from the optical axis

- astigmatism – asymmetry between horizontal and vertical planes leads to variation in the focus and the magnification, with distance from the optical axis
- field curvature – field-dependent differences in focus
- distortion – field-dependent differences in magnification

The analytical approach of Seidel has been revisited and refined several times since its inception, such as in [13–15], and it remains a useful tool in lens design even today. An alternative approach, considering wave propagation and diffraction in an optical system, was established by Hopkins in the middle of the 20th century and it allows for easier calculation of high order aberrations [16]. In practice, real aberrated wavefronts are usually a combination of many different types of aberrations, but they are dominated by the lower order ones.

Often times in adaptive optics it is desirable to describe the shape of the wavefront at a specific point, such as the pupil. Zernike polynomials provide a convenient way to do this for circular pupils. They were initially conceptualised by Zernike in the first half of the 20th century [17]. They were later explicitly defined, and proven to be the only sequence of polynomials with the three properties of orthogonality, rotational invariance, and existence of polynomial of every degree [18]. There are many ways to order, index, and normalise the Zernike polynomials, as reviewed in [19]. One of the most popular strategies is that of Noll [20], where they are defined as:

$$Z_j(r,\theta) = \begin{cases} \sqrt{n+1}R_n^m(r) & \text{when } m = 0 \\ \sqrt{2(n+1)}R_n^m(r)\cos(m\theta) & \text{when } m \neq 0 \text{ and } j \text{ is even} \\ \sqrt{2(n+1)}R_n^m(r)\sin(m\theta) & \text{when } m \neq 0 \text{ and } j \text{ is odd} \end{cases} \tag{2.3}$$

$$R_n^m(r) = \sum_{s=0}^{(n-m)/2} \frac{(-1)^s(n-s)!}{s!(\frac{n+m}{2}-s)!(\frac{n-m}{2}-s)!} r^{n-2s} \tag{2.4}$$

$$n = \left\lfloor \sqrt{2j-1} - 0.5 \right\rfloor \tag{2.5}$$

$$m = \begin{cases} 2 \cdot \left\lfloor \dfrac{2j+1-n(n+1)}{4} \right\rfloor & \text{when } n \text{ is even} \\ 2 \cdot \left\lfloor \dfrac{2(j+1)-n(n+1)}{4} \right\rfloor - 1 & \text{when } n \text{ is odd} \end{cases} \tag{2.6}$$

Just like any linear dynamic system can be described by the superposition of its

normal modes, so can any continuous surface with a circular support be decomposed to a linear combination of enough Zernike polynomials. Mathematically, the wavefront phase ϕ at polar coordinates (r, θ) for a pupil of radius R can be expressed as:

$$\phi(R \cdot r, \theta) = \sum_{j=0}^{J} a_j \cdot Z_j(r, \theta) \tag{2.7}$$

This generally presents an overdetermined linear system of equations, as usually there are more (r, θ) points than modes J. In this case, the Zernike coefficients a_j can be found with least squares or other optimisation methods. Alternatively, due to the orthogonality of the Zernike polynomials, the coefficients can be expressed as:

$$a_j = \frac{\int_0^1 \int_0^{2\pi} \phi(r, \theta) Z_j(r, \theta) r \, dr d\theta}{\int_0^1 \int_0^{2\pi} r \, dr d\theta} \tag{2.8}$$

With this approach, simple numerical integration can yield the coefficient values.

One of the greatest advantage of Zernike modes is that they correspond to classic aberrations, such that low order aberrations can be adequately represented by low order modes. For example, modes 4 and 11, using Noll indexing, correspond to defocus and first order spherical aberrations, respectively. It should be noted, however, that Zernike polynomials are defined for a wavefront at a single field point (r, θ) and therefore cannot capture well field-dependent aberrations like curvature or distortion. At a point, they appear as defocus and as tilt, respectively. A selection of Zernike modes, their shapes, and their physical interpretations are listed in Figure 2.3. Besides Zernike polynomials, other orthonormal bases have been used to describe optical aberrations, such as Chebyshev polynomials [21] and polar Walsh functions [22].

All the aberrations described so far in this subsection have been monochromatic, in the sense that a single wave, of any wavelength, is subject to them. The use of light of different wavelengths can induce a class of inter-wavelength errors called chromatic aberrations. They are a consequence of dispersion and more precisely the fact that the refractive index (RI) of a material depends on the wavelength. For positive focusing lenses, this means that light waves of different wavelengths travel at different velocities, over different distances, and converge at different foci. There are two types of chromatic aberrations: axial and lateral. The former corresponds to differences in focus

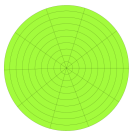
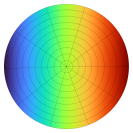
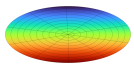
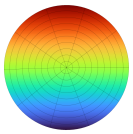
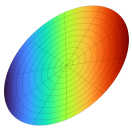
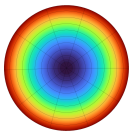
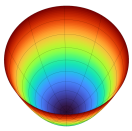
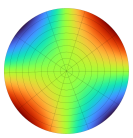
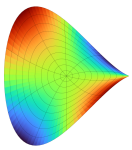
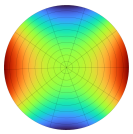
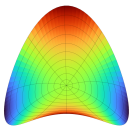
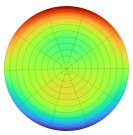
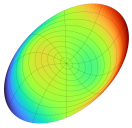
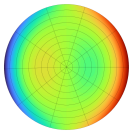
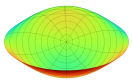
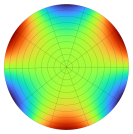
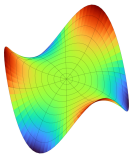
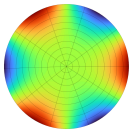
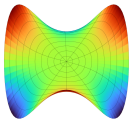
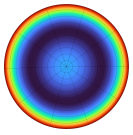
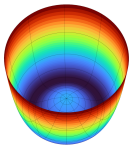
Visualisation 2D	Visualisation 3D	Noll index	Name	Polynomial
		1	Piston	1
		2	Horizontal tilt	$2r\cos(\theta)$
		3	Vertical tilt	$2r\sin(\theta)$
		4	Defocus	$\sqrt{3}(2r^2-1)$
		5	Primary oblique astigmatism	$\sqrt{6}r^2\sin(2\theta)$
		6	Primary vertical astigmatism	$\sqrt{6}r^2\cos(2\theta)$
		7	Primary vertical coma	$\sqrt{8}(3r^3-2r)\sin(\theta)$
		8	Primary horizontal coma	$\sqrt{8}(3r^3-2r)\cos(\theta)$
		9	Primary vertical trefoil	$\sqrt{8}r^3\sin(3\theta)$
		10	Primary oblique trefoil	$\sqrt{8}r^3\sin(3\theta)$
		11	Primary spherical	$\sqrt{5}(6r^4-6r^2+1)$

Figure 2.3: The first 11 Zernike polynomials. The 3D view is rotated 90 degrees clockwise with respect to the 2D view.

between different wavelengths and it is usually impossible to correct once the image has been formed. Lateral chromatic aberrations lead to radial differences in magnification, similar to coma, and they can be corrected computationally in a post-processing step.

Optical aberrations can also be categorised according to their origin. The ones induced externally by the environment, such as the atmosphere in astronomy and the sample (including medium) in microscopy, are caused by RI mismatch. The rest are then a consequence of the imaging system itself. Such system aberrations are caused either by flaws in the components and their arrangement, or they are an inevitable consequence of the design of the optical system, such as the use of curved mirrors in reflecting telescopes.

2.4.4 Wavefront sensing

There are different ways to detect the wavefront of a light wave. One of the most commonly used devices is a Shack-Hartmann sensor [23]. It consists of an array of small lenses (lenslets) focused on a regular light detector such as a CMOS camera. In the absence of aberrations, the planar wavefront will result in a grid of spots in the detector. When there are aberrations, however, the positions deviate from this ideal array and create an irregular arrangement of spots. These deviations are proportional to the local slopes of the wavefront, making it possible to reconstruct its entire surface and in this way to directly measure the aberrations.

The wavefront can also be sensed with certain types of interferometer, such as Twyman-Green [24]. The typical layout of such an instrument is shown in Figure 2.4. The light from a laser is split by a beamsplitter into two paths: reference and measurement. The latter is focused on the plane at which the wavefront is to be detected. Both paths are reflected back into the beamsplitter where they interfere with each other. The resulting fringe pattern is imaged onto a camera. The intensity of this fringe pattern is expressed mathematically as:

$$I(x, y) = a(x, y) + b(x, y) \cos(2\pi f_0 x + \phi(x, y)) \tag{2.9}$$

The $a(x, y)$ and $b(x, y)$ terms represent the amplitudes of the background and the interference signals, respectively, and $\phi(x, y)$ is the phase. Using Euler's formula, this

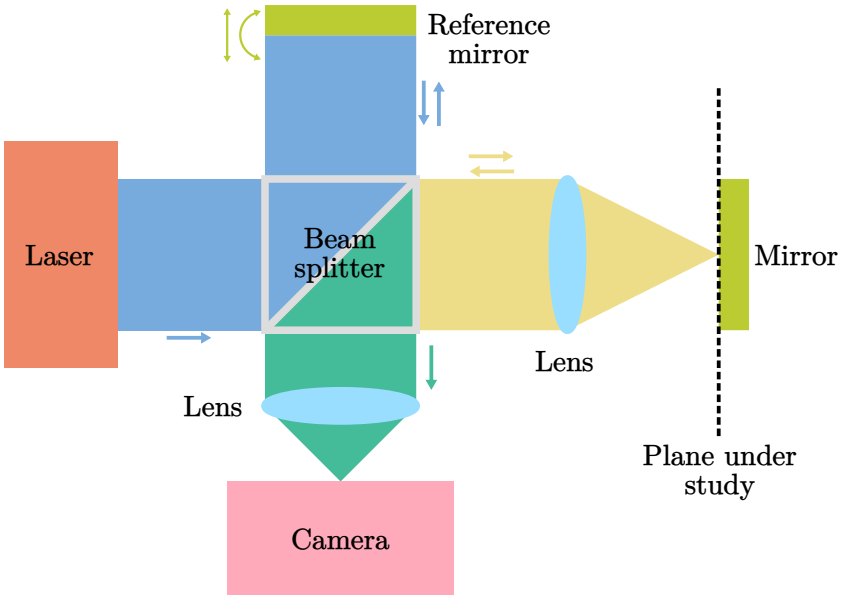


Figure 2.4: Layout of a Twyman-Green interferometer. The beam splitter separates the light into reference path (light blue) and measurement path (light yellow). The two are recombined (mint) and imaged onto a camera.

can be rewritten as:

$$I(x,y) = a(x,y) + c(x,y)e^{2\pi i f_0 x} + c^*(x,y)e^{-2\pi i f_0 x}; \quad c(x,y) = \frac{1}{2}b(x,y)e^{i\phi(x,y)} \tag{2.10}$$

Reconstruction of the phase from the fringe pattern can be achieved with several different algorithms. Common class of such algorithms is phase-stepping [25, Chapter 6], where multiple images of different fringe patterns are captured, each corresponding to different reference phase, which can be achieved by varying the rotation angle of the reference mirror or by translating it axially, as indicated by the arrows next to it in Figure 2.4. For example, four different intensity images with steps $\{0, \pi/2, \pi, 3\pi/2\}$ can be used to find the phase with:

$$\phi(x,y) = \arctan\left(\frac{I_4(x,y) - I_2(x,y)}{I_1(x,y) - I_3(x,y)}\right) \tag{2.11}$$

Alternatively, Fourier-based algorithms can be used to reconstruct the phase from a single image, without the need to vary the position or orientation of the reference mirror [26]. This procedure is illustrated in Figure 2.5. In the Fourier domain (Figure 2.5b),

equation 2.10 takes the form:

$$I(u,v) = A(u,v) + C(u - f_0, v) + C^*(u + f_0, v) \tag{2.12}$$

Both spectra C and C^* contain the same signal content, and they are symmetrically distributed around the 0 point (and around A). Therefore, the phase can be retrieved by filtering one of them (Figure 2.5c), shifting it to the origin, taking the inverse Fourier transform, and then finally taking the argument of the resulting complex signal (Figure 2.5d). Mathematically, this is expressed as:

$$\phi(x,y) = \arg(c(x,y)); \quad c(x,y) = \mathcal{F}^{-1}\{C(u - f_0, v)\} = \mathcal{F}^{-1}\{C^*(u + f_0, v)\} \tag{2.13}$$

Unfortunately, the results of equations 2.11 and 2.13 are ambiguous because they come down to the arctangent of a sinusoidal function, which is periodic:

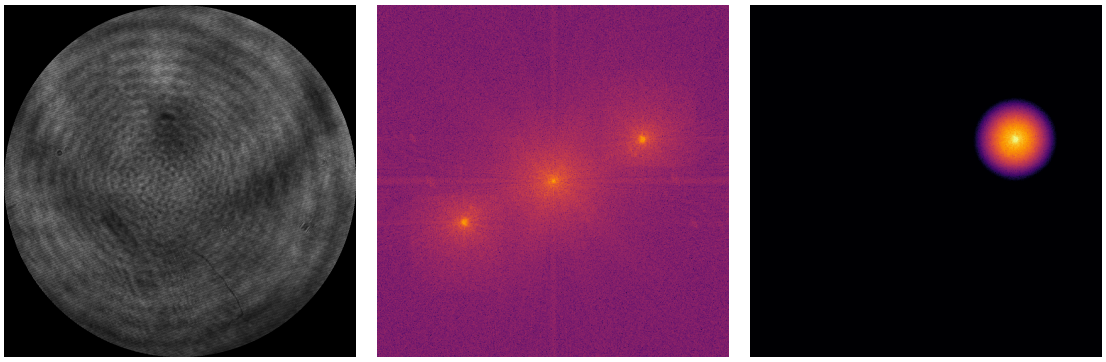
$$\arctan(\cos(x)) = \arctan(\cos(x + 2n\pi)) \quad \forall n \in \mathbb{Z} \tag{2.14}$$

For example, a phase of 5π will be wrapped twice to a final value of π. To deal with this issue, the phase can be unwrapped with a standard algorithm such as [27] (Figure 2.5e). The unwrapped phase can be converted from radians to length by multiplying it with $\lambda/2\pi$, where λ is the wavelength of the light.

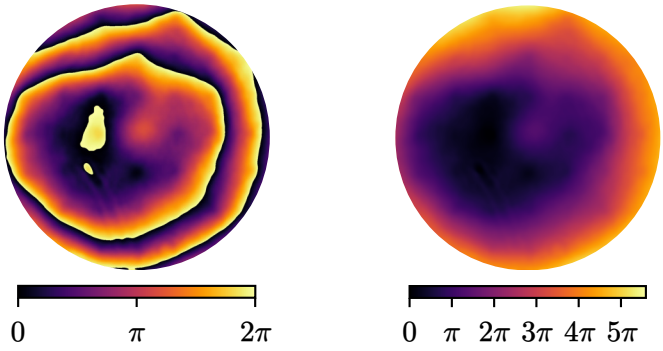
2.4.5 Aberration correction algorithms

Optical aberrations are corrected by reshaping the wavefront so that it as close to its ideal shape as possible. This is achieved with the help of special devices called adaptive elements, which can modulate the phase of incident light and they can work either by transmitting or reflecting the light. The goal is to compensate the distortions by applying a shape of opposite direction and suitable amplitude, as shown in Figure 2.6. Conceptually, this is straightforward to do with a wavefront sensor such as Shack-Hartmann which measures the aberrations directly, making it easy to determine how to reshape the light with the adaptive element.

In situations where it is not possible or not desirable to use a wavefront sensor, various indirect correction methods can be used. The goal of such sensorless approaches

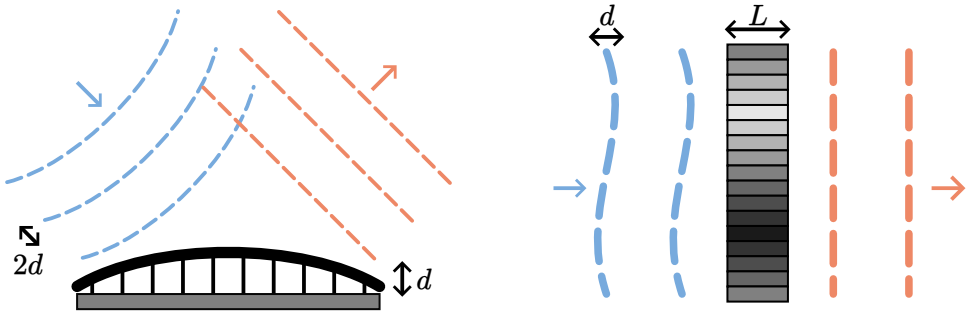


(a) Interferogram. (b) Fourier transform. (c) Filtered spectra.



(d) Wrapped phase. (e) Unwrapped phase.

Figure 2.5: Wavefront sensing with interferometry. Panels (b) and (c) show the logarithm of the power spectrum of the 2D Fourier transform of (a).



Reflective Transmissive

Figure 2.6: Compensation of wavefront aberrations with reflective and transmissive devices. On the left, the curvature of the mirror needs to be half of that of the incoming wavefront, because it will be affected twice, once on incidence and once on reflection. On the right, the different shades of grey indicate the different RIs, n, of the device. In this case, successful compensation requires the change in displacement of the wavefront to be matched to the change in RI, such that $\Delta d = \Delta n L$.

is usually to determine the necessary phase compensation from regular (aberrated) intensity images. The approach closest to the direct methods is zonal control, where the adaptive element is partitioned into segments. This technique has been applied to various types of optical systems [28, 29]. Imaging of each segment individually and comparing the result to a reference allows to measure the local slope of the wavefront, in a similar manner to how a Shack-Hartmann sensor operates. Alignment of the segmental images to the reference compensates the slopes and flattens the wavefront, thus correcting the aberrations.

An alternative to the zonal methods is modal indirect methods. The idea is to represent the aberrations as a combination of modes, like in equation 2.7 for the case of Zernike polynomials. The orthogonality of the modes is important, because it allows to interrogate them one by one, effectively scanning for aberrations. This algorithm is illustrated in Figure 2.7 where seven points are examined for Z_8, equally distributed in the range $a_8 \in [-10; 10]$ rad. For each of these scanning points, the adaptive element adds the mode to the wavefront, with a corresponding coefficient, and the resulting intensity image is recorded. In the presence of aberration of the same mode within the scanning range, applying a coefficient value of opposite sign produces an image of higher quality. This variation in quality is symmetric around the aberration (or around 0 in the absence of aberrations) and can be modelled with a Gaussian distribution, with its mean representing the magnitude of the aberration. In Figure 2.7, the Z_8 aberrations are corrected with $a_8 = -2.79$ rad. The procedure can then be repeated for as many modes as necessary. Various image quality metrics exist, measuring properties such as intensity or contrast, or performing Fourier-domain analysis to capture features such as sharpness and signal content.

Although zonal and modal indirect methods were developed separately as alternative approaches, they could be unified in a single framework where the pupil is segmented and each segment (zone) is represented by a set of modes [30]. This framework includes purely zonal method with multiple zones and a single mode, purely modal method with a single zone and multiple modes, and hybrid methods with multiple zones and multiple methods.

Modal-based sensorless methods cannot compensate very well field-dependent aberrations, a problem known as anisoplanatism. The situation is further complicated

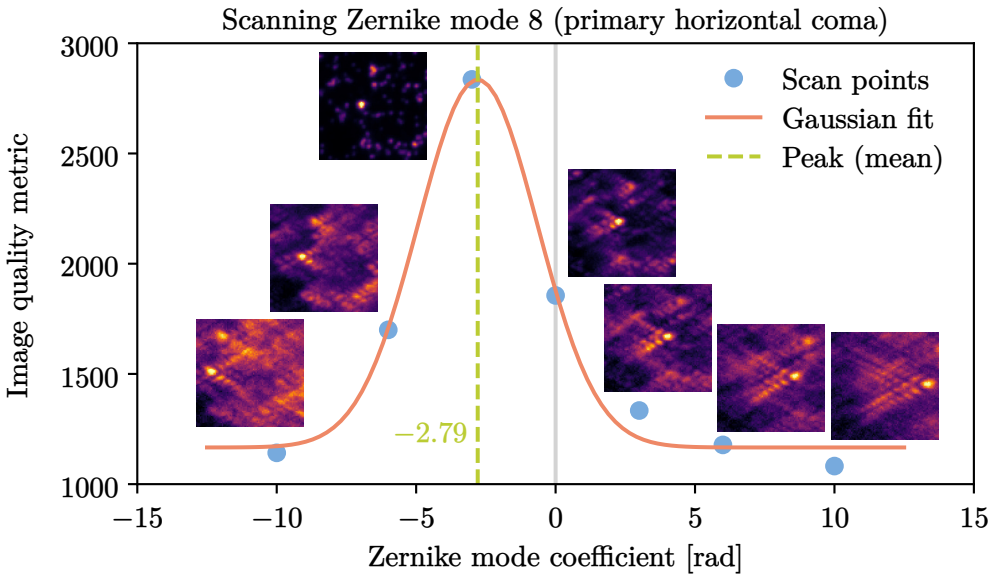


Figure 2.7: Modal wavefront sensorless aberration correction routine for Z_8. Different coefficients are applied and the corresponding image quality metric values are fitted to a Gaussian curve. The peak of the curve corresponds to the magnitude of the aberrations present for this mode.

when the modes are Zernike polynomials, which are only defined for a single point of the wavefront, as discussed in subsection 2.4.3. Field-dependent aberrations are usually largely corrected by the optical components, particularly by the objective lens where they are a major design consideration. However, certain optical systems exhibit significant residual field-dependent aberrations, and this has been reported for both water-immersion objective lenses and remote focusing [31, 32]. Zonal-based algorithms, that use parallel correction of multiple isoplanatic subregions, have shown promising results in laser-scanning confocal microscopy [33].

In subsection 2.4.2 it was described how the PSF and OTF can be directly derived from the pupil function. The opposite is not trivial to do because the two contain only intensity information, and they lack the proper amplitude and phase information necessary to describe a wave like the pupil function. The problem of reconstructing the phase from intensity images is called phase retrieval and it is yet another method of indirect wavefront sensing [34]. A variation of it is phase diversity, which is an approach that eases the computational difficulty by taking two or more images with a known phase difference between them, such as simple defocus [35].

To make corrections, both direct and indirect methods need to compare the wavefront of the incoming light to some ideal (reference) shape. Usually, bright diffraction-

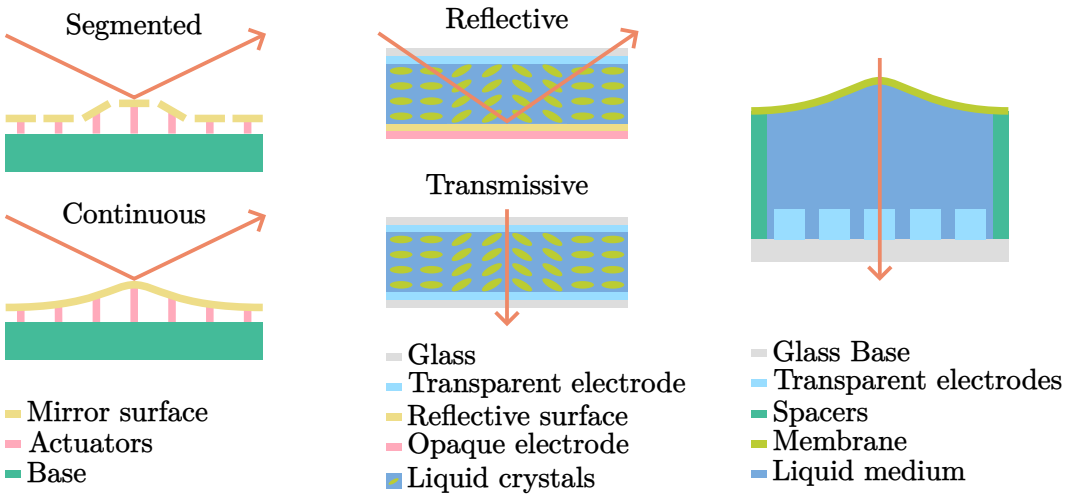
limited (point-source) objects are used for this purpose. They emit spherical waves which turn planar after collimation, which then makes the calculation of the wavefront error trivial. The presence of such objects, referred to as guide stars, is not guaranteed to occur naturally in many applications, especially in microscopy. This is why they are often introduced artificially, such as by focusing a laser into the field of view [36]. In microscopy, the same effect can be achieved by adding fluorescent beads or proteins to the sample [37]. It is also possible to use structured illumination to create high-frequency features in the image, taking on similar role as point sources [38].

2.4.6 Types of adaptive elements

Adaptive elements can be distinguished by three key parameters:

- spatial resolution – encompasses the size of the pixels or actuators of the adaptive element, as well as their number and arrangement; higher resolution allows for finer control over the shape of the wavefront and therefore the correction of higher-order aberrations which have increasingly higher intricacy.
- dynamic range (stroke) – measures the maximum possible amplitude of wavefront distortion; the stronger the aberration, regardless of its order, the more stroke is necessary to correct for it.
- speed – the rate at which actuators can be displaced or pixels can be reconfigured; it is usually the application which dictates how often the light needs to be reshaped and thus what speed is required.

There are several different types of optical devices that can be used to directly change the shape of light. Figure 2.8 illustrates some of the most common such devices. DMs [39] (Figure 2.8a) are reflective devices where a reflecting surface is made either of an array of small mirrors or of a continuous membrane. The overall shape is determined by a set of actuators, each individually controlled by applied voltage. This could be based on electromagnetic force, in which case the actuators can be pulled and pushed around a neutral position. Devices like these offer good dynamic range and speed. The actuators can also be controlled with electrostatic attraction or with piezoelectric and electrostrictive materials. Spatial light modulators (SLMs) [40], shown in Figure 2.8b are based on liquid crystal technology, thus requiring polarised light. They could be



(a) DM (b) SLM (c) DPP

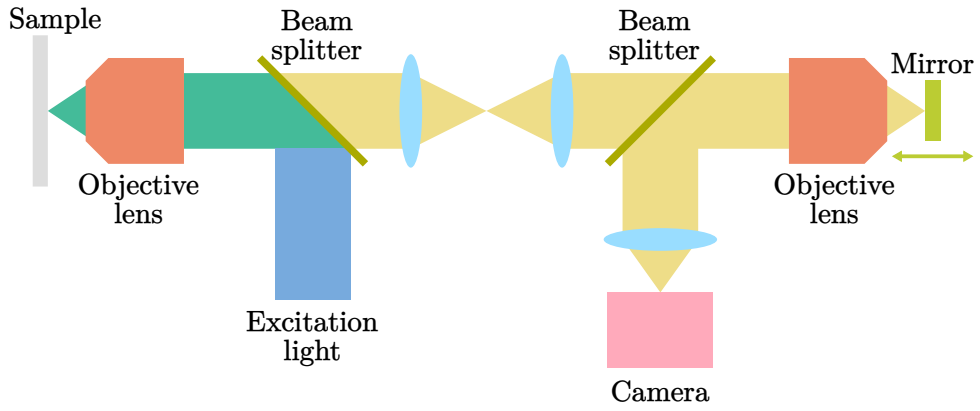
Figure 2.8: Three different types of adaptive elements, devices capable of shaping the wavefront of light. The olive arrows indicate the way light is reflected or transmitted by the device.

either transmissive or reflective. Instead of actuators, the optical area is divided into pixels, where the RI of each pixel is individually controlled. They have good resolution and dynamic range but typically slow speed. Another class of devices are deformable phase plates [41], which are similar to DMs, but they are based on optofluidics and they are transmissive (Figure 2.8c). Because of this they are compact and easy to integrate in an optical system, but they have lower resolution and dynamic range than other types of devices. Finally, digital micromirror devices [42] consist of an array of miniature mirrors and therefore they are reflective. Each mirror can be switched between two fixed positions, which makes the devices binary. They excel in speed and resolution.

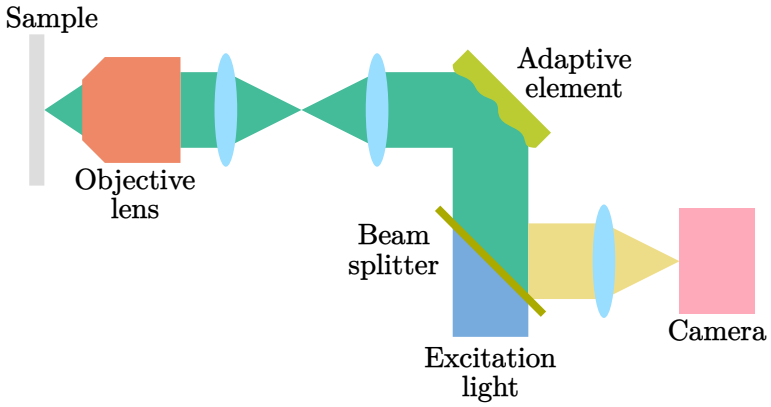
2.4.7 Remote focusing

The goal of some optical imaging systems, such as microscopes, is to produce a 3D volume of the object. This is usually achieved with optical sectioning, where the volume is composed of a set of sections, taken while the object is translated axially with respect to the focus plane. There are two difficulties with this approach: the speed is usually slow and the motion could induce disturbances to the object.

With the addition of extra optical components to the system, it becomes possible to create a 3D replica of the object, which can then be scanned by a simple mirror



(a) Two-objective layout. Refocusing is achieved by translating the mirror on the right.



(b) Single-objective layout. Refocusing is achieved by reshaping the wavefront with the adaptive element.

Figure 2.9: Remote focusing optical configurations.

[43]. This configuration, shown in Figure 2.9a, keeps the object stationary, while the mirror is translated axially, which is a lot simpler and faster. The disadvantages of this approach are the increased complexity and cost of the system, the reduced optical efficiency which is inversely proportional to the number of optical components, and the high sensitivity to magnification mismatch [44].

Alternatively, as noted in [43], it is possible to use adaptive optics to change the shape of the light in such a way that it results in a focal shift. The layout of such an optical system is shown in Figure 2.9b. The caveat to this approach is that the image quality is dictated by how accurately the adaptive element can represent the phase associated with focal shift. In this sense, such optical refocusing is different process than its mechanical counterpart, but the two could achieve similar results. When the change in phase is expanded in a common basis such as Zernike modes, it includes

terms for both defocus and spherical aberrations of increasing order [45]. This could present a problem, as adding Z_4 to the system will result in a more aberrated image. An adjustment to the definition of Zernike modes can decoupled these two quantities, allowing independent control of focus and aberration correction [46].

2.5 Denoising of live fluorescence images

2.5.1 Ill-posed nature of the denoising problem

There are different ways an image can be corrupted with noise. The one of interest to fluorescence microscopy is additive noise. Mathematically it can be described in the following way:

$$y_{i,j} = x_{i,j} + n_{i,j} \tag{2.15}$$

This expression tells us that the measured intensity of a single pixel $y_{i,j}$, where i and j denote indices along the vertical and horizontal directions of the image, respectively, is a combination of the true signal $x_{i,j}$ and a noise component $n_{i,j}$. For convenience the pixel subscript will often be omitted or flattened to a single dimension when spatial information is not important for the argument at hand.

A realistic model of the noise would combine Gaussian and Poisson distributions to capture the statistics of the different types of noise in an imaging system, as described in subsection 2.3.4. Foi et al present such a model in [47]:

$$n = n_g + n_p(x) \tag{2.16a}$$

$$n_g \sim \mathcal{N}(0, \sigma^2) \tag{2.16b}$$

$$\alpha(x + n_p(x)) \sim \mathcal{P}(\alpha x) \tag{2.16c}$$

This model is parameterised by the variance of the Gaussian distribution $\sigma^2 \geq 0 \in \mathbb{R}$ and the gain of the imaging system $\alpha > 0 \in \mathbb{R}$. It accounts for the fact that the signal x and one of the noise components n_p are drawn from a Poisson distribution. It follows that the mean of both noises is zero and that the variance of the Poisson noise scales



(a) Noisy image. (b) Three plausible clean images.

Figure 2.10: Three different plausible explanations, amongst many, of the same noisy image.

with the signal (number of detected photons):

$$\mathbb{E}\left[n_g\right] = 0 \tag{2.17}$$

$$\mathbb{E}\left[n_p(x)\right] = 0 \tag{2.18}$$

$$\mathrm{Var}\left(n_p(x)\right) = x/\alpha \tag{2.19}$$

Often, for simplicity and convenience, the noise is modelled as just a Gaussian distribution, where $n \sim \mathcal{N}(0, \sigma^2)$. This approximation is based partially on the central limit theorem and partially on the fact that $\mathcal{P}(\lambda) \approx \mathcal{N}(\lambda, \lambda)$ for high values of λ.

Regardless of how the noise is modelled, the clean image can always be explained by multiple noisy realisations, as illustrated in Figure 2.10. The presence of multiple solutions makes the problem underdetermined and therefore ill-posed, as it fails the second of Hadamard's criteria for a well-posed problem [48]:

1. [Existence] A solution exists
2. [Uniqueness] There is one solution
3. [Stability] The solution depends continuously on the inputs

Ill-posed problems need to be modified to be solved. This is usually done via regularisation, which is a technique that expands the problem statement by adding prior information about x, such as what the clean image is expected to look like. It helps to constrain and reduce the solution space, and to add stability. Broadly speaking, there are two ways to then solve the problem: deterministic and probabilistic. In the deterministic case, the objective is to estimate a point, a singular value. For images, this point estimation is usually performed for each individual pixel and it can be expressed

in the following generic form:

$$\hat{x} = \underset{x}{\operatorname{argmin}} \{F(x, y) + \lambda R(x)\} \tag{2.20}$$

$F(x, y)$ is the data fidelity term and it describes the closeness of the estimate to the true value. One of the most common choices is the squared L^2 norm, which for 2D images yields:

$$F(x, y) = \|x - y\|_2^2 = \sum_i \sum_j \left(\sqrt{(x_{i,j} - y_{i,j})^2}\right)^2 = \sum_i \sum_j (x_{i,j} - y_{i,j})^2 \tag{2.21}$$

The second term is the regularisation term, the strength of which is controlled by the λ parameter. To give an example, one of the oldest and most well-established regularisation methods is that of Tikhonov [49]:

$$\begin{aligned} R(x) &= \|\nabla x\|_2^2 \\ &= \sum_i \sum_j \left(\sqrt{(x_{i+1,j} - x_{i,j})^2 + (x_{i,j+1} - x_{i,j})^2}\right)^2 \\ &= \sum_i \sum_j \left[(x_{i+1,j} - x_{i,j})^2 + (x_{i,j+1} - x_{i,j})^2\right] \end{aligned} \tag{2.22}$$

On the other hand, the probabilistic approach to the denoising problem can be realised with Bayesian inference, which leads to an expression for a posterior probability $P(x \mid y)$:

$$P(x \mid y) = \frac{P(y \mid x) \cdot P(x)}{P(y)} \propto P(y \mid x) \cdot P(x) \tag{2.23}$$

The posterior represents an elaborate distribution of all clean images consistent with the noisy image. Similarly, the likelihood $P(y \mid x)$ corresponds to all possible noise realisations of the clean image, and effectively serves as a noise model. The prior $P(x)$ is the probability of observing the clean image in the first place and it captures pre-existing knowledge about the problem. Various estimators, such as maximum likelihood and maximum a posteriori, can then be used to estimate the clean image. In the case of the latter, the optimisation problem can be formulated in the following way:

$$\hat{x} = \underset{x}{\operatorname{argmax}}\, P(x \mid y) = \underset{x}{\operatorname{argmin}} \{-\log P(y \mid x) - \log P(x)\} \tag{2.24}$$

As was the case with equation 2.20, the first term describes the fidelity and in this context it is called negative log likelihood. The second term acts as regulariser and it is called negative log prior.

Probabilistic approaches provide more complete description of the denoising problem, incorporating knowledge and uncertainty of arbitrary complexity, but only the simplest of formulations have an analytical solution. This drawback is addressed by various numerical methods, such as sampling [50] and variational inference [51]. However, they are often computationally demanding. The accuracy of probabilistic approaches also depends on the aptness and accuracy of the noise model, amongst other things, which is not always easy or possible to get right. Because of all of these difficulties, probabilistic denoising methods have been less popular than their deterministic counterparts.

2.5.2 Traditional image denoising algorithms

Traditionally, practical denoising algorithms have been hand-crafted and tailored to a specific problem or a domain thereof. After all, a concrete mathematical description is required, including domain knowledge in the form of regularisation. This is in contrast to learning-based methods which have gained a lot of popularity recently [52].

One of the main differences between the traditional point estimators is the choice of prior, determining what information about the clean image will be used to narrow down the possible solutions and to reach them in a stable manner. For example, the Tikhonov method mentioned in subsection 2.5.1 is based on smoothness and the idea that natural images do not have abrupt changes in intensity. The interpretation on the right in Figure 2.10b would be penalised in this case. Unfortunately, this also makes it difficult to preserve sharp features and the Tikhonov method is known for its tendency to over-smooth [53]. An alternative prior is total variation regularisation [54], which in its discrete isotropic (L^2) form is defined as:

$$R(x) = \|\nabla x\|_2 = \sum_i \sum_j \sqrt{(x_{i+1,j} - x_{i,j})^2 + (x_{i,j+1} - x_{i,j})^2} \tag{2.25}$$

This is essentially the same as Tikhonov (equation 2.22), but without the square. It is precisely the square root that manages to preserve sharp features, such as edges, better.

However, it is not differentiable which makes it harder to implement and compute, and it comes with its own drawbacks, such as loss in contrast [55].

One of the most well-established priors used for regularisation, and the basis for many state-of-the-art traditional denoising algorithms, is non-local self-similarity (NSS) [56]. It is based on the fact that natural images have regions (or blocks), which are similar to each other (self-similarity) and not necessarily close to each other (non-local). This property could be used, for example, to group and filter these blocks. This is the idea behind one of the most popular and successful algorithms for removal of zero-mean Gaussian noise, routinely used as a benchmark even today: block-matching and 3D filtering (BM3D) [57]. It works in two steps, both of which do similar things. The first step takes the noisy image and produces a preliminary clean estimate, using hard thresholding for the filtering of groups of blocks. The second step then repeats this process, but using the preliminary estimate as an input and using Wiener filtering. The hard thresholding in the first step requires knowledge of the noise variance, which makes the algorithm non-blind, in the sense that it needs both an image and its noise statistics as an input. Unfortunately, this limits severely the usefulness of the algorithm for real-world applications, where the noise statistics are usually unknown.

The concept of filtering of non-local groups of blocks is further extended to multi-dimensional image data in nD-SAFIR [58], which is specifically designed for denoising of live fluorescence images. Unlike many other algorithms, BM3D included, it is completely blind, working solely on noisy images, and it accounts for the fact that the noise distribution is not just Gaussian but Poisson-Gaussian. The signal-dependency is decoupled with a variance-stabilising transformation and all filtering is done in this domain before applying an inverse transformation to produce the final estimate. The grouping of blocks (called patches in this context) is based not only on similarity but also on the noise variance. Patches are either 2D or 3D regions. In the latter case, a group of patches becomes a 4D structure. As the algorithm is designed around time-series experiments, patches are sought not only spatially but temporally, too. The groups of patches are denoised by weighted averaging of the intensity values. No motion estimation is involved.

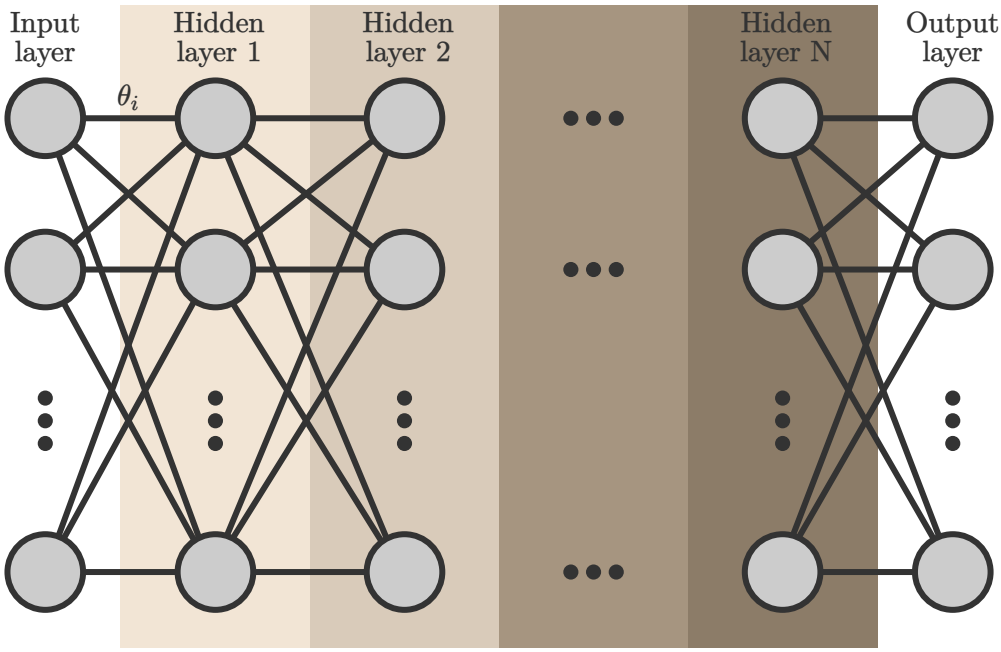


Figure 2.11: Diagram of a fully connected feedforward neural network.

2.5.3 Neural networks for image denoising

Neural networks are flexible representations of functions of many dimensions, which makes them highly parameterised models that can be shaped (trained) for various tasks. Their performance has been unmatched in many domains of image processing, such as image classification where on popular challenges like ImageNet [59] they have been outperforming humans for a while now [60]. Therefore, their use for denoising is a promising direction in the field and an active area of research.

Neural network is a collections of nodes, also called neurons, organised in different ways. The most fundamental type is the feedforward network where the nodes are arranged in at least two layers and information flows in one direction, from the first layer (input) to the last layer (output). Intermediate layers, if any, are called hidden. Usually, each node is connected to every other node from the adjacent layers, with each connection having a certain weight θ_i associated with it. This is illustrated in Figure 2.11. In the case of the denoising problem (equation 2.15), such a feedforward neural network, parameterised by weights $\boldsymbol{\theta}$, could be trained to take a noisy image y and estimate its clean counterpart $\hat{x} = f(\boldsymbol{\theta}, y)$.

A problem with feedforward neural networks is that the number of connections

grows exponentially with the dimensionality of the input. For example, a 10×10 pixel image requires 100 connections for each node in the first hidden layer, and therefore that many weights, too. This amount quickly grows to $10,000$ for an image of size 100×100 pixels. If there is an equal amount of nodes in the hidden layer as in the input layer, namely $10,000$, there are 100 million weights in total. In addition, because of the full connectivity, the network is prone to overfitting, the learned features are not transformationally invariant, and the performance is susceptible to translations and distortions.

Convolutional neural networks (CNNs), originally called neocognitron [61], were designed to address these problems. One or more of their hidden layers are convolutional, whose nodes are called a feature map and they are produced by convolving the previous layer with a filter. The values of the filter are the weights of the layer and they are shared across all nodes of the feature map. Stride is the amount the filter is moved every time it is shifted, and therefore it has a minimum value of 1. Reusing the last example, a 100×100 pixel image would require just 25 weights for a convolutional layer with a single filter of size 5×5, which is 4 million times less than the fully connected layer. A single node in the feature map is only connected to the previous layer nodes that were covered by the filter. This is called the receptive field of the layer and the concept is illustrated for the 1D case in Figure 2.12. As more convolutional layers are used and the depth of the network increases, more sophisticated abstract features can be built and the receptive field grows. In this sense, convolutional neural networks can learn hierarchical features in the data. To give an example with the popular AlexNet [62], the receptive field is 11×11 pixels in the first convolutional layer, meaning that each node in this layer sees only that many pixels of the input, and the receptive field then grows to 163×163 pixels in the fifth convolutional layer.

One of the big disadvantages of CNNs is the fact that spatial information is lost. The output layer is just a vector of nodes. This is fine for classification, but not useful for denoising and many other domains of image processing where it is desired that the network outputs a full image. In addition, downsampling (called pooling) is often used in practice to make the network more robust against transformations and to speed up computation, which reduces the feature size progressively. Even if downsampling is not used directly, the strides of the convolutional layers can be changed to achieve the same

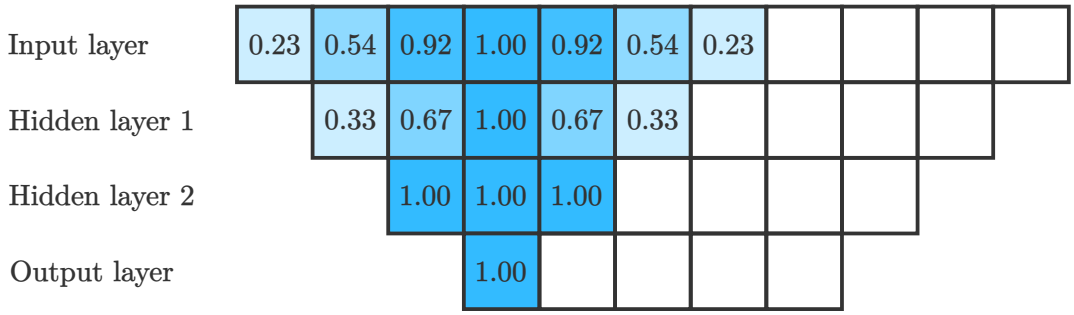


Figure 2.12: Receptive field in a CNN. A single node in the output layer sees 7 nodes from the input layer. Layers 1 and 2 see 5 and 3 nodes, respectively. All convolutions are done with a kernel of size 3, with a stride of 1 and no padding, meaning that the kernel is shifted one pixel at a time without changing its shape. The scores in each node represent the relative contribution within the layer, with a range of $[0; 1]$. Scores of 0 are omitted.

result – contraction with depth [63]. This is the case in Figure 2.12. Using AlexNet again as an example, the feature size in the first layer is 55×55 pixels, whereas in the last layer it is 6×6 pixels. The U-Net architecture [64] addresses this problem of contraction by matching each downsampling operation with an upsampling one, in a consecutive expansion path. The spatial context is preserved via long skip connections, which fuse information from early layers with high spatial content to later layers with high semantic content. It also uses input data augmentation, via rotation and flipping, which reduces 8-fold the required size of the training data. This is important because in many biomedical applications it is challenging or even impossible to obtain the necessary training data.

Neural networks approach the denoising problem differently than traditional methods. Instead of looking for an estimate of the clean image directly, as in equations 2.20 and 2.24, they first look for the parameters of a function that can produce such an estimate, in a learning process called training. The actual denoising is then done in a separate inference step.

Training of neural networks requires some sort of reference (supervisory) signal. The direct approach is to provide this in the form of pairs of clean and noisy images. Mathematically, the objective of the training is to find the weights θ of a network f_θ which minimise a loss function L between the estimate based on the noisy image y and the true clean image x:

$$f_\theta^\star = \underset{f_\theta}{\operatorname{argmin}}\, L(f_\theta(y), x) \tag{2.26}$$

Usually the mean squared error (MSE) is chosen as a loss function because then the optimal network $f_\theta^\star$ learns to produce the mean of all clean images, consistent with the noisy data:

$$f_\theta^\star = \underset{f_\theta}{\operatorname{argmin}}\, \mathbb{E}\left[(f_\theta(y) - x)^2\right] = \mathbb{E}[x \mid y] \tag{2.27}$$

Proof of equation 2.26 is given in Appendix B.1. It is important to note that although the equation implies pixel-wise correspondence between the input and the output, in the case of CNNs this is not entirely correct. Because of the shared weights of filters and because of the receptive field of the output layer, the estimation of the network is based not on just a single pixel but the entire neighbourhood. In this sense, a more explicit formulation would be $\mathbb{E}[x \mid y_{\mathrm{RF}}]$, where y_{RF} are all the pixels in the receptive field. Of course, as illustrated in Figure 2.12 the central pixel is the most important and it has the strongest contribution.

While regularisation is of crucial importance in the case of ill-posed problems, it takes a more subtle role in the case of neural networks. Amongst the many regularisation methods, one worth mentioning specifically is early termination [65], mainly due to its popularity. The idea behind it is to set aside part of the training data, for example 10%, and use it to validate the performance of the network after each training step. The validation error initially decreases, then reaches a stable level, and finally begins to increases again. This is when the network starts to overfit and the training needs to be stopped.

Chapter 3

Developing software for the control of imaging systems

3.1 Introduction

Modern microscopes are complex imaging systems which require precise and robust control, monitoring, and synchronisation of many devices of different kinds, such as cameras, light sources, stages, and filters. Microscope manufacturers usually offer integrated solutions that come with proprietary software which was specifically developed to drive their own instruments. Although suitable for some applications, such software also often has several disadvantages. One of the major issues is that of inflexibility. Making modifications and extending the software is difficult or impossible, usually because it is closed source and proprietary. This also makes it impossible to inspect, preventing users from learning from it or from understanding how it works. Another significant issue is expensiveness.

Commercial microscopes are only one part of the whole. There are many research instruments and imaging systems which are built from separate components. Individual components may have their own software but there is no single cohesive unit and therefore coordination is missing. There are two options in this case – development of custom software or use of third-party solutions.

Developing custom software is the apt approach sometimes, but there are many obstacles that make it difficult to go this way. It is not only a laborious endeavour but requires an advanced level of technical knowledge and experience which is especially challenging in microscopy, where a significant proportion of the users have a more

biological or medical background. Custom software could be seen as wasteful, too, because often times such software remains unavailable to the general public and misses the opportunity to be reused, leading to needless duplication of effort. Of course, there are also many examples of the opposite, where software packages are released to other people and they see continual support and use from people from all over the world. An example of this is ZeroCostDL4Mic [66], which provides a platform for deep-learning that could be used by researchers from all kinds of background, without the need to know any programming at all. This theme of democratisation is important for disseminating technological innovation and for improving the way researchers work, and it is a core element of science.

The need for a unified microscope control software has been recognised before. There have been efforts to develop third-party solutions which are not tied to a particular manufacturer or a specific piece of equipment, and this category of software represents the second option for the control of bespoke microscopes. Initial foundational work in this domain emerged in the early 90s and it was pioneered by David Agard and John Sedat, amongst other contributors [67, 68]. In the modern landscape, Micro-Manager is one of the most well-established tools for driving microscope hardware. It is free, open-source, functionally comprehensive, and it supports a large selection of devices. Despite all this, there are some fundamental limitations to the way in which it interacts with devices and orchestrates experiments [69]. It is generally difficult to use with complex system, including limited support for multiple cameras and inability to synchronise device operation during experiments. Except for its graphical interface, most of its code is developed in C++, which is an advanced language and it introduces an entry barrier for any modifications. In addition, incorporating any such modifications requires the rebuilding of the software, which is an elaborate procedure that requires yet again advanced technical knowledge and experience. A Python extension that builds on top of Micro-Manager was developed [70]. Although it make it easier to interact with the software and to automate some of its functionality, it ultimately does not address its fundamental flaws. Micro-Manager also has limited support for adaptive optics.

An alternative third-party solution is offered by the Cockpit software suite, hereafter referred to as simply Cockpit, which is also the basis of the software developments

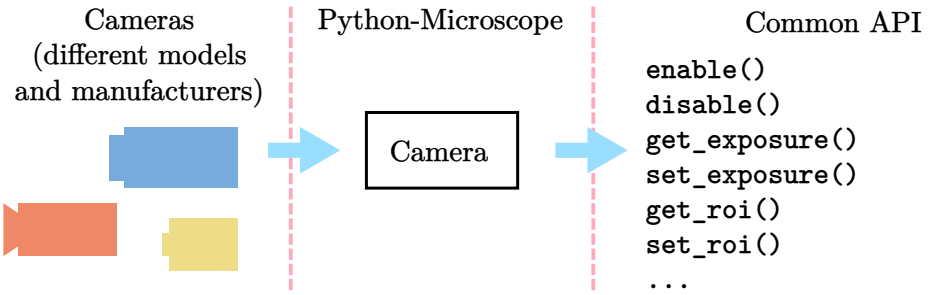


Figure 3.1: Device abstraction in Python-Microscope.

described in this chapter. It was based on the software foundation and philosophy developed in John Sedat's group, especially around the OMX platform [71], and it consists of three Python packages, each taking on different responsibilities. The control of the hardware devices is carried out by Python-Microscope [72], which provides a layer of abstraction such that different physical devices of the same type can be treated as a generic software object with unified interface. This is illustrated in Figure 3.1 for cameras where Python-Microscope provides the abstract `Camera` object which uses the same application programming interface (API) regardless of the underlying hardware device.

The second package is Microscope-Cockpit [73], and it delivers a graphical user interface (GUI), the ability to carry out experiments, and various other high level functionality that builds on top of Python-Microscope, such as coordinated stage control and tiling of images. The third package is Microscope-AOtools [74] and it provides support for adaptive optics by extending both of the other two packages. It add a new abstract adaptive element device, capable of calibration and basic sensorless aberration correction, and it also adds new GUI elements.

One of the main advantages of Cockpit is that it addresses the major limitations of Micro-Manager. All three packages are designed around a client-server interaction where the overall control and the GUI can be split and distributed amongst multiple computers, as along as they have network connection, allowing to work with imaging systems of arbitrary scale and complexity. Moreover, Microscope-Cockpit can use a real-time controller type of device to synchronously control other devices for fast and robust experiments. All three packages are written in Python, creating a better foundation which makes the software suite more accessible and easier to work with, both for users and developers. Coupled with the open-source and community-driven nature of

the packages, this makes them easy to modify and to extend. Another advantage specifically related to the developments described in this chapter is that the three packages were developed locally and therefore they already supported most of the hardware of the bespoke microscopes.

There are many other third-party software tools, but all of them are less suitable. This includes proprietary paid products, such as SlideBook (Intelligent Imaging Innovations), which are not as comprehensive and yet not extendable or modifiable like the open-source alternatives, which makes their application rather limited. Finally, ImSwitch was recently released and it offers similar functionality [75] to Cockpit but it is less mature, supporting fewer hardware devices and having less complete functionality.

3.2 Aims and objectives

The overall aim was to extend Cockpit with the functionality required to do all the experimental work described in this thesis, as well as to improve the existing toolkit with extensive practical testing and with the enhancements that come from the lessons learnt during this experience. The objectives were:

- Add the ability to control adaptive element devices with external electronic signals, and integrate them into the image acquisition framework of Microscope-Cockpit
- Implement remote focusing in a way that it is independent from yet compatible with the already existing aberration correction
- Extend the GUI to provide live feedback and results, such as mode-metric plots during aberration correction, as well as more ways to interrogate and diagnose the performance of the adaptive optics
- Develop a new extended focus image acquisition method, as part of the Microscope-Cockpit's experiment infrastructure, where the remote focusing is used to increase the depth of field

3.3 Results

3.3.1 Functional software changes

Microscope-Cockpit

Microscope-Cockpit is mature and packed with features, and there was not much that needed to be done in order to have all the required functionality. Most of the changes were about extending, instead of creating from scratch, and pertained almost exclusively to the experiments module of the software. The main goal was to accommodate the adaptive optics and particularly the remote focusing.

One of the greatest strengths of Cockpit was the ability to perform fast, robust, and complex experiments. This was difficult to do in normal circumstances on a regular desktop computer running modern general-purpose operating systems such as Microsoft Windows because they were optimised for smooth user experience and did not allow reliable time control of the tasks they ran. This meant that tasks could be interrupted at any time, which led to huge variation in the timings. For example, if messages needed to be sent over USB every 50 ms, they could take anywhere between 50 ms to 80 ms, with occasional delays of up to a second or longer. On shorter time scales, such as microseconds, the timings were even worse, further confounded by the imprecise clock mechanisms of general-purpose operating systems.

To deal with these issues, Cockpit made use of *executor* class of devices which were capable of (pseudo) real-time control. This was usually either a digital signal processor, attached directly to a desktop computer, or an entirely separate embedded system, including single-board computers such as Raspberry Pi. Cockpit performed experiments by first creating a table of actions, specifying what signals needed to be sent to what devices and at what times, and then this table was sent to an executor which carried out the specified actions. In order to guarantee the best response times, these control signals were physical electric signals, as opposed to a software signals such as commands sent over serial port. They were either analogue or digital in nature and they triggered an action in the receiving device, which is why they were referred to as trigger signals. However, in practice only analogue triggers were used to control stages during experiments.

Most microscope Z stages intended for fast and reliable experiments are based on

piezoelectric actuation. They are usually controlled by an analogue voltage, which is mapped to the position of the stage. For example, a travel range of 0 µm to 200 µm is mapped to a voltage range of 0 V to 10 V, such that 0 V correspond to the bottom position, 2.5 V correspond to 50 µm, 5 V correspond to 100 µm, and so on. Because of this, Cockpit was designed under the assumption that the axial change in position during experiments was determined by an analogue signal.

Ignoring the fact that this excluded all digital Z stages, which are not common to begin with, it presented an obstacle for the remote focusing, where it was desired to iterate through a set of control patterns, such as DM shapes or SLM images, by sending digital signals to the adaptive element. As such, Cockpit had to be modified to support digital Z stages during experiments and consequently remote focusing. This opportunity to revise the logic of the experiments was used to also add general support for the triggering of adaptive elements, for other purposes such as aberration correction.

This was implemented as a three-step process, based on concept initially developed by Ian Dobbie, and it also needed changes to the way stages were treated, to distinguish between analogue and digital ones. First, the control patterns were calculated from the parameters set by the user in the GUI during experiment configuration. These patterns were then uploaded onto the adaptive element device. Finally, special provisions were added to the calculation of the action table to ensured that the right type of signal was sent out. This kind of changes were made both to regular experiments and to structured illumination microscopy (SIM) ones. The logic behind the choice of signal is illustrated in Figure 3.2 and it demonstrates how it was possible to independently support remote focusing and generic AO with physical stage, analogue or digital.

The improvements to the experiments module also included the addition of a new type of experiment that enabled extended focus. This was a technique that varied the focus during exposure, resulting in reduced axial spatial resolution but increased depth of focus and temporal resolution. It was implemented only with the remote focusing for start, because of concerns that the response time of the piezo stage was not good enough for fast experiments. The implementation is illustrated in Figure 3.3, which shows an example of the timings that were used in such experiments.

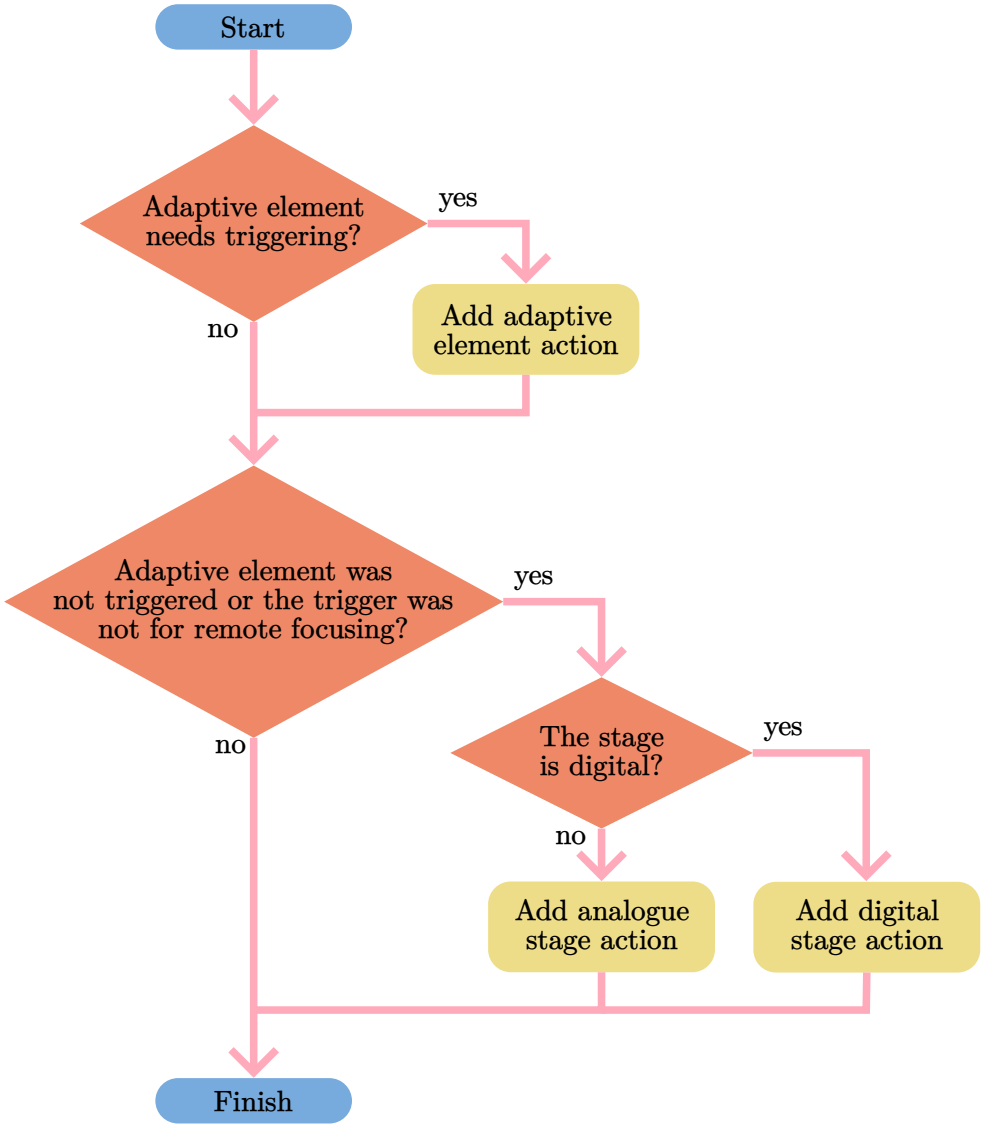


Figure 3.2: Combined logic for the triggering of digital stages and adaptive elements in Microscope-Cockpit experiments.

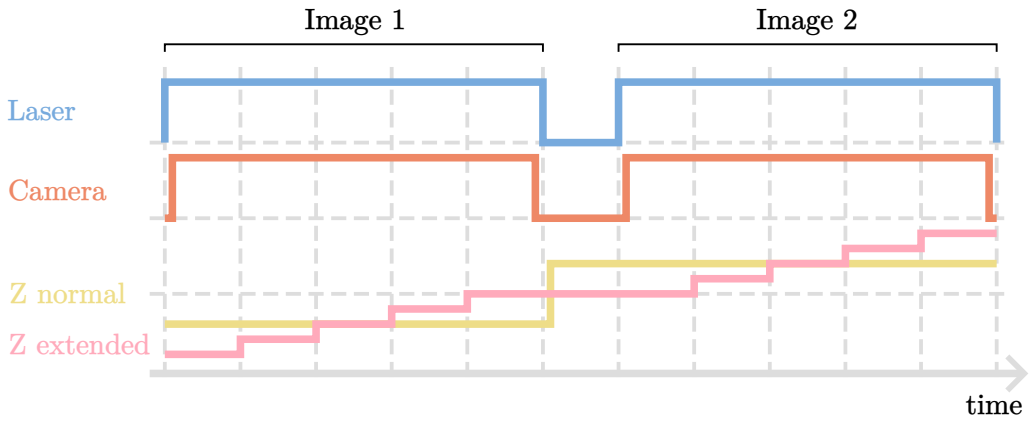


Figure 3.3: Example timing diagram, comparing normal stage (Z) movement with that of extended focus.

Table 3.1: Correction layers.

Layer name	Purpose
mode control	Manual mode specification
remote focus	Remote focusing
sensorless	Sensorless aberration correction
system_flat	Correction of system aberrations

Microscope-AOtools

Microscope-AOtools needed significant changes in order to accommodate remote focusing and to include all the required functionality. Originally, the software had a simple design where the Zernike mode coefficients, derived with the sensorless aberration corrections algorithm, were applied to the adaptive element and this was the extent to which the adaptive optics was controlled. This was impractical and limiting, and had to be reworked completely.

The new system operated with the concept of layering. It allowed to define separate correction layers, each having its own set of Zernike mode coefficients or actuator voltages. The final shape of the adaptive element was then derived as the sum of all of the correction layers. This allowed not only to have a clear separation between different types of correction, but it also allowed to toggle them on demand and to use an arbitrary combination of them (or none). The different aspects of the software had their own individual correction layers, listed in Table 3.1.

The remote focusing was implemented as corrections at two different points. In order to get all intermediate values there was a need to fit them to a function. Section 4.4.1 describes this in detail. From the software perspective, there was a need for a common way of fitting corrections and interpolating between them. The default layers associated with remote focusing and sensorless aberration correction were changed to no longer be a static set of correction values, but to be functions that took a position and mapped it to a set of corrections. For the remote focusing, this position was defined as an offset around the piezo stage and due to this relative nature it remained invariant to stage drift over time. For the sensorless aberration correction, the position was absolute, taking into account all physical Z stages. Due to stage drift this was not a

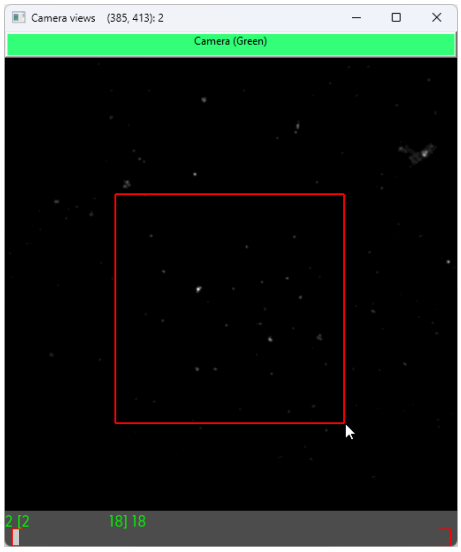
reliable mapping over long periods of time, and therefore it had to be recalculated. If the mapping was defined for a single point, no fitting or interpolation took place and the corrections of the layer remained the same, regardless of the axial position, which effectively bypassed the mapping functionality. This was also the case if there were no points associated with the sensorless aberration correction. In contrast, the remote focusing required at least two points to work.

3.3.2 Improvements to the GUI

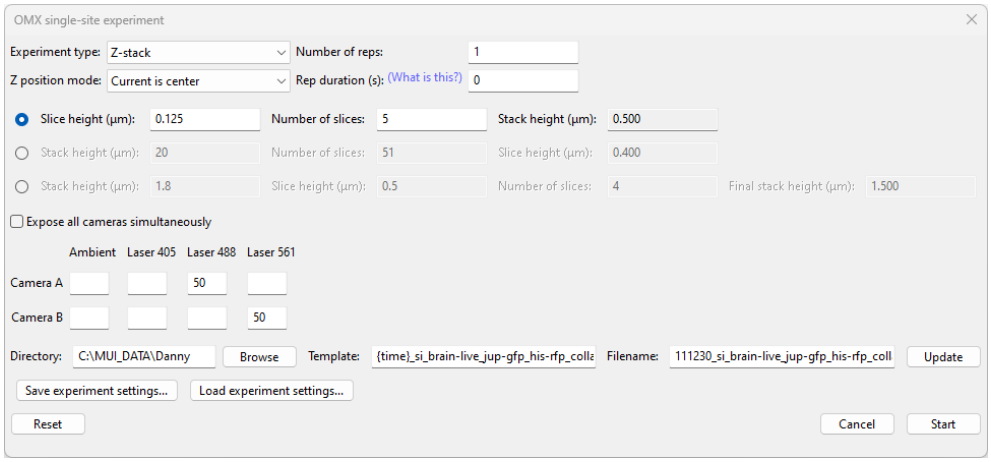
Microscope-Cockpit

The GUI of Microscope-Cockpit had everything necessary to perform the experiments described in this thesis, but two changes were made that notably improved the interaction of the user with the software and thus also the workflow. The first such change concerned the way cameras were configured. A common practice in microscopy is to use a region of interest (ROI) instead of the entire field of view of the imaging system, mainly to increase the acquisition speed, but Microscope-Cockpit had no way to do this visually. It required specifying a tuple of integers (x, y, w, h), where (x, y) was the top right corner of the ROI rectangle, w was the width, and h was the height. These values were difficult to derive just from looking at an image and they took significant amount of time to calculate, which was impractical during experiments, especially with live samples. To get around this issue, the camera image viewer was extended with the ability to draw a rectangle with the mouse, reducing the time it took to specify ROI down to seconds. This change is shown in Figure 3.4a.

In a similar vein, the window for specifying the experimental parameters had several issues related to the way the stage was configured. It only allowed to specify a step size and the total height, without giving any indication about the number of steps. This was not only impractical, but also misleading because with a step of 0.5 µm and total height of 2.2 µm the number of steps were six and so the actual height was 2.5 µm. In order to make this discrepancy explicit and to allow greater flexibility in the specification of the stage parameters, the experiment window was changed to have a choice between three different configurations, as shown in Figure 3.4b. The first configuration allowed specifying the height and the number of slices, and it then showed the total height. The second configuration specifying the total height and the number of slices, and it showed



(a) Interactive ROI selection.



(b) Revised experimental settings.

Figure 3.4: Changes to the Microscope-Cockpit GUI.

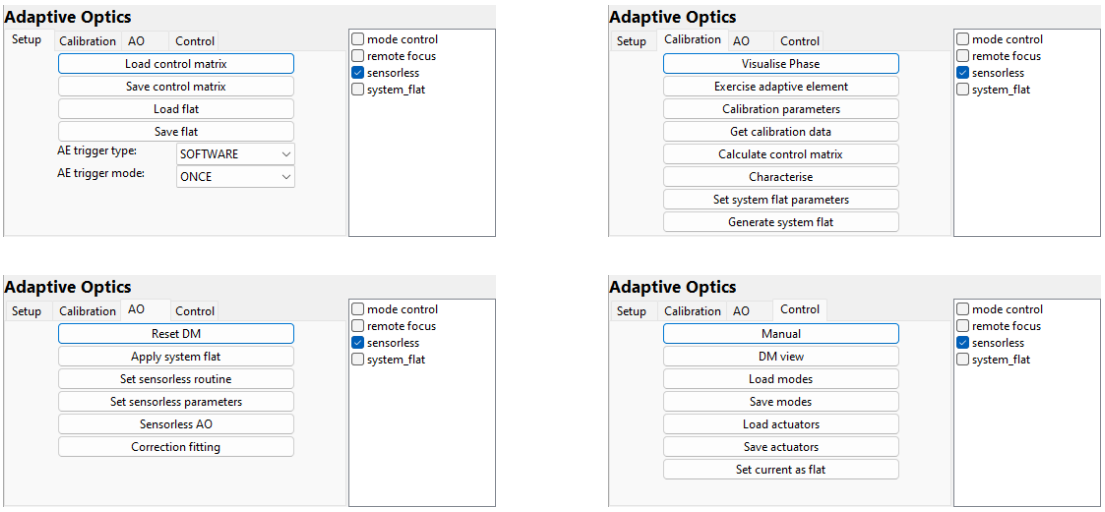


Figure 3.5: Revised main panel. The four images show the different tabs.

what the slice height would be. Finally, the third configuration allowed to specify all three parameters, but it then showed what the actual total height would be.

Microscope-AOtools

The software was designed with the one-button-does-all concept, aiming to streamline the use of adaptive optics for biologists and other non-technical users by doing all the work behind the scenes. Unfortunately, this did not work well in practice because there was no way to tell if or what went wrong when something inevitably did. This was not just from developer perspective, as users were unable to tell, for example, how well their selected metric performed, nor even the Zernike mode coefficients resulting from the correction. Therefore, significant efforts were spent on adding diagnostics to the software, as well as other ways of interaction. The GUI reflects these changes the most. Figures 3.5 to 3.10 shows updated interface.

3.4 Discussion and conclusions

All the software developments described in this chapter are open-source and publicly available on the GitHub online platform. They had to be implemented in an accelerated manner, in order to meet project time frames, and for this reason they are still in the process of being integrated into the main software branches. They are located at https://github.com/dstoychev/cockpit/tree/deepsim and https://github.

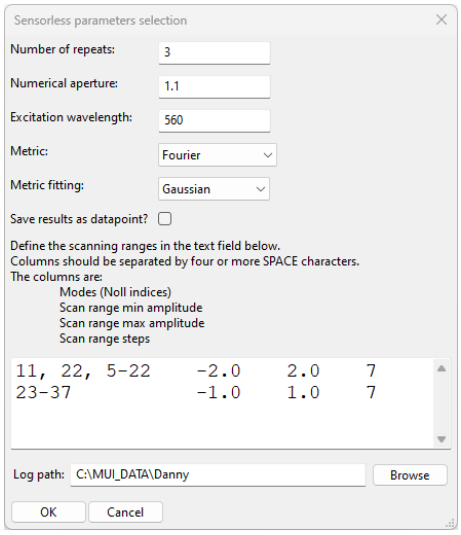


Figure 3.6: Updated sensorless aberration correction setup window.

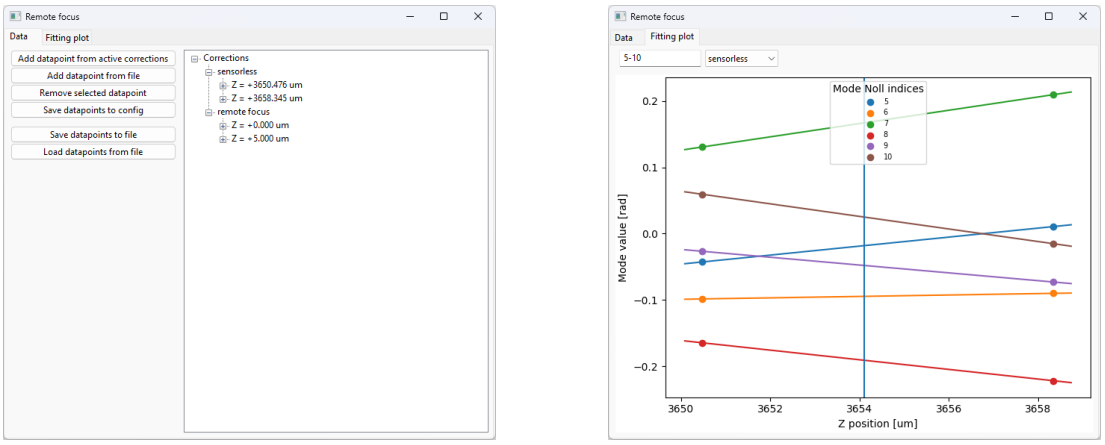


Figure 3.7: New correction fitting window.

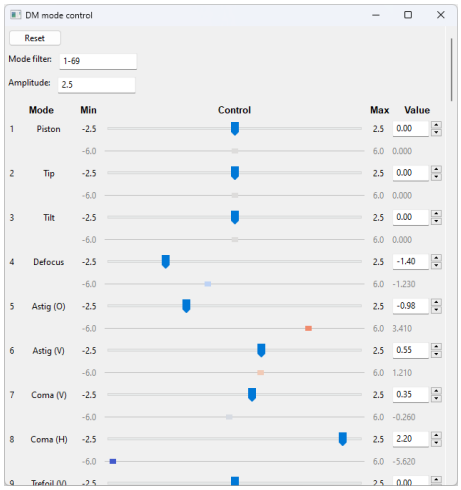


Figure 3.8: New manual mode setting window.

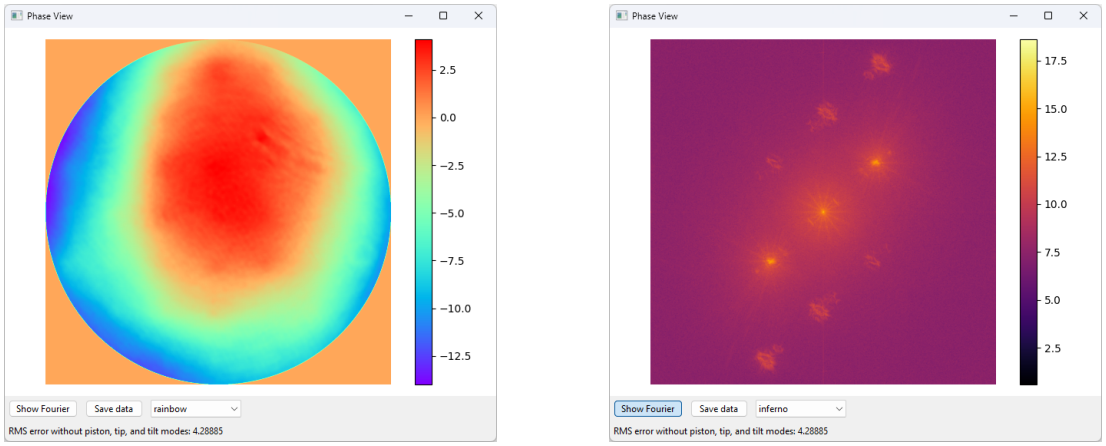


Figure 3.9: Revised phase visualisation window.

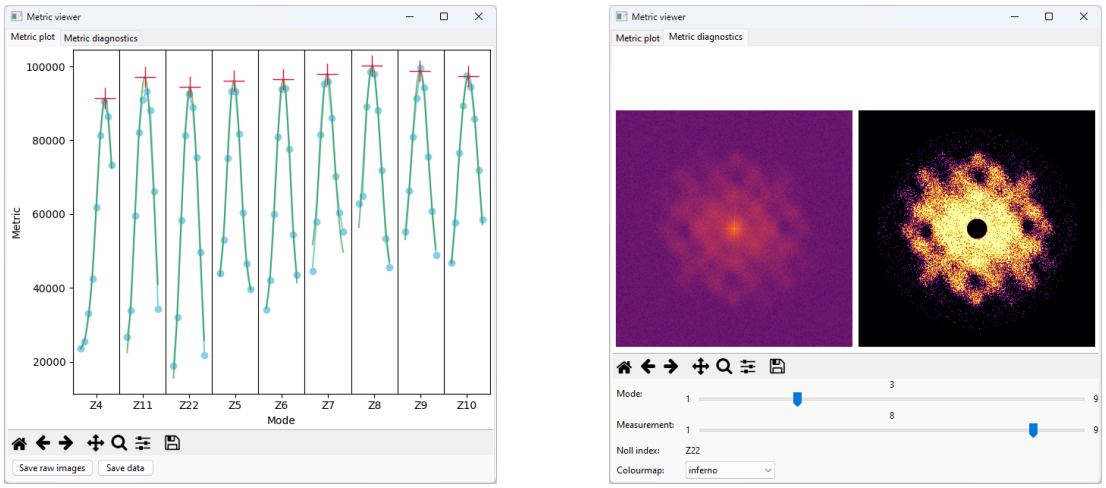


Figure 3.10: New aberration correction results and diagnostics window.

com/dstoychev/microscope-aotools/tree/deepsim-matthew-danny for Microscope-Cockpit and Microscope-AOtools, respectively. The two also need a custom version of the Python-Microscope package, which contained minor bugfixes related to a specific DM device and were not described in this chapter. Appendix A.1 contains a description of the bug.

Chapter 4

Adaptive optics for live fluorescence imaging

4.1 Introduction

All forms of light microscopy produce images which are optically aberrated to some degree. This could distort significantly the results of quantitative imaging and therefore it should be minimised. One way to compensate some of these aberrations is with adaptive optics where a special device, such as a DM or a SLM, is used to change the shape of the light.

Adaptive optics traces its origin to the fields of astronomy and military engineering, starting with the ideas of Horace Babcock in the middle of last century, which were then refined and put into practice in the following decades [76]. Application to microscopy emerged a bit later, around the turn of the new millennium [77]. Depending on whether a wavefront sensor is used or not, aberration correction methods are either direct or indirect. Section 2.4.5 outlines some of the most common approaches.

Direct wavefront sensing methods are fast and straightforward to use, but also expensive and complex to implement, and they only work well with samples that do not scatter the light too much, such as thin, transparent, or other weakly scattering samples [78]. Unfortunately, this limits severely the repertoire of biomedical samples. Indirect wavefront sensorless methods are preferred in these cases. They have been demonstrated in broad range of microscope modalities, including widefield [79], confocal [80], multiphoton [81], and super-resolution [82] microscopy.

In addition to the correction of aberrations, adaptive optics can endow an op-

tical imaging system with remote focusing, which is the ability to axially scan the sample while keeping it stationary. The advantages are speed and the elimination of the disturbances to the sample caused by traditional mechanical actuation. The first application of this technology to microscopy was achieved with the aid of a second objective lens [43] and additional optical components, as discussed in section 2.4.7. This arrangement, however, not only increases the complexity of the imaging system but it also significantly reduces the optical efficiency, due to increased amount of optical surfaces in the light path.

Alternatively, the focus can be varied with an adaptive element, first demonstrated in a widefield microscope with a DM and direct wavefront sensing [83]. This type of remote focusing has been implemented in various microscopes [84–86], including for particle tracking [87] which is closely related to the study of the dynamics of mRNA transport. Most such applications do not perform independent aberration correction or they employ devices such as electrically tuneable lenses and tuneable acoustic gradient lenses, which have much more limited dynamic range and spatial resolution than other types of adaptive elements. The combination of wavefront sensorless aberration correction and remote focusing, with a DM or SLM, has only been demonstrated in multiphoton microscopes [88, 89].

Despite its many advantages, adaptive optics in microscopy has not been able to achieve great widespread presence and commercial adoption, partly because of the complexity involved and the inconveniences it presents to biologists, to other researchers, and to clinicians [90, 91]. Recently, a generalised sensorless adaptive optics framework was developed to address some of these concerns [74, 92]. Although it was a step in the right direction it did not achieve the accessibility, completeness, and flexibility necessary to be easily and reliably deployed.

In this chapter I describe the technical work associated with the development of a super-resolution upright imaging platform, intended for fast live imaging. It builds on the ideas of accessible and democratised adaptive optics, it evaluates its performance in several different applications, and it demonstrates novel results in both aberration correction and remote focusing. The software aspect of this work is described in chapter 3.

4.2 Aims and objectives

The aim was to develop a novel comprehensive adaptive optics system, capable of both aberration correction and remote focusing. The idea was to understand how much benefit adaptive optics can bring to the study of the mechanisms of RNA transport, what are its strengths and weaknesses, and how it can be used to improve the state of the art. It was an attempt to further bridge the gap between engineering and biology, and to develop the methodology to use adaptive optics in a robust and reliable way, which is technically challenging. The emphasis was on fast live fluorescence imaging, which necessitated the use of widefield microscopy because of its speed and gentleness. The independent application of aberration correction and remote focusing with an adaptive element had not been shown before in such an instrument.

Building a microscope is a laborious and time-consuming task, which is certainly outside the scope of this work. Therefore, to make this project feasible the adaptive optics were implemented in a pre-existing super-resolution imaging platform, capable of both widefield and structured illumination microscopy. This work was carried out in collaboration with Dr Jingyu Wang from the departments of biochemistry and engineering, amongst other contributors. The combination of remote focusing and SIM had not been shown before either, and presented another exciting research direction. Moreover, imaging very deep with such an instrument, confounded partially by optical aberrations, was yet another previously not well explored area and thus an excellent application for the adaptive optics that can push the boundaries of the field, although similar work was published in the meantime [93]. The objectives were then as follows:

- Implement remote focusing separately and independently from aberration correction, in a way where it could be used as a direct replacement of traditional mechanical axial actuation, such as with a piezo stage
- Evaluate the suitability (e.g. speed and light exposure) and impact of aberration correction and remote focusing for live fluorescence imaging
- Investigate how far can aberration correction push the depth limit of 3D-SIM super-resolution imaging, and whether remote focusing can still function well at such depths

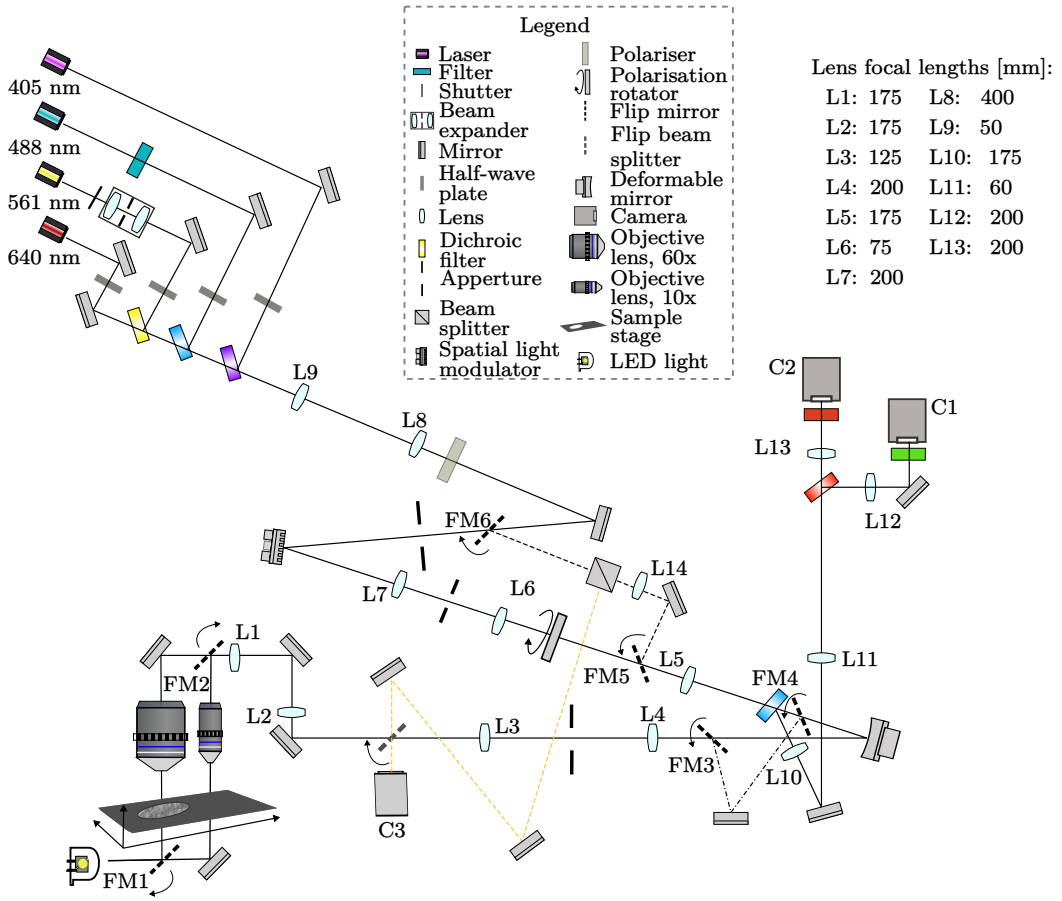


Figure 4.1: Optical diagram of the system in which the adaptive optics were developed and then evaluated. Based on a drawing by Jingyu Wang.

4.3 Methods and materials

4.3.1 Microscope design

The optical diagram of the system is shown in Figure 4.1. The light for reflected fluorescence imaging was delivered by a custom laser combiner, shown in the top left corner, which included four standard laser lines: 405 nm, 488 nm, 561 nm, and 640 nm. Transmitted illumination, used for brightfield imaging, was delivered by a white LED from below the sample stage, shown in the bottom left corner. Structured illumination was created with an SLM (Boulder Nonlinear Systems, HSP512-488-800), conjugated to the image plane. Half-wave plates were used to condition the light and to polarise it linearly. Polarisation rotator (Meadowlarks Optics, LPR-100-λ) was used to change the angle of the structured illumination pattern during SIM experiments.

A DM (Alpao DM69-15) was used for the adaptive optics, conjugated to the pupil plane and used to correct both the excitation and emission paths. The system was

design mainly around a single objective lens, the Olympus LUMFLN60XW 1.1NA, but it also included a secondary low-magnification objective, Olympus LMPLFLN10X 0.25NA, which was important for many of the biological experiments. The main objective lens was selected to be the water-based objective lens with the highest NA offered by this manufacturer, which can be used both in a dipping configuration and with a cover slip using water immersion. The choice was dictated by the application of the system – live fluorescence imaging, particularly related neurobiology. The change from water immersion to water dipping was done via the adjustment collar on the barrel of the objective lens. This collar could also be used for compensation of spherical aberration to some level, but this is inferior to adaptive optics and it even poses the danger of making the image quality worse [94]. For this reason, the collar was mostly unused, besides ensuring it is in the right position for the way the objective lens was used.

The sample mount consisted of three mechanical stages: a heavy-duty linear stage (Aerotech PRO115) used for coarse Z translation, and a combination of two piezo stages – Physik Instrumente M-687 and P-736 – providing fine fast motorised control in all three dimensions. The emission and transmitted light were collected by two parallel cameras (Andor iXon Ultra 897 EMCCD), capable of simultaneous imaging of green and red fluorescence. This was another choice dictated by the application. In theory, the system could be extended with more cameras covering all four laser lines or different wavelengths entirely. The only limits are the availability of suitable dichroic filters and working wavelength range of the SLM.

Finally, there were three pairs of flip mirrors used to make alterations to the light path. The first pair, FM1 and FM2, was used to change between the two objective lenses. The second pair, FM3 and FM4, was used to switch between the DM and a regular planar mirror. This alternative path was used mainly for calibration and development purposes, and it is shown with a black dash-dotted line in Figure 4.1. It was referred to as “bypass” mode because it bypassed the adaptive optics. The third pair, FM5 and FM6, was used to take out the SLM. Again, this was used for calibration and it included extra optics for interferometry. This path is shown with black dashed line in Figure 4.1, with the reference path of the interferometer shown with a yellow dashed line. The principle of the interferometer was the same as the Twyman-Green layout shown in Figure 2.4, with the difference that two beam splitters

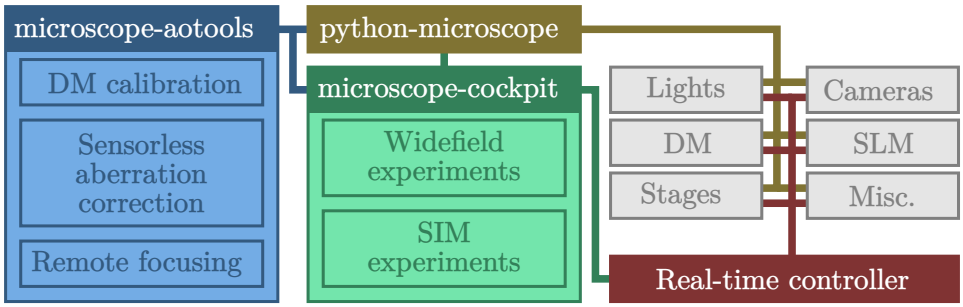


Figure 4.2: Control configuration of the imaging system.

were used instead of one.

The optical system was controlled with a software stack consisting of three key pieces of software, developed as Python packages: python-microscope, microscope-cockpit, and microscope-aotools. Their interaction is shown in Figure 4.2. In summary, python-microscope provides the means to communicate with the hardware, microscope-cockpit provides GUI and experiment functionality, and microscope-aotools uses both of the aforementioned packages to provide support for adaptive optics. In order to facilitate fast, accurate, and complex experiments, microscope-cockpit uses one or more external timing devices, which are capable of coordinating other devices by means of electrical signals, analogue or digital, as opposed to software signals which are usually slow and unreliable. This ability to respond to external electrical control signals is referred to as hardware triggering, as it usually results in the initiation of a one-off action, such as camera exposure or stage displacement. The optical system employed one such timing device (RedPitaya 125-14 v1.0). More information about the structure of the software packages and their interaction is given in chapter 3.

4.3.2 Control of Alpao DM69-15 deformable mirror

The Alpao DM69-15 works by electromagnetic actuation and it was chosen for its large dynamic range, high speed, and external triggering. The large stroke of the device also made it particularly suitable for certain applications such as remote focusing. It consists of 69 actuators distributed in a circular shape, as shown in Figure 4.3.

During operation, we discovered that the device exhibits unusual and unreliable behaviour, which can be described by two phenomena: drift and temperature dependence. The first has been reported in the literature before and it has been attributed to the properties of the polymer materials used to attach the actuators to the mirror

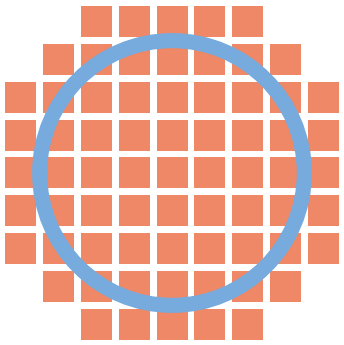


Figure 4.3: Actuator layout of the Alpao DM69-15. The actuators (orange) are spaced 1.5 mm apart and the aperture (light blue) has diameter of 10.5 mm.

membrane [95]. The temperature dependence has been observed before too, and it relates to membrane deformations caused by the raise in temperature from the high currents used in the electromagnetic coils [96].

Two experiments were performed to understand better the dynamic behaviour of the device, inspired by the previous work on characterising similar effects [95, 96]. The first test involved monitoring the change in phase, and hence mirror surface, over a long period of time. The DM was first flattened using the procedure described in section 4.3.5, producing a set of normalised actuator control voltages S in the range $[0, 1]$. The goal was then to observe how this shape is going to evolve over time, for a total duration of 16 hours. The experiment was divided into 1-minute cycles. For the first 50 seconds of each cycle, random offsets were added to the control signals and applied to the device with a frequency of several hundred hertz, to simulate normal operation of the device. These offsets were uniformly distributed in the range ± 0.05, so on average the actuator pattern remained the same. After this initial period of 50 seconds, the device was set to the original flat shape S, without any actuator offsets, for the remainder of the cycle. Interferograms were captured at the end of each cycle. These interferograms were later processed to obtain phase maps for each of the 1 minute marks. This data is presented in Figure 4.4. It shows a rapid increase in the phase error within the first few minutes, followed by more gradual increase which shows an upward trend even after 16 hours.

A second experiment was performed to distinguish between the drift and temperature effects. This time the device was warmed up and prepared by holding a specific shape for 4 hours. Afterwards, six different shapes were applied and held, while monitored, for 20 minutes each. Finally, an eighth shape was held and monitored for 60

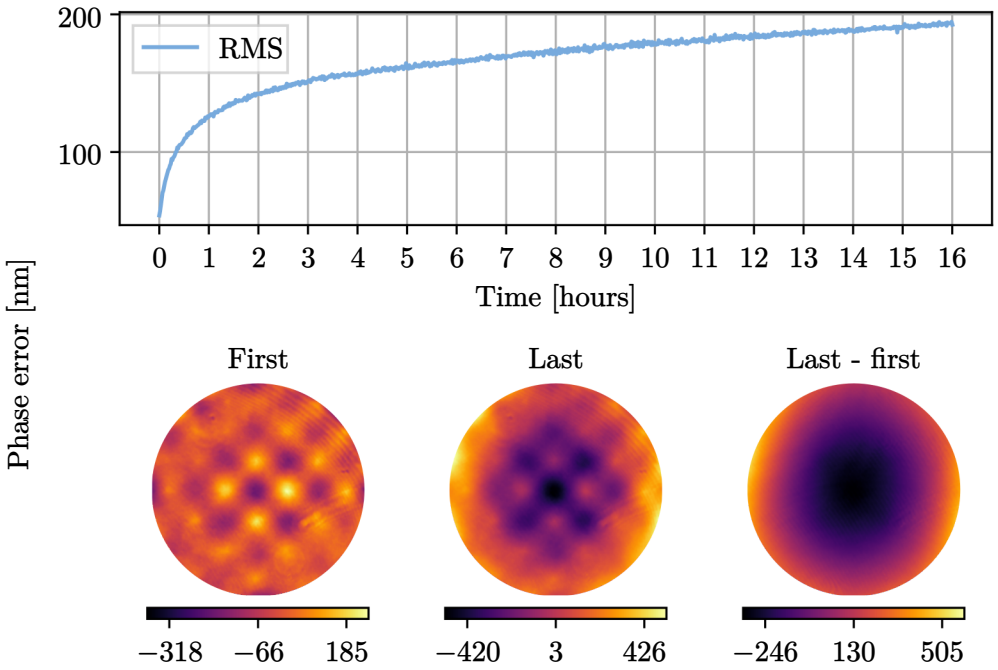


Figure 4.4: Long-term stability test of the Alpao DM69-15. The top row shows the RMS phase error as function of time. The bottom row shows the 2D phase error maps for the first and last timepoints, as well as their difference.

minutes. Unlike the first experiment, the DM was not used during the monitoring periods. In other words, once the device was set to a specific shape it was undisturbed for the remainder of the monitoring period. All of the shapes were individual Zernike modes with coefficients of 1. They are listed in Table 4.1, together with the time for which they were held. The results are shown in Figure 4.5. Besides some spurious measurement errors, the device had reached steady state by the first measurement at $t = 1\text{min}$. This confirmed that the behaviour observed in the first experiment was largely due to the temperature dependence, whereas the drift had a dominant short-term effect, working in the order of seconds. The bias of the drift was determined by the average of the last few shapes of the mirror. For example, if the surface was kept flat and then changed to parabolic shape such as defocus, the final shape at the 1 minute mark would be pulled back towards flatness.

As a result of all of these observation, the DM was only used after it had been properly warmed-up first. It involved thorough exercising of the device at least every day that it was used, certainly before each imaging session. In addition, the device was exercised before calibration, otherwise the control deteriorated severely with the increase of temperature. The exercise routines should use large amplitude, spanning

Table 4.1: Shapes used in the second experiment to test the dynamic behaviour of Alpao DM69-15.

Shape	Zernike mode (Noll index)	Holding time [min]
S_0	Z_4 (defocus)	240
S_1	Z_5 (primary oblique astigmatism)	20
S_2	Z_6 (primary vertical astigmatism)	20
S_3	Z_7 (primary vertical coma)	20
S_4	Z_8 (primary horizontal coma)	20
S_5	Z_{11} (primary spherical)	20
S_6	Z_{22} (secondary spherical)	60

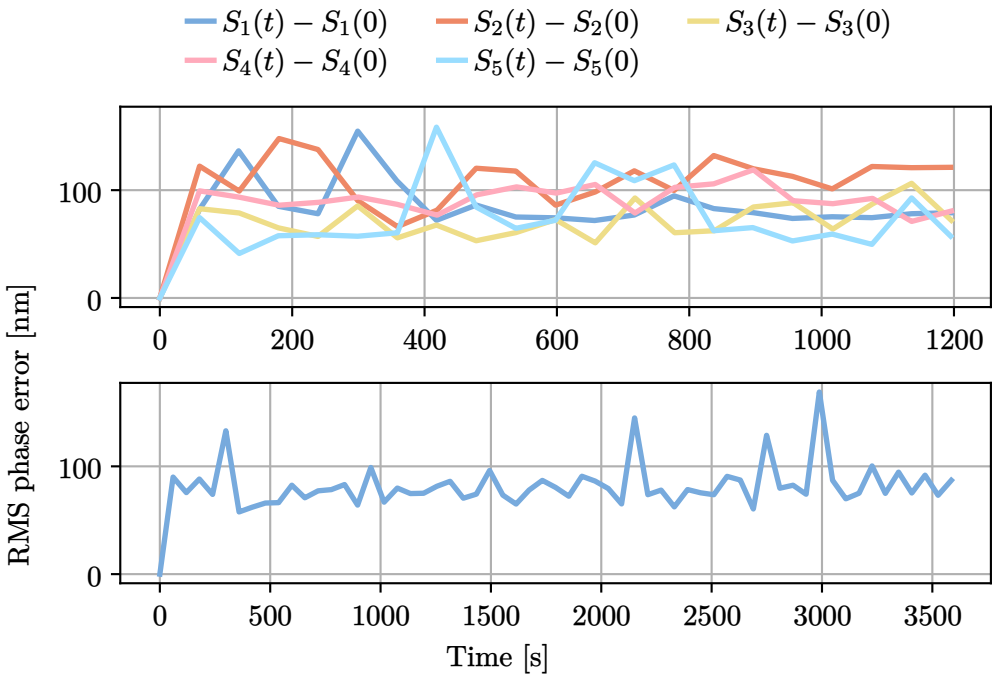
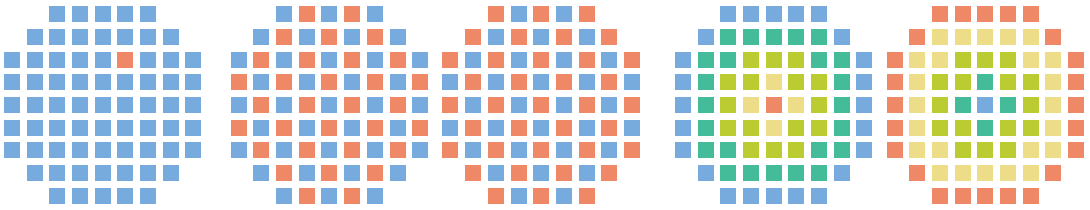


Figure 4.5: Short-term stability tests of the Alpao DM69-15. The top row shows the monitoring of individual shapes S_1 to S_5 for 20 minutes, while the bottom row shows S_6 for 60 minutes.



(a) Poking. (b) Checkerboard patterns. (c) Defocus.

Figure 4.6: Three different types of DM exercise patterns.

a significant proportion of the actuator control voltage range, and it should cover all actuators. To be on the safe side and to diversify the used shapes, often a combination of different exercises were performed. Figure 4.6 illustrates three different exercise routines which were often used. The first exercise was poking of each actuator individually, where the control voltage was set to several different values spanning 80% of the total range and held for 2-3 seconds at each step. This was usually repeated once with 21 steps or twice with 11 steps. The second exercise was an alternating binary checkerboard patterns, holding each for 3 seconds, for a total of 50 repetitions. The two values of the patterns usually spanned a range of 80%, meaning one of them was 0.1 and the other 0.9 of normalised actuator control voltages. Finally, the third exercise was varying defocus, Zernike mode Z_4, with coefficients in the range $[-5; 5]$, and typically in 11 steps from one end to the other. Using random actuator voltages as an exercise was an option, too, but it needed to be done carefully and gently because having too large discrepancies in control signals between neighbouring actuators could have negative effects on the device. A range of 10 to 30% and holding each pattern for several seconds were sensible parameters for this kind of exercise.

4.3.3 Calibration of the DM

Even though the system was designed to perform sensorless aberration correction, the wavefront created by the DM was still sensed for calibration purposes. This was achieved with an integrated interferometer, as explained in section 4.3.1. The optical arrangement is shown in Figure 4.1. The excitation light was redirected by FM6 and then separated into reference (orange dashed line) and measurement light paths. The latter was reflected by the DM and then recombined with the reference at the flip beam splitter (50:50 pellicle) between lenses L2 and L3. The resulting interference fringes were then detected by a CMOS camera (Ximea xiQ MQ042MG-CM) C3. The phase

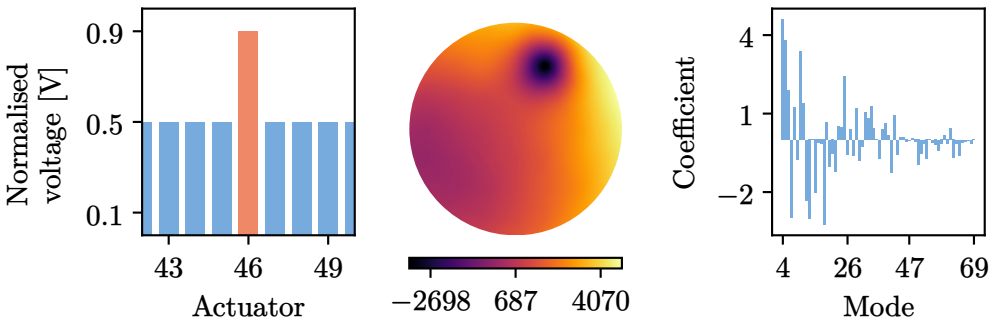


Figure 4.7: Probing the influence of actuator 46 of Alpao DM69-15. The units of the phase map are nm. The modes are given with Noll indexing.

was recovered from a single intensity image by using the Fourier-based algorithm outlined in section 2.4.4. The phase was then decomposed to a set of Zernike modes using numerical integration, as described by equation 2.8.

This method of wavefront detection was used for calibration, following standard procedures [97, 98]. Pushing or pulling of a single actuator deformed part of the surface that spanned more than just the area above the actuator. The impact of each actuator on the shape of the surface was described by a set of influence functions. To measure these functions, each actuator was set to several different control voltage levels, while keeping the rest in their neutral position, and the resulting phases were decomposed to Zernike coefficients, as shown in Figure 4.7. The influence functions were then derived by linear fitting of actuator voltages and Zernike coefficients. An example of such fitting is shown for Z_5 in Figure 4.8. The influence functions were aggregated in an influence matrix $\mathbf{\Omega}$ of size N by M, where N was the number of Zernike modes used in the decomposition and M was the number of actuators. The number of modes was kept at 69 to make the matrix square, which is not strictly necessary but it is neat and convenient. Each element of this matrix corresponded to different influence function (line slopes) and an example matrix is shown in Figure 4.9 where the symmetry around the centre of the DM surface (actuator 35) is readily observed.

The influence matrix related actuator voltages κ to Zernike mode coefficients a, such that $\mathbf{a} = \mathbf{\Omega}\kappa$. The control of the device, however, required the inverse – to get κ, given $\mathbf{a}$. This was achieved by creating a control matrix $\mathbf{C} = \mathbf{\Omega}^{+}$, where the pseudo-inverse was used to ensure the existence of a unique solution. It was constructed via singular value decomposition, where $\mathbf{\Omega} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^{\top}$ and $\mathbf{C} = \mathbf{V}\mathbf{\Sigma}^{-1}\mathbf{U}^{\top}$. The singular

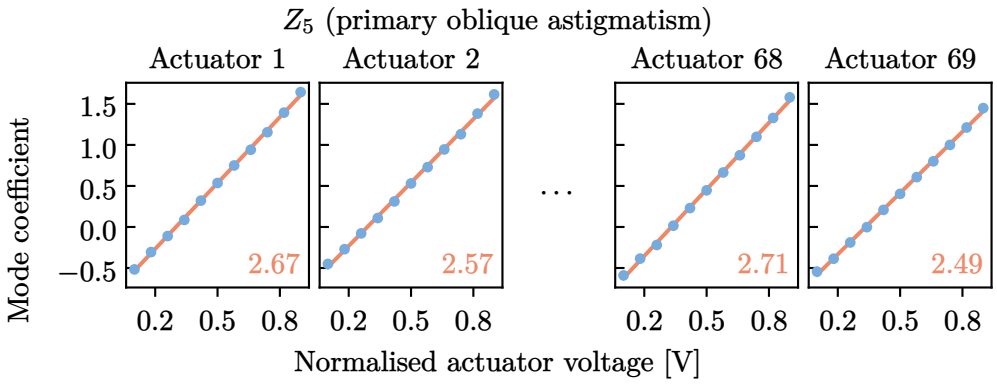


Figure 4.8: Influence functions for Zernike mode 5. The corresponding line slopes are shown in the bottom right corner of each plot.

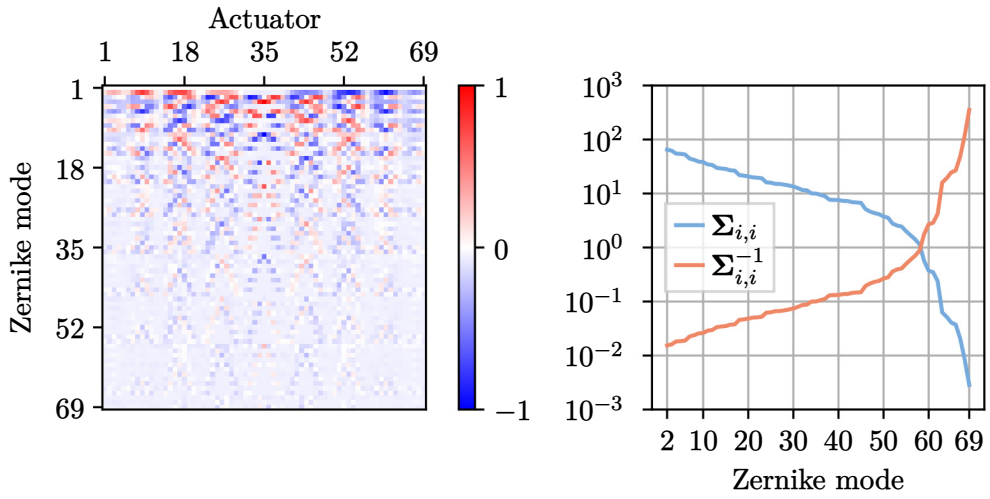


Figure 4.9: Influence function matrix (left) and its singular values (right).

values $\mathbf{\Sigma}_{i,i}$ also carried meaningful information about the control of the device, as they described the changes in Zernike mode coefficients as a result of changes in actuator voltages. Naturally, $\mathbf{\Sigma}^{-1}_{i,i}$ corresponded to the inverse ratios, but they represented the same thing. The two are plotted in Figure 4.9, showing how higher order Zernike modes required increasingly large actuator voltages. The two graphs crossed the 1.0 line around Z_{58}, followed by steep changes, which demonstrated that the device may struggle to reproduce well modes beyond this threshold.

The accuracy of reproducing Zernike modes was further evaluated with a characterisation assay. Each mode was individually applied to the DM with a coefficient of 1.0. The resulting phase was recorded with the interferometer and decomposed into Zernike modes. The errors between the applied and the measured Zernike modes were collected into a matrix, as shown in Figure 4.10. The diagonal values showed generally

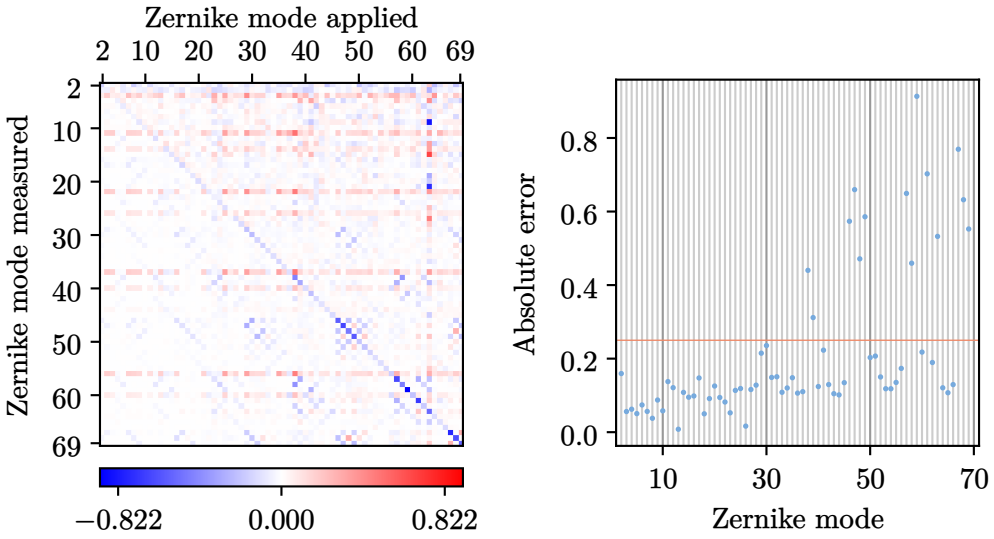


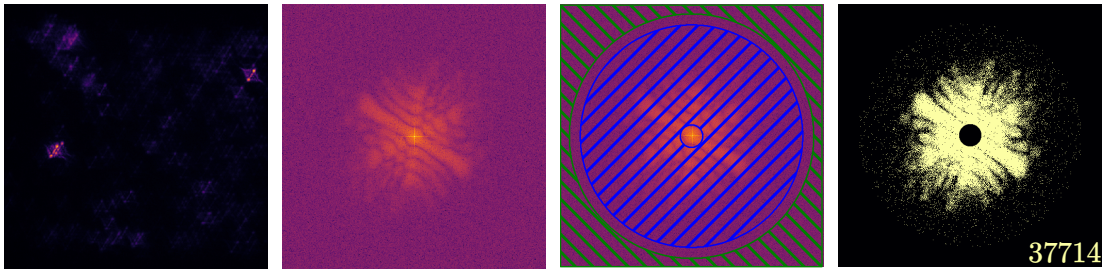
Figure 4.10: Assessment of the calibration quality. The ability to reproduce each mode of the calibration matrix was tested and the results gathered in a matrix of errors (left). Each column is a separate test and each row is decomposition of modes for that test. The diagonal entries of the matrix are shown on the right.

good reproduction of most modes, with most high errors spread around the higher order modes. Besides noise and measurement errors, all the other entries in the matrix corresponded to the coupling between modes, which is particularly noticeable in the rows for defocus (Z_4) and spherical aberrations ($Z_{11}, Z_{22}, Z_{37}, \dots$). The absolute diagonal values are also shown in Figure 4.10, where an arbitrary threshold of 0.25 was used to discriminate between modes that can or cannot be reproduced.

4.3.4 Sensorless aberration correction

The aberration correction was carried out in a wavefront-sensorless way, as described in section 2.4.5. The scanning usually covered Zernike modes 4 to 22 (incl.) using Noll indexing, with coefficient ranges from ± 1.0 to ± 3.0. Higher order modes did not seem to bring noticeable improvement in the image quality, and thus they were ignored for the most part. One round of correction was often enough, but two or more rounds were performed whenever possible, to ensure that the optimum correction was reached.

The metric which determined the image quality in a quantitative manner was based on Fourier analysis. Its calculation consisted of several different steps, as illustrated in Figure 4.11. An intensity image $I(x, y)$ was transformed to frequency space via Fourier transform and then converted to logarithmic power spectrum $S(u, v)$, with $u \in U$ and



(a) Intensity image. (b) Power spectrum. (c) Filtered regions. (d) Metric pixels.

Figure 4.11: Fourier metric calculation algorithm. Panel (a) shows an intensity image $I(x, y)$ that suffers from strong astigmatism. Panel (b) shows the logarithmic power spectrum of the Fourier transform of (a). Panel (c) shows the high-pass filtering (green hatched area) used to isolate the high-frequency noise and the band-pass filtering (blue hatched area) used to isolate the signal. Panel (d) shows the final set of pixels used to calculate the metric, listed in the bottom right corner.

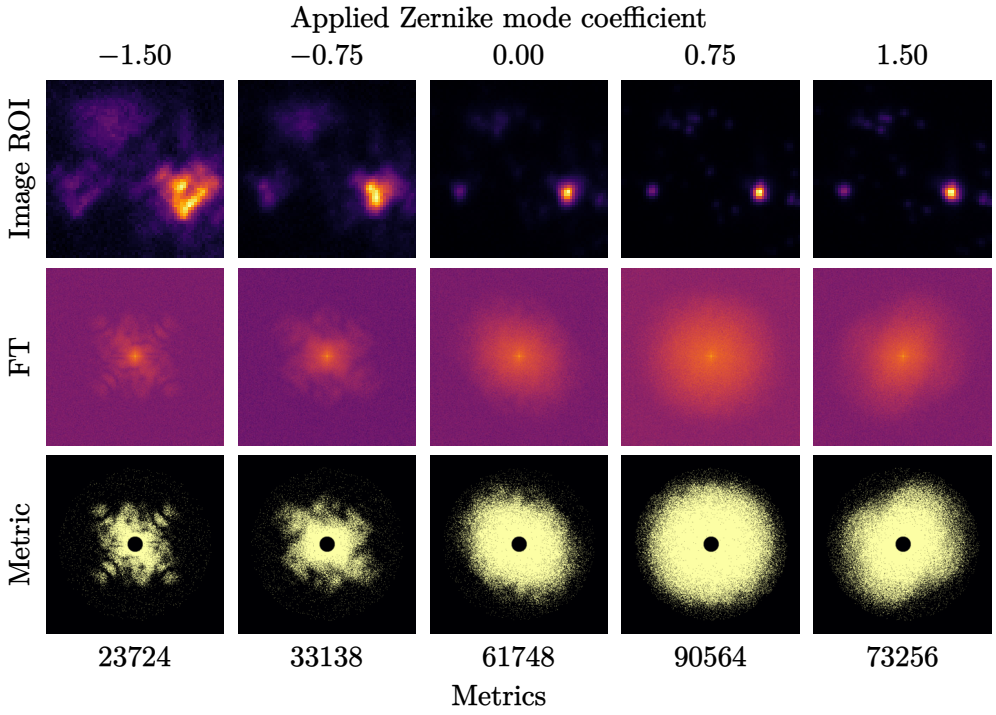


Figure 4.12: Evolution of the Fourier metric during scanning of defocus. The top row shows a cropped region of the intensity image $I(x, y)$. The middle row shows the logarithmic power spectrum of the Fourier transform of the image. The third row shows the final set of pixels used to calculate the metric values, which are listed underneath.

$v \in V$ being the spatial frequencies in each dimension:

$$S(u,v) = \ln\left(\left|\mathcal{F}\{I(x,y)\}\right|^2\right) \tag{4.1}$$

The resolution r and pixel size p of the imaging system were then used to filter the power spectrum in two different ways, first with a high-pass and then with a band-pass filters. The spatial cut-off frequency of the imaging system ω_c was calculated as the ratio $1/(r\omega_{\max}) = 2p/r$ times the maximum spatial frequency in the image, which for square images was $U = V$. The two filters F_H and F_B were then defined as:

$$F_H = \begin{cases} 1 & \text{when } \|(u,v)\|_2 > 1.1\omega_c \\ 0 & \text{otherwise} \end{cases} \tag{4.2}$$

$$F_B = \begin{cases} 1 & \text{when } \omega_c > \|(u,v)\|_2 > 0.1\omega_c \\ 0 & \text{otherwise} \end{cases} \tag{4.3}$$

The high-pass filter was used to separate the high-frequency noise in the image, which skews the calculation of the metric. This high-frequency noise was the subset of pixels $S_n(u,v) = S(u,v) \cdot F_H$, shown as green hatched area in Figure 4.11c. This was further used to define noise threshold as the mean of all pixels:

$$\bar{S_n} = \frac{1}{UV}\sum_{u}^{U}\sum_{v}^{V} S_n(u,v) \tag{4.4}$$

In a similar way, the band-pass filter was used to isolate the signal content, while blocking out the lowest frequencies because they contained background fluorescence and general unspecific fuzziness which were not part of the desired signal. This defined another subset, $S_s(u,v) = S(u,v) \cdot F_B$, which is shown with blue hatched area in Figure 4.11c. It was further filtered by the noise threshold to yield the final set of pixels associated with the metric, shown in Figure 4.11d:

$$S_m(u,v) = \begin{cases} 1 & \text{when } S_s(u,v) > \bar{S_n} \\ 0 & \text{otherwise} \end{cases} \tag{4.5}$$

The numerical value of the metric, shown in the bottom right corner of Figure 4.11d,

is simply the amount of non-zero pixels in this subset: $|\{s \mid s \in S_m(u,v) \text{ and } s = 1\}|$. In essence, the metric quantified the spatial frequency content of the signal and it was at its maximum in the absence of aberrations. Figure 4.12 shows an example of how the metric changed when a Zernike mode was scanned during the sensorless aberration correction routine.

The metric values and the corresponding Zernike mode coefficients were fitted to a parameterised Gaussian function with offset o, amplitude a, centre (mean) c, and width (standard deviation) w:

$$(o - a) + ae^{-\frac{(x-c)^2}{2w^2}} \tag{4.6}$$

In the rare situations when the fitting failed there were two mechanisms to still provide a sensible value, either by looking for sharp peaks in the metric values or by defaulting to the mean of the coefficient range.

In the presence of anisoplanatic aberrations, the correction algorithm described in this section will produce an average compensation across different points of the field of view. If the spatial frequency content is uniformly distributed in the image, such as when imaging a monolayer of fluorescent beads, then the aberration correction is the average across the entire field of view. In case of non-uniformities, the correction is biased towards parts of the field of view where the image metric is highest.

4.3.5 Correction of system aberrations

The compensation of system aberrations was carried out in two stages. The first was to flatten the surface of the DM to avoid introducing unnecessary aberrations. This was done with a simple negative feedback controller where the measured Zernike modes (using the integrated interferometer) were fed back into the controller with a negative gain, but only if the corresponding RMS wavefront phase error was lower than in previous iterations. The procedure stopped either after a certain amount of iterations or when a threshold in the RMS phase error was crossed. The gain was kept above -1.0, and typically a large gain of -0.9 was enough to ensure convergence within 10 iterations. Figure 4.13 shows example phase maps of the neutral shape (neutral actuator voltages) and the shape after flattening.

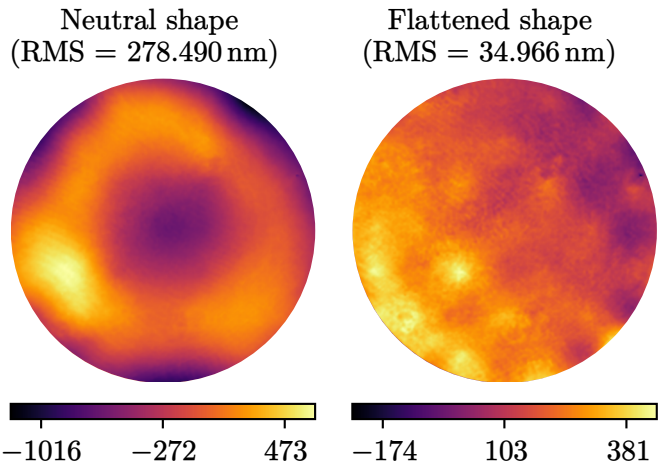
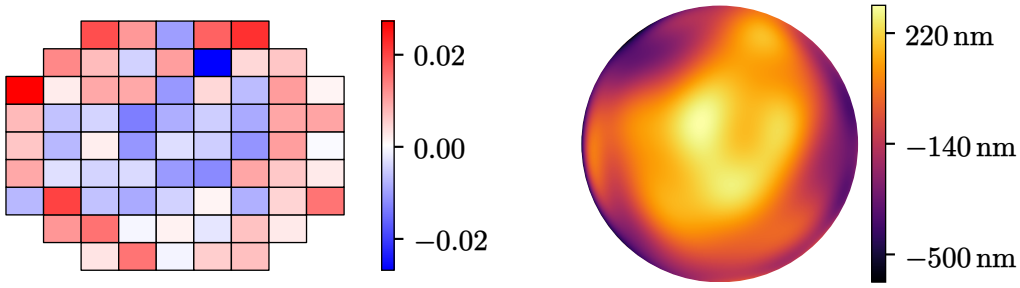


Figure 4.13: Wavefront phases before and after flattening. The DM was exercised before the flattening.



(a) Actuators.

(b) Wavefront phase.

Figure 4.14: Example actuator voltages and wavefront phase resulting from the compensation of system aberrations.

Despite flattening the DM there were still noticeable aberrations in the system, mainly a bit of defocus and astigmatism. They were attributed to the uncommon path differences between the interferometry and the main imaging paths. Regular sensorless aberration correction, using a sample with 100 nm fluorescent beads, was used to correct for those residual aberrations. Zernike modes 4 to 22 were usually scanned, with coefficient range of ± 1.5. The resulting corrections were treated as a baseline and all other corrections, related to sample aberrations or to remote focusing, were added on top of them. Figure 4.14 shows an example of the actuator values and phase associated with the system correction. They used a small fraction of the available dynamic range of the DM and they did not hinder the rest of the adaptive optics functionality.

The compensation of system aberration slightly improved the PSF of the imaging system, compared with bypass mode where the DM was switched with a regular planar

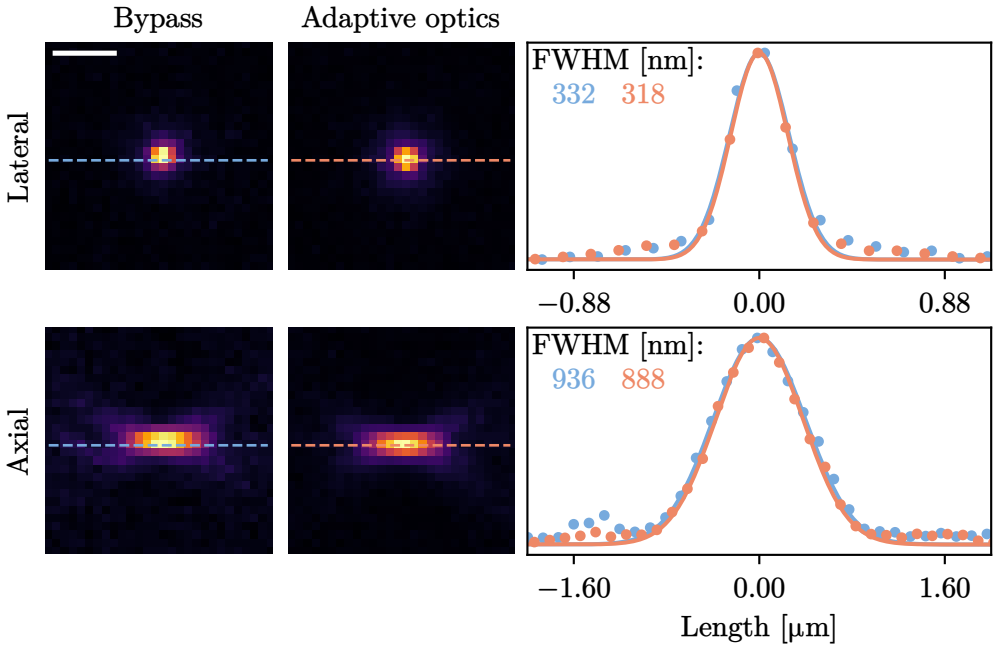


Figure 4.15: Lateral (XY) and axial (XZ) orthogonal projections of images of a red 100 nm fluorescent bead, with emission centred around 605 nm. The graphs on the right show the normalised intensity profiles along the dashed lines. The scale bar is 1 µm.

mirror. This is demonstrated in Figure 4.15 which shows lateral and axial orthogonal projections of images of a fluorescent bead. The full width at half maximum (FWHM) resolutions are given in the top left corner of the respective graphs.

4.3.6 Structured illumination microscopy

The 3D-SIM used an analogue SLM to generate sinusoidal light patterns (five phases) in three different orientation angles. A polarisation rotator was used to preserve the S polarisation of each orientation. The structured illumination was created from the interference of the 0 and ±1 order beams at the sample plane. The raw data was reconstructed with softWoRx (GE Healthcare).

4.3.7 Fly stocks

The fly lines used in the experiments described in this chapter are summarised in Table 4.2. All stocks were raised on standard cornmeal-based medium at 25C.

Table 4.2: List of fly lines, their source, and the sample in which they were used.

Line	Source	Sample
Nrx-IV::GFP (CA06597)	Kindly gifted by the Rita Teodoro lab	NMJ
Cno::YFP (CPTI000590)	DGRC 115111	Brains
Jupiter::GFP, His2A::RFP	Davis Lab stock, described in [99]	Embryos
Eb1::GFP (two copies)	Davis Lab stock	Macrophages

4.3.8 Sample preparations and fluorescent labelling

The NMJ samples were prepared in a similar way to the protocol described in [100]. The fixation was done in 4% paraformaldehyde in phosphate-buffered saline with 0.1% Triton X-100 (PBSTX) for 30 minutes, followed by two washes in PBSTX (0.1% Triton X-100) for 20 minutes each. For the antibody staining, the samples were then incubated with Cy3-conjugated anti-horseradish peroxidase (HRP) antibody (1:500, Jackson ImmunoResearch 123-165-021) for 1 hour. Finally, the samples were washed in PBSTX (0.1% Triton X-100) for 45 minutes and then mounted in Vectashield. All steps were done at room temperature.

The brain samples were prepared by dissecting L3 *Drosophila* larvae in phosphate-buffered saline (PBS). Their brains were removed and fixed in 4% paraformaldehyde in PBSTX (0.3% Triton X-100) for 25 minutes. Afterwards, they were rinsed three times in PBSTX (0.3% Triton X-100), further permeabilised with two 20-minute washes, and then mounted in PBS.

The embryo samples were prepared in the same way as described in [101], except that they were mounted in PBS instead of halocarbon oil.

The macrophages samples were prepared by making a small puncture in the cuticle of L3 *Drosophila* larva in a petri dish in Schneider's medium. The wounded larva was removed after dragging it around the dish for several seconds. This was repeated 5

times. The dish was then covered and left to incubate for 1.5 hours at room temperature in a humid environment.

4.4 Results

4.4.1 Remote focusing for fast motionless volumetric imaging

The objective of the remote focusing was to provide an alternative to the Z piezo stage, especially in experiments. The rapid synchronised use of devices during experiments relied on a real-time controller and the DM needed to be integrated into this configuration. It needed to respond to digital trigger signals, which instructed the device to perform an action. The drive electronics for the Alpao DM69-15 had a memory storage where a collection of actuator voltage patterns was loaded in advance, and then each digital trigger switched to the next pattern in the queue. Every trigger received after the final pattern had been reached did not have any effect and the last pattern remained in place. This worked well for the most part, besides a bug in the device's firmware which is discussed in Appendix A.1.

The remote focusing module was designed to move the focal plane up or down by reshaping the light with the DM. Because the device was controlled with Zernike modes, the module had to associate mode coefficients with concrete axial displacements. In other words, it was necessary to define a function that mapped any given offset Δd to a set of Zernike mode coefficients a_j. The imaging system had an upright configuration and the convention was that a positive offset Δd meant going into the sample or, equivalently, moving the sample up with the stage. Conversely, $-\Delta d$ was displacement away from the sample or moving the sample down.

The mapping $\Delta d \to a_j$ was determined in a calibration step, using a sample that contained a layer of 100 nm red fluorescent beads. The offset Δd was introduced with the Z piezo stage, illustrated for $\Delta d = -5\,\mu\text{m}$ in Figure 4.16. Regular sensorless aberration correction (as described in section 4.3.4) was then used to refocus the beads. In theory, the modes related to the focus, namely defocus and spherical aberrations, were supposed to be symmetric around the zero position, such that the corrections for $-\Delta d$ and Δd differed only in sign. In this case the mapping $\Delta d \to a_j$ was the line equation best fitted to the points $\{-a_j, 0, a_j\}$ and thus it could be determined from

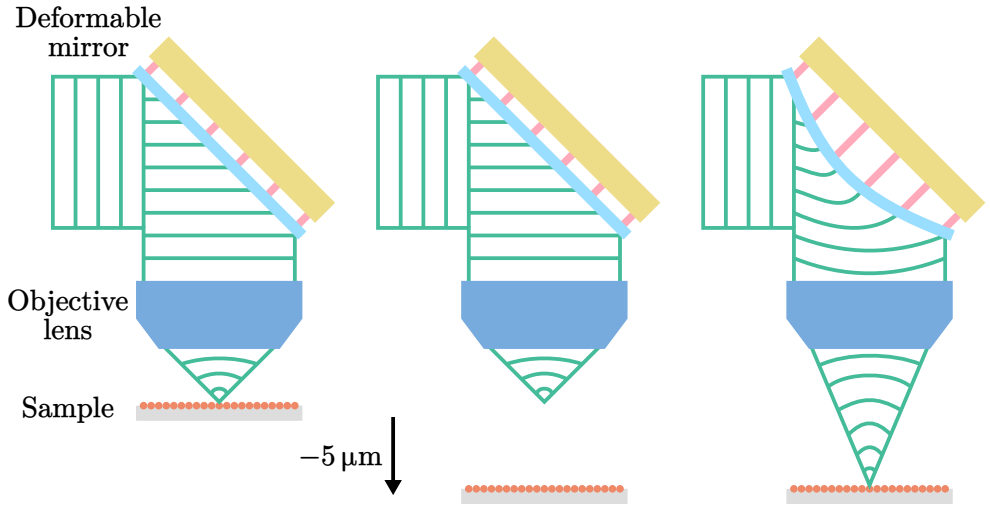


Figure 4.16: Using a DM to refocus on a layer of fluorescent beads after they have been moved downwards by 5 µm.

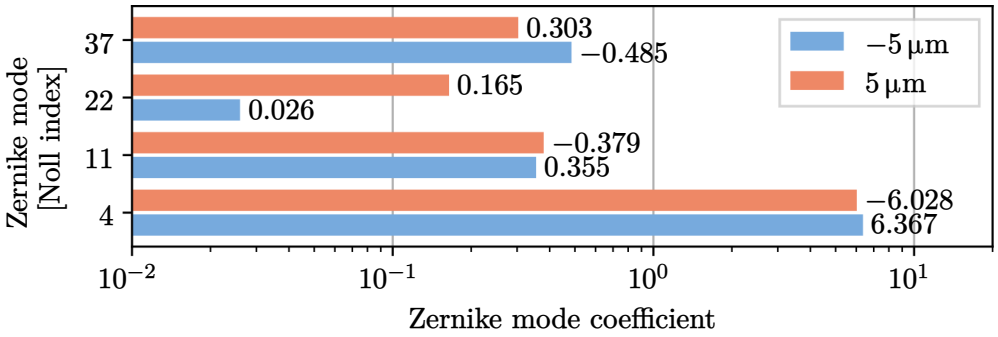


Figure 4.17: Corrections for defocus and primary to tertiary spherical aberrations for two offsets −5 µm and 5 µm during remote focusing calibration.

any one of the two points alone. In practice, however, there was a noticeable difference between the two, which is shown in Figure 4.17 for defocus and primary to tertiary spherical aberrations.

The performance of the remote focusing was investigated further by looking at the image quality and the accuracy of the axial displacement at several different positions. The idea was to use the piezo stage to move up or down a sample with a layer of fluorescent beads, and then to see if the image of the beads shifted by the same amount in a z stack volume. Four axial offsets were used in this experiment: −5 µm, −2.5 µm, 2.5 µm, and 5 µm. Two volumes were acquired for each offset – one with the piezo, used as a reference, and another with the remote focusing. Orthogonal projection followed by Gaussian fitting was used to determine the axial positions of the beads. Volumes without any offset were acquired before each change of position in order to

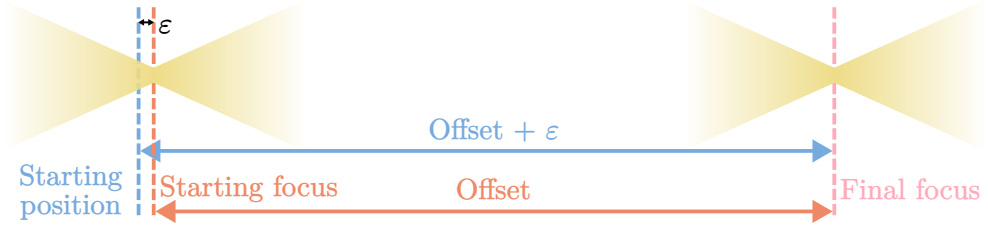
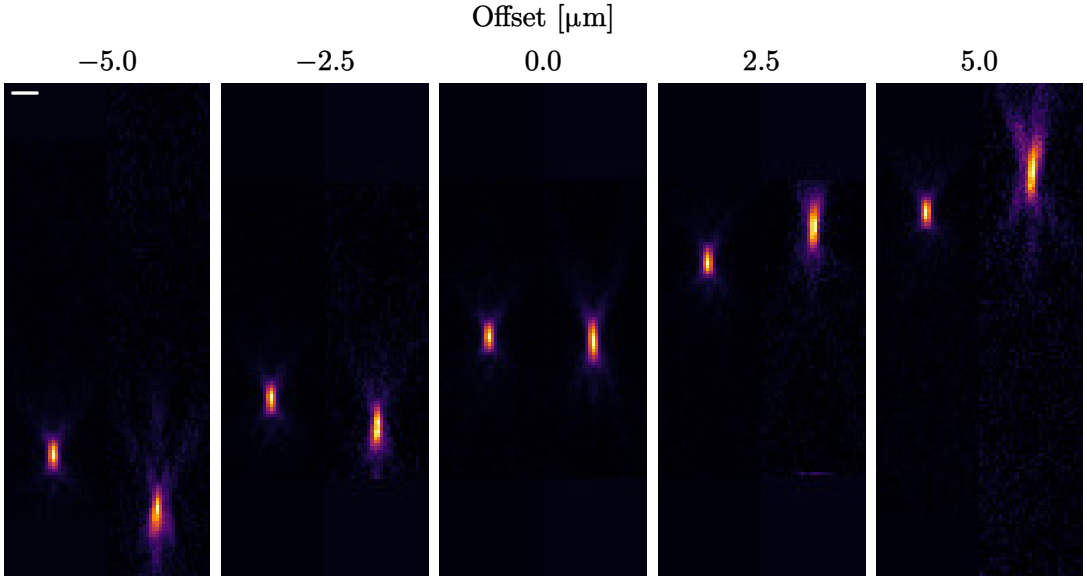


Figure 4.18: Focus error in offset calculation. Not taking note of the starting focus before moving the sample (represented by the PSF in yellow) could lead to wrong estimation of the real offset. Acquiring a volume at the starting position, in addition to the end point, provides a measurement of ε.

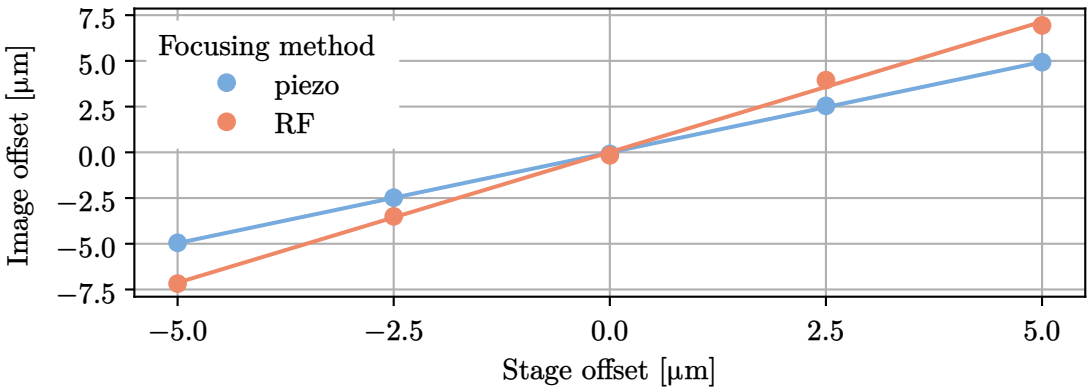
avoid errors in the calculation of the actual offset lengths, as illustrated in Figure 4.18. The volume acquisition with remote focusing was repeated ten times for each offset and ten beads were arbitrarily selected from the field of view for further processing. The average standard deviations between repeats and between beads were 78.78 nm and 64.13 nm, respectively, demonstrating the high precision of the DM. Despite this, there were problems of inaccuracy of the axial displacements and of increased aberrations, specifically spherical (axial elongation), as shown in Figure 4.19a.

The cause of these two issues was interpreted to be the drift effect described in section 4.3.2. During the calibration of the remote focusing, the correction at Δd resulted in high quality images but this was achieved over the span of minutes, where the bulk of the drift had already happened. Therefore, the bias from this drift was already compensated for in the correction, leading to larger mode coefficients. In contrast, during experiments the shape of mirror changed on the scale of milliseconds, before the drift had taken its full effect. In this situation the previously derived large mode coefficients pushed the focus further away, as well as introducing other aberrations. In summary, the calibration of the remote focusing included bias which was absent in experiments because of the speed at which the mirror shapes were changed. This behaviour also explained the differences in Figure 4.17, as doing the correction on one side created a bias, such that the opposite side was slightly pulled towards it.

A second calibration step was required to address the issues caused by the drift. The mode coefficients for the remote focusing needed to be scaled down by a factor, which was determined from the same experimental data (Figure 4.19). It was derived as the ratio between the slopes of the offset of the bead images and the actual axial



(a) YZ orthogonal projections of bead images for various axial offsets. Two volumes were acquired for each offset, one with the piezo stage (left) and another with remote focusing (right). The scale bar is 1 µm.



(b) Comparison of bead image offsets between the piezo stage and the remote focusing.

Figure 4.19: Calculation of scaling factor for remote focusing.

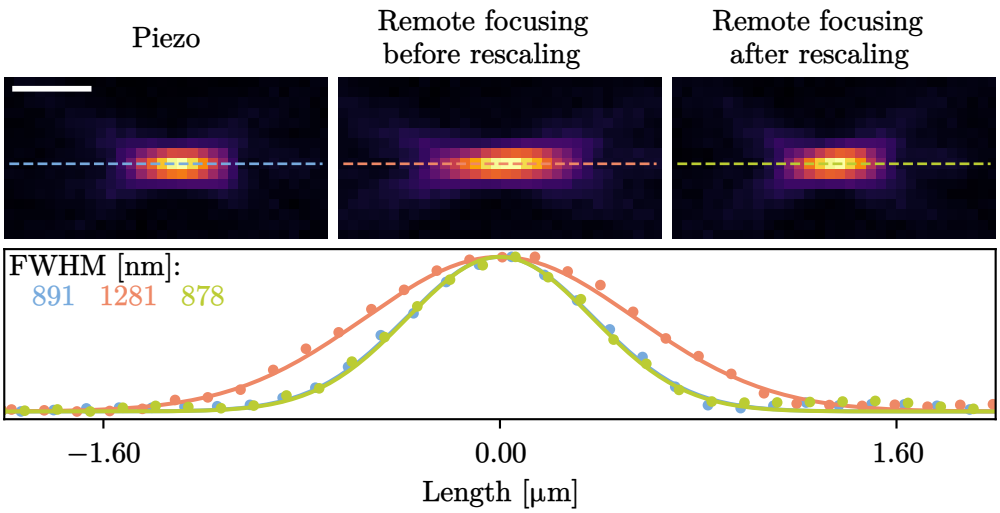


Figure 4.20: Impact of rescaling on the remote focusing image quality. Top row shows YZ orthogonal projections of a 100 nm fluorescent bead. The graph on the second row shows normalised intensity profiles along the dashed lines. The scale bar is 1 µm.

position (Figure 4.19b). For example, if the beads were offset by 2.47 µm but they appeared 3.61 µm away in the remote focusing images, then the scaling factor would be $2.47/3.61 \approx 0.684$. Figure 4.20 demonstrates how the scaling method restored the image quality and even improved the axial resolution.

In agreement with [95, 96], the drift effect in the Alpao DM69-15 was deterministic, predictable, and consistent. As such, the scaling factor was determined mainly by the exposure time. There was no opportunity to explore this relation in great detail, but the experiments with this particular imaging system used mostly the same low exposure times, in the range 10 ms to 100 ms. In addition, the use of the remote focusing was restricted to a range of 10 µm, partially because this was more than enough for most experiments but also because the PSF started to degrade beyond this range and it was desired to have optimum image quality, especially for SIM experiments.

4.4.2 Sensorless aberration correction in the *Drosophila* PNS

The performance and the impact of the sensorless aberration correction was assessed in the *Drosophila* NMJ because it is one of the most common systems for studying RNA transport, amongst other applications such as RNA localisation and synaptic plasticity. The samples were prepared in a way that exposed the brain and the motor neurons, as well as the synapses that they form with the muscles. The imaging configuration of such a fillet preparation is shown in Figure 4.21.

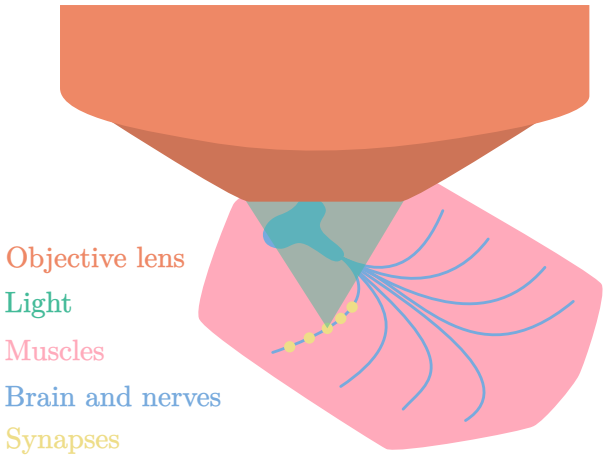


Figure 4.21: Upright imaging configuration of *Drosophila* larval fillet dissection. Motor neurons from the central brain extend axonal projections that connect to muscle fibres all along the body of the larva. They form many synapses and some of them are illustrated with yellow dots.

Despite some variation in altitude, NMJs were mostly situated on the surface of the muscles and therefore they were very accessible with the fillet dissection. However, this also meant that there was not much tissue or mounting medium in the way, so there was less opportunity for RI mismatch and thus the aberrations were less impactful. Depending on the sample preparation it was even possible to have some of the NMJs pressed right against the coverslip.

Larvae expressing Bruchpilot (Brp) fused to GFP were first imaged in widefield mode, as shown in Figure 4.22. Brp is expressed at the NMJ in the presynaptic active zones and often appears as diffraction limited spots, which made it a good model for tracking of single molecules and for quantitative imaging in general.

The same kind of sample was then imaged with SIM (Figure 4.23), because like other super-resolution methods it was sensitive to the image quality and therefore it highlighted better the advantages of aberration correction at shallow depths. The larvae expressed Neurexin IV (Nrx-IV) fused to GFP, localised to septate junctions of glial cells, and they were stained with anti-HRP to visualise the neuronal membranes.

4.4.3 Sensorless aberration correction for deep imaging

The sensorless aberration correction was also evaluated for deep SIM imaging using explanted fixed *Drosophila* third instar larval brains. A line with YFP protein trap in Canoe (Cno) was used because it was expressed at the adherens junctions between

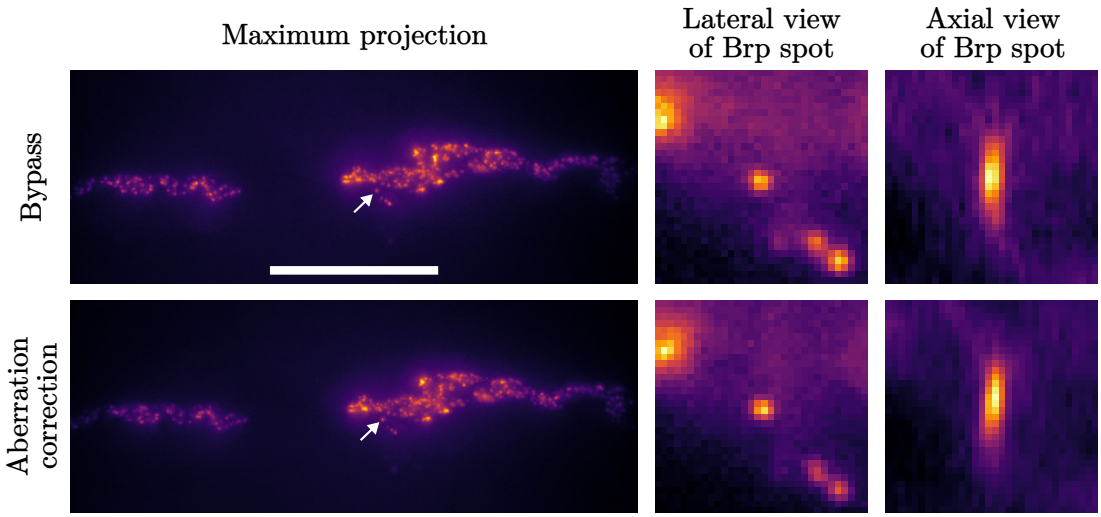


Figure 4.22: Impact of aberration correction in widefield mode. The images show expression of Brp::GFP in a third instar larva. The spot used for the lateral and axial views is indicated with a white arrow in the maximum projections. The scale bar is 20 µm.

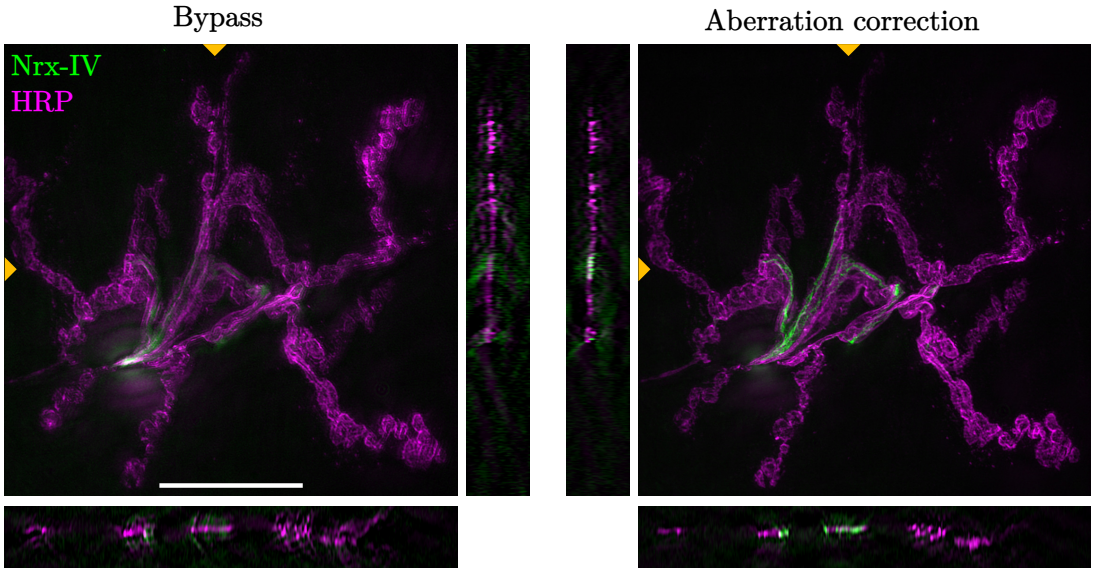


Figure 4.23: Impact of aberration correction in SIM. The images show expression of Nrx-IV::GFP in a third instar larva, which is also stained with anti-HRP. The orthogonal projections are along the directions indicated by the golden arrowheads. The important feature to highlight is the improved contrast in the green channel, the reduced reconstruction artefacts, and the improved resolution. Particularly note the improvement in the shape of the green signal in the orthogonal projections. The scale bar is 20 µm.

neuroepithelial stem cells in the outer proliferation centres in the optic lobes. Therefore it provided similar structures all around the brain lobes' circumferences, allowing both shallow imaging at the surface of a lobe and deep imaging on the opposite side through the entire lobe. A schematic of this arrangement is shown is Figure 4.24a, while the corresponding data is presented in Figures 4.24b and 4.24c. The deep region, imaged about 130 µm below the bottom surface of the cover slip, was reached by mechanical translation the Z stages.

4.4.4 Sensorless adaptive optics for live imaging

The improvements brought by adaptive optics were difficult to quantify and demonstrate in live imaging experiments where two datasets inevitably showed different events. An example of such comparison is given in Figure 4.25. It shows syncytial mitotic division in the blastoderm of early *Drosophila* embryos. The microtubules were tagged with GFP and the histones were tagged with RFP. This took place around 5 µm under the surface of the embryo, but the images were blurry because of the vitelline membrane and the waxy layer around the embryo.

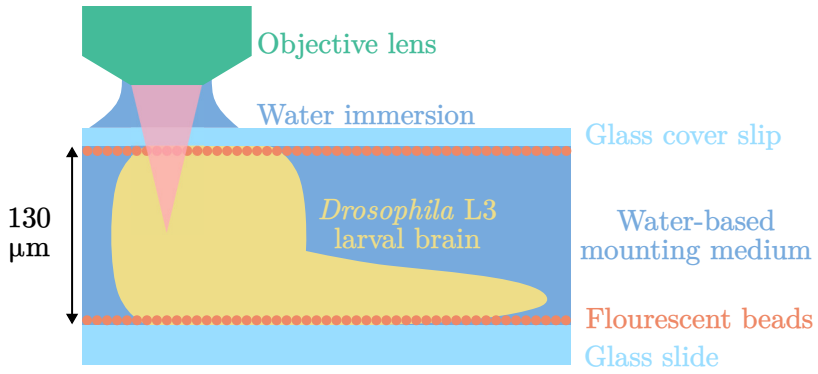
The full capability of the adaptive optics module was demonstrated by performing the same experiment, but using both aberration correction and remote focusing with SIM. The results are shown in Figure 4.26.

4.4.5 Multi-plane aberration correction

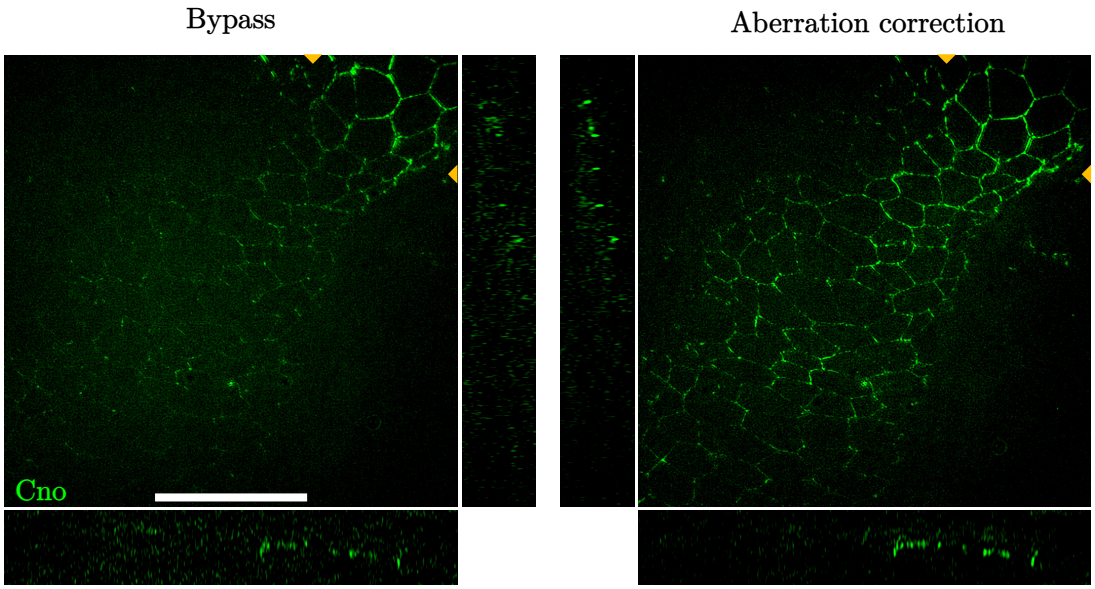
Inspired by the results from the remote focusing, and borrowing the already developed functionality, multi-plane aberration correction was performed. The results are shown in Figure 4.27. The correction was done at the top and at the bottom of the stack, then the resulting Zernike mode coefficients were fitted to a straight line, and finally correction was performed for each stack slice by interpolating.

4.4.6 Extended focus

One of the advantages of the remote focusing was the ability to shift the focal plane rapidly without perturbing the sample. This was exemplified with extended focusing where the focus was changed during the camera exposure, leading to increased depth of field. Figure 4.28 shows an example of this.

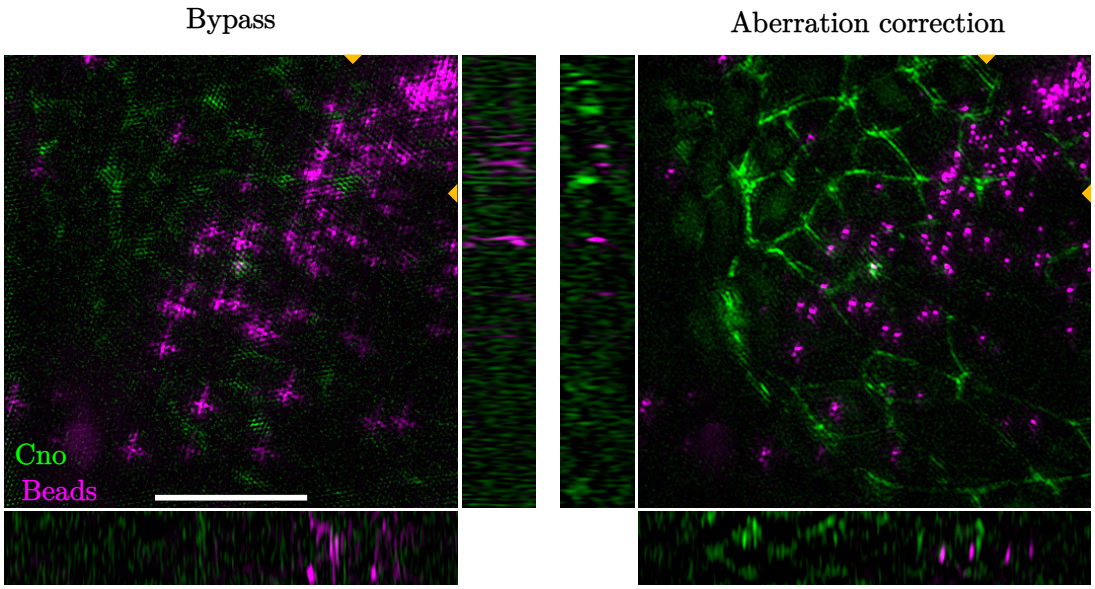


(a) Imaging configuration of fixed explanted brains.



(b) Expression of Cno::YFP at the top of the brain, around 3 µm below the coverslip. Aberration correction was performed on a layer of 100 nm red fluorescent beads on the cover slip. The orthogonal projections are along directions indicated by the golden arrowheads. The aberration correction improves the resolution and the contrast, and this is most noticeable in the middle and the bottom parts of the images, where the honeycomb-like structure of the expressed protein suddenly becomes visible. This is also noticeable in the orthogonal views where the diffuse fuzzy cross-section of junctions become proper bright spots. The scale bar is 20 µm.

Figure 4.24: Aberration correction for deep 3D-SIM.



(c) Expression of Cno::YFP at the bottom of the brain, around 130 µm below the coverslip and 1 µm above the glass slide. Aberration correction was performed on a layer of 100 nm red fluorescent beads on the glass slide. The orthogonal projections are along directions indicated by the golden arrowheads. The beads provide a good representation of the poor PSF of the microscope at this depth in bypass mode. The protein expression is also almost not visible without the aberration correction, which is the thing enabling super-resolution imaging at this depth. The scale bar is 10 µm.

Figure 4.24: Aberration correction for deep 3D-SIM.

4.5 Discussion and conclusions

Adaptive optics holds the potential to bring many improvements to fluorescence imaging, live or fixed. The more aberrating the sample, the greater the benefit. It is particularly challenging to live imaging, as the dynamics of the sample and phenomena like photo-bleaching can severely impact the correction algorithms in negative way. Moreso, the algorithms struggle to work without a guide that is flat and with good contrast. The results presented in this chapter were achieved by using fluorescent beads which was a robust and reliable approach, but it is not always an option, such when the structure of interest is deep into a tissue where beads cannot be deposited. IsoSense [38] is an option, but it is a complex and potentially expensive mechanism which is still only good in a restricted set of situations.

Even for shallow samples and others that have little to no aberrations, adaptive optics can still provide many advantages in the form of remote focusing, which is broadly advantageous. It brings speed (the response of DM used in this chapter was around 2ms and it is compared to the piezo stage in the Appendix A.2), motionlessness,

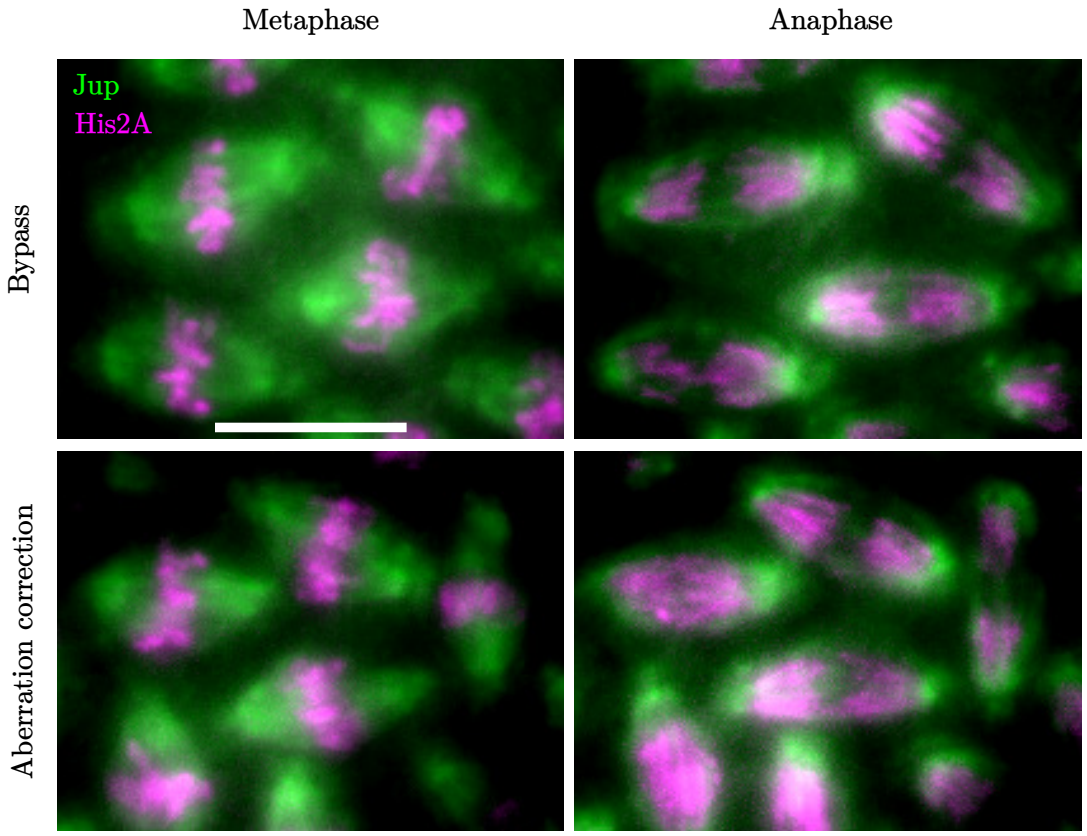


Figure 4.25: Aberration correction for live imaging in widefield mode. The images show the metaphase and the anaphase of the mitotic division in the syncytial blastoderm of early *Drosophila* embryos. Microtubules and chromosomes are visualised with transgenic Jup::GFP and His2A::RFP, respectively. The images are maximum projections. The aberration correction brings very little noticeable improvement, mainly in the contrast. The reason for this is that the blastoderm is close to the surface of the sample, the aberrations are not that strong to begin with, and widefield is much more forgiving modality compared to super-resolution. The scale bar is 10 µm.

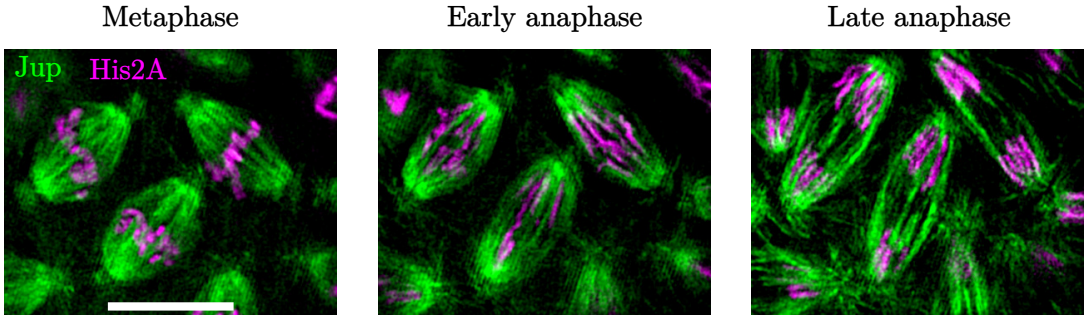


Figure 4.26: Aberration correction and remote focusing in live SIM. The images show metaphase, early anaphase, and late anaphase during mitosis in the syncytial blastoderm stage of early *Drosophila* embryonic development. Microtubules are labelled with Jup::GFP and chromosomes with His2A::RFP. The images are maximum projections. The good performance of the remote focusing is demonstrated by the lack of strong reconstruction artefacts, which appear as ringing around sharp edges or as grid-like pattern. 3D-SIM is particularly sensitive to aberrations, demanding as perfect imaging performance as possible, and therefore the absence of prominent artefacts is a testament to quality of the remote focusing. The scale bar is 10 µm.

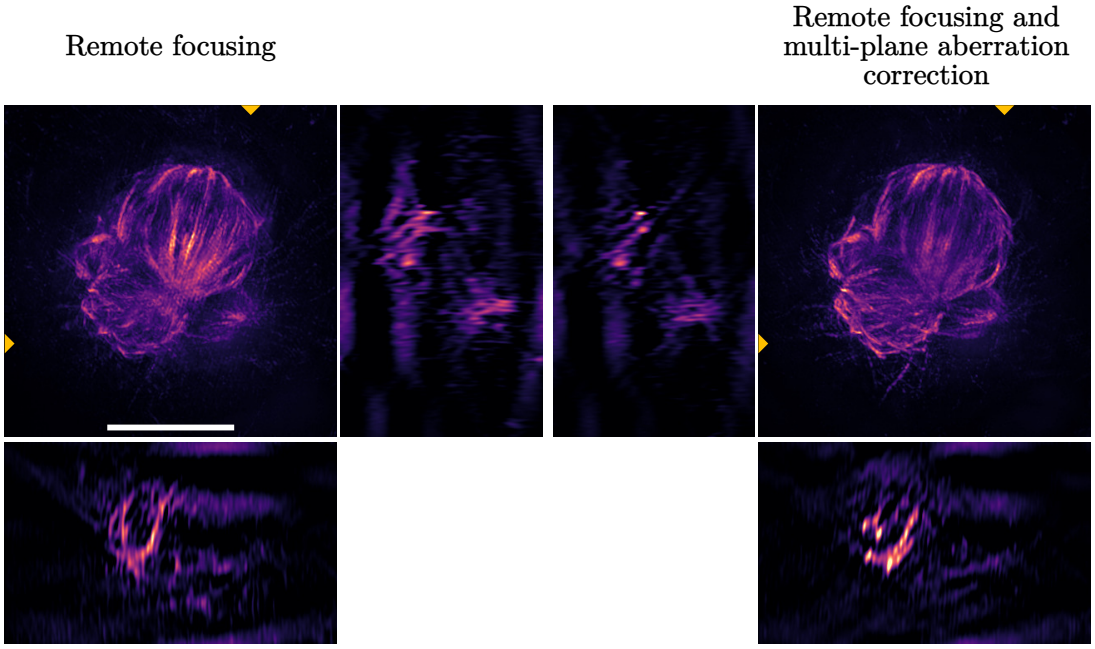
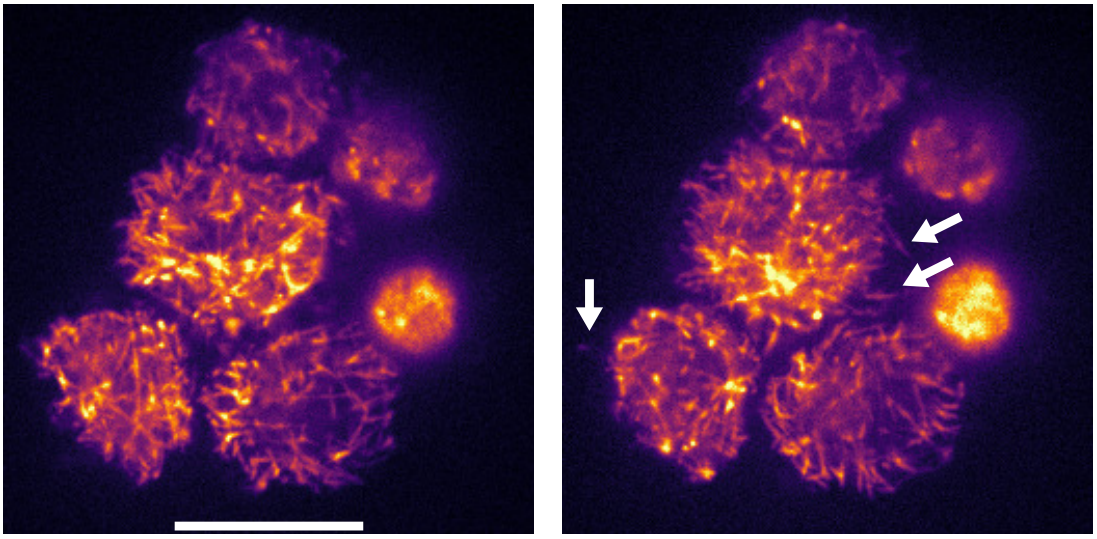


Figure 4.27: Remote focusing and multi-plane aberration correction in 3D-SIM. The images show the spindle fibres in fixed COS-7 cells. The microtubules were labelled with Alexa Fluor 488. The orthogonal projections are along directions indicated by the golden arrowheads. Despite being interpolated, the corrections drastically improved the contrast and resolution, which is demonstrated very well by the cross-sections of microtubules present in the orthogonal projections. The scale bar is 10 µm.



(a) Single plane. (b) Extended focus.

Figure 4.28: Comparison between single plane and extended focus time-lapse data. The images show standard deviation projections (in the time axis) of macrophages expressing Eb1::GFP. The white arrows indicate how the larger depth of field has captured more detail. The scale bar is 5 µm.

and it could even be used in applications where a traditional microscope stage cannot be used, such as with large or heavy specimen. Despite problems with the DM, excellent precision, accuracy, and range were demonstrated in the results. Newer devices should do even better.

The deep imaging in Figure 4.24, although shown on fixed samples, applies equally well to live samples. It was not demonstrated on such only out of convenience. The sample was already mounted in PBS and it did not go through any optical clearing procedures. This is a scenario in which adaptive optics really excels, and it is the thing that enabled SIM very deep, far exceeding the recently shown 80 µm in *C. elegans* [93].

Besides the remote focusing, two other interesting applications of the developed methodology were demonstrated: multi-plane aberration correction and extended focus. Both show impressive performance and they could be very useful in relevant applications.

Chapter 5

Denoising of fluorescence images

5.1 Introduction

Noise is an unavoidable part of fluorescence microscopy. Even a perfect imaging system would still experience fluctuations in the photons it detects, in the form of shot noise. The situation is worsened by the fact that despite its usefulness and its inevitability in the formation of images, light has deleterious effects both to living matter and to fluorophores, in the related processes of phototoxicity and photobleaching. Limiting these negative effects necessitates the use of low dose of light, which could also be imposed by the experimental requirements, such as short exposure time for the sake of high frame rate. The low light conditions lead to images with low signal-to-noise ratio (SNR). Denoising methods could be used to restore some of signal, as well as enable experiments over longer periods of time and under more physiological conditions. In some situations they can even make possible the use of new microscopy modalities for live imaging [102].

Recent technological developments have enabled the practical implementation of neural networks for image processing [103, 104]. Their potential for image denoising has been known for a while and the first such application of a CNN, designed to remove Gaussian noise, demonstrated the potential to match or even outperform state-of-the-art traditional methods [105]. It was nearly a decade later that denoising with CNNs gained interest again within the broader scope of image restoration [106], although other types of neural networks had been used in the meantime [107–109].

Despite their impressive performance, all of these methods operate in a supervised

manner, which requires clean-noisy pairs (x, y) of images for training. This kind of data is not always possible to acquire or it is challenging to do so, which prompts the need to train with just noisy data. However, if the same noisy data were used to construct pairs (y, y), then the network would overfit and it would simply learn the identity function. To prevent this from happening, a different piece of noisy data z needs to be introduced, resulting in pairs (z, y). The difference between training with or without clean data is mathematically expressed as:

$$\underset{f_\theta}{\operatorname{argmin}} \mathbb{E}\left[(f_\theta(y) - z)^2\right] = \underset{f_\theta}{\operatorname{argmin}} \left\{\mathbb{E}\left[(f_\theta(y) - x)^2\right] - 2\mathbb{E}\left[f_\theta(y)(z - x)\right]\right\} + \text{const.} \tag{5.1}$$

The proof of equation 5.1 is given in Appendix B.2.

The first to take advantage of this relation is Noise2Noise [110]. The authors made three key assumptions about the training data. The first assumption is that y and z are statistically independent noisy images, given the clean image x. The second assumption is that both noisy images have the same underlying signal x, following the corruption model from equation 2.15. The third assumption is that the two noise components have an expected value of zero, such that $\mathbb{E}\left[y - x \mid x\right] = \mathbb{E}\left[z - x \mid x\right] = 0$. The second term in equation 5.1 vanishes under these three assumptions and the resulting expression shows that training with only noisy data is the same as training with clean data, up to a constant. The network learns to give the mean of the noisy data z, given the other noisy data y, which is equivalent to the supervised case of mean clean data (equation 2.27) under the second and third assumptions: $f_\theta^\star = \mathbb{E}[z \mid y] = \mathbb{E}[x \mid y]$.

Relaxing the constraints from noisy-clean training pairs to just noisy pairs is certainly a step in the right direction, but it is still difficult to even not possible, especially for dynamics scenes. For this reason various ways have been developed to make it easier to obtain the training data required for Noise2Noise type of denoising. One of the most suitable applications is time-series image data. If the frequency at which the frames are acquired is higher than the frequency at which the scene changes, then every few consecutive frames have the same signal but different noise realisations and therefore pairs of frames can be used for training [111]. Another approach is down-sampling, used to produce a pair of noisy images from a single one. In this case the underlying signals of the two images differ slightly and the second term in equation

5.1 needs to be considered, unlike the original Noise2Noise approach where this term vanishes. The difference between the two images is usually small enough that it can be ignored [112], but it is also possible to compensate for it [113]. Finally, there are various generative methods that can synthetically create training data, but they are usually not blind, meaning that they require either clean images or prior information, such as noise statistics [114–116]. Therefore, generative methods are not considered here.

A slightly different approach than Noise2Noise, allowing self-supervised training on single (unpaired) noisy images without resorting to downsampling or synthetic data generation, is taken by the blind spot class of methods, pioneered by Noise2Void [117, 118] and Noise2Self [119]. They rely on the fact that CNNs make predictions based on the entire receptive field and not just individual nodes, as illustrated in Figure 2.12. Therefore, the network can be prevented from overfitting to the identity function by making it blind to the pixels it tries to predict. In this case, the predictions are based entirely on the rest of the pixels in the receptive field. This is still mathematically described by equation 5.1, but crucially z is the pixel in the centre of the receptive field and y are all the remaining pixels in the same receptive field. In this sense, y and z are partitions of the receptive field, as well as complements of each other, such that $z = y_{\mathrm{RF}} \backslash y$.

One of the greatest limitations of blind spot methods is the fact that the most important part of receptive field – the central pixel – is absent from the prediction, which reduces the effectiveness of the network. One way to bring back this information is by Bayesian inference:

$$p(x \mid y, z) \propto p(y \mid x, z) \cdot p(x \mid z) \tag{5.2}$$

Laine et al [120] approach this expression by explicitly specifying the distributions on the right-hand side. The prior is taken to be a multivariate Gaussian $p(x \mid z) \sim \mathcal{N}(\mu, \Sigma)$ and the likelihood (noise model) $p(y \mid x, z)$ can take the form of various common distributions such as Gaussian or Poisson. Two networks are then trained to estimate the parameters of these two distributions. A clean image is estimated in a second step, where the mean posterior is computed analytically, using the outputs of the two

networks.

Krull et al [121] and later Prakash et al [122] take a similar approach that differs in two major ways. First, the prior is sampled instead of being explicitly defined. This means that the network trained to estimate the prior will output a number of samples (800 in this case) for each pixel, instead of the parameters of the distribution. The second difference is the need for clean-noisy image pairs for the training of the network that estimates the likelihood. This makes it necessary to have a separate calibration step.

The presence of all of these novel CNN-based denoising methods, from Noise2Noise to the probabilistic blind spot networks, motivates the question whether they are suitable for live fluorescence imaging and how do they compare to the state-of-the-art traditional methods, such as nD-SAFIR. The development and validation of many of them were heavily based, or even exclusively done, on public datasets such as Kodak [123], BSD300 [124], and Set14 [125], which are very different from the type of data typically used in biomedical imaging. This makes it difficult to judge how well such methods will work in practice and it creates the need for proper evaluation and comparison.

In this chapter, I compare the performance of different denoising methods on different types of data. I highlight their advantages and disadvantages, and discuss which ones are most suitable for *in vivo* live imaging of single mRNA molecules.

5.2 Aims and objectives

The overall aim is to investigate how well deep-learning denoising methods can perform on fluorescence image data. Of particular interest is the case of tracking of mRNA particles in *Drosophila*, but also other kinds of data, such as the ones generated by fluorescence in situ hybridisation and hybridisation chain reaction experiments. Generalisation and performance over broader range of samples and specimens is important, too. The deep-learning methods should be compared to well-established traditional ones, serving as a baseline. There are three main objectives:

1. Acquire rich image data sets which are representative of real fluorescence imaging data and which can be used to develop, train, and evaluate mathematical models

based on machine learning

2. Select different categories of deep-learning image denoising methods, implement them and train them on the previously acquired data sets, and compare their performance to traditional denoising methods

3. Investigate the strengths and weaknesses of the different deep-learning denoising methods, their impact on the study of live mRNA dynamics, and identify specific algorithms that are suitable for this task

5.3 Methods and materials

5.3.1 Microscopy

The data was acquired on a custom Olympus IXplore spinning-disk microscope, equipped with Photometrics Prime BSI camera and PI P-736.ZR1 piezo stage. The hole size of the spinning disk was 50μm. The camera was operated in 16-bit mode (HDR combined gain) and default settings, which included despeckling with bright thresholds 115 and 120 (200 min ADU), and dark thresholds 85 and 85 (200 min ADU). The objective lens was Olympus UPlanXApo 100× 1.45NA. All experiments used a 488 nm laser line, whereas multi-channel (green and red) experiments also included a 561 nm line.

5.3.2 Computational environment

The data processing was done on a standard desktop computer, equipped with an AMD Ryzen 9 5950X processor, 128GB of memory, and a NVIDIA GeForce RTX 3090 discrete graphics card. The computer had the Windows operating system, but all software was ran in a Linux environment, using the Windows Subsystem for Linux. The Linux distribution was Ubuntu 22.04 (Jammy Jellyfish) with version 5.15 of the kernel. The NVIDIA computing platform included CUDA version 11.8 and cuDNN version 8.6. All software was developed and ran with Python version 3.10.

5.3.3 Datasets and sample preparation

The fixed 3D data was acquired from an Invitrogen FluoCells Prepared Slide #1 (catalogue number F36924), which contained fixed bovine pulmonary artery endothelial cells.

Table 5.1: Selected denoising methods

Method	Description	Label
nD-SAFIR	baseline traditional method	nD-S
Noise2Noise	semi-supervised method (pre-trained model)	N2N
Noise2Void v2	blind spot self-supervised method	N2V
Probabilistic Noise2void	probabilistic blind spot self-supervised method (with calibration)	PN2V
Noise2Fast	downsampling self-supervised method	N2F

Nuclei were stained in blue with DAPI, actin filaments were stained in green with Alexa Fluor 488 phalloidin, and mitochondria were stained in red with MitoTracker Red CMXRos.

The live time-series (2D+t and 3D+t) data was acquired in explanted *Drosophila* larval brains, using a fly line with two insertions of EB1, each fused to GFP: EB1-GFP;;EB1-GFP. Crawling L3 larvae, grown on standard cornmeal-based medium at 25C, were dissected in Schneider's medium and immediately mounted on a glass slide with a #1.5 thickness glass cover slip, again in Schneider's medium.

5.3.4 Selection of deep-learning denoising methods

The deep-learning methods comprise a broad selection of blind denoising techniques. Priority was given to methods which were specifically designed for fluorescence microscopy. Table 5.1 lists them and provides brief description for each of them. To be able to include evaluation of a semi-supervised method such as Noise2Noise, pre-trained models from [126] were used. Noise2Void v2 [118] and Probabilistic Noise2Void [122] were chosen as the representative blind spot methods, with the latter being the only one in the probabilistic category with open source code and readily available tools. Noise2Fast [112] was picked amongst the downsampling methods because of its excellent trade-off between speed and quality.

5.3.5 Baseline denoising method

The baseline denoising method was chosen to be nD-SAFIR, mainly because it is one of the best traditional blind denoising methods for fluorescent microscopy data, able to remove mixed Gaussian-Poisson noise very effectively. The Linux version 2.2 of the software was used and all data was passed to it in the form of separate 3D volumes. Each channel was denoised separately for data with multiple channels. The software was used with arguments `-iter 3`, which sets the number of iterations of the algorithm to $2^3 = 8$, and `-p 2` which sets the size of the neighbourhood to 5×5.

5.3.6 Neural network training

Noise2Noise configuration

There was no training stage because pre-trained models were used. For the prediction stage, images were normalised to range $[-0.5; 0.5]$ and split into tiles of size 256×256 pixels. After prediction, the tiles were assembled, with overlapping areas being averaged, and denormalised to the original intensity ranges. All channels were processed separately.

Noise2Void configuration

Each noise level and each channel were trained separately, leading to 5 models for the 2D+t data, 10 models for the 3D data, and 5 models for the 3D+t data. The 2D+t and 3D+t data was augmented in the standard U-Net way, which is $8\times$ via flipping and rotation. All extracted patches were randomly shuffled and then 10% were withheld for validation. The model hyperparameters were chosen mainly according to [118] and they are listed in Table 5.2.

Table 5.2: Noise2Void hyperparameters.

Parameter	Value
Depth	3
Kernel size	3 px
Initial filters	64

Continued on next page

Table 5.2: Noise2Void hyperparameters. (Continued)

Parameter	Value
Epochs	200
Steps per epoch	256
Initial learning rate	0.001
Batch size	80
Blind spots (as percent of random pixels per patch)	0.198%
Patch size (2D data)	(64, 64) px
Patch size (3D data)	(8, 64, 64) px
Manipulator	"median"
Manipulator neighbourhood radius	2 px

Probabilistic Noise2Void configuration

Similarly to Noise2Void, the training was done separately for each noise level and channel, producing 5, 10, and 5 models for the 2D+t, 3D, and 3D+t datasets, respectively. The hyperparameters were chosen from [122] and included default values for everything, except for 5 steps per epoch and an initial learning rate of 0.001. The noise was modelled as a mixture of 3 Gaussian distributions and 2 coefficients per parameter, and it was trained with a batch size of 250,000, for 4000 epochs, with a learning rate of 0.1.

The calibration data was acquired with the same spinning-disk microscope as the rest of the data, and it contains images of an acrylic fluorescent calibration slide, excited with the 488 nm laser line, with a slightly closed aperture in the spinning-disk unit (down to 11mm) to create a gradient in intensity levels. This was done in order to generate signal content within the entire dynamic range of the camera, so that the subsequent model can accurately describe the noise of all possible intensity values. The data consisted of 100 frames, taken at 200ms exposure and 100% laser power. An average of all frames and its histogram are shown in Figure 5.1.

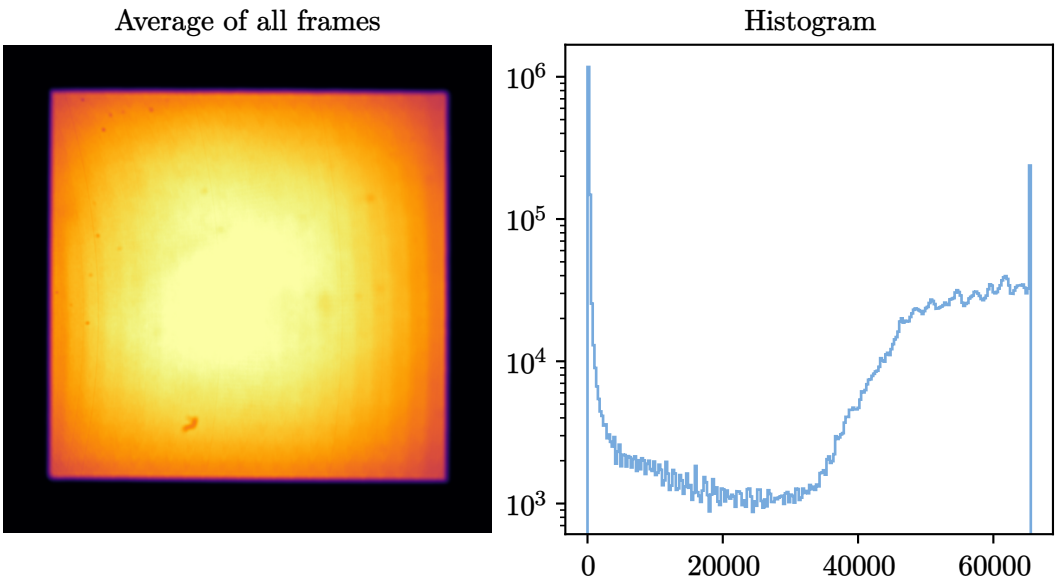


Figure 5.1: PN2V calibration data overview. The histogram was created with 256 bins and over the entire dynamic range of the 16-bit image on the left: $[0, 65535]$.

Noise2Fast configuration

There are no hyperparameters associated with this method. All data was passed directly to the provided tool, which processes each image separately, making no distinction between frames, channels, or sections.

5.3.7 Image quality metrics

The denoised images were compared to their reference (ground truth) counterparts by two image quality metrics. The first was peak signal-to-noise ratio (PSNR) , which is based on the MSE. For a reference image x and second image y, it is defined as the ratio between the powers of the noise and the peak signal:

$$10 \log_{10} \left(\frac{x_{\text{peak}}^2}{\text{MSE}(x, y)} \right) = 10 \log_{10} \left(\frac{(x_{\max} - x_{\min})^2}{\mathbb{E}\left[(x - y)^2\right]} \right) \tag{5.3}$$

Higher values are better, as they correspond to smaller difference between the two images.

The second metric was structural similarity index measure (SSIM) [127], which assesses a combination of luminance, contrast, and correlation, to account for perceived

image quality in the human visual system. Given two images x and y, it is defined as:

$$\frac{(2\mu_x\mu_y + C_1)(2\sigma_{xy} + C_2)}{(\mu_x^2 + \mu_y^2 + C_1)(\sigma_x^2 + \sigma_y^2 + C_2)} \tag{5.4}$$

The coefficients C_1 and C_2 are small constants and they are included just to avoid instability. The value of SSIM is in the range $[-1; 1]$, where 0 corresponds to no similarity, while 1 and -1 correspond to perfect direct or inverse similarity, respectively.

There is no single metric which can adequately discriminate all the different ways that an image can be distorted. The two chosen metrics complement well each other for the task of denoising. The PSNR is more sensitive to additive Gaussian noise than the SSIM, but the latter captures better other forms of corruption, such as blocking artefacts [128].

Z-score normalisation was performed before the two metrics were calculated. The reason was to account for the heterogeneity of the experimental datasets, as well as the fact that noisy images generally have lower signal intensities than the corresponding ground truths. This normalisation was done by first subtracting the mean of the image and then dividing by its standard deviation: $(x - \mu_x)/\sigma_x$.

The combined standard deviation of the metric values across different groups, such as channels or scenes, was calculated as the sum of the mean standard deviation and the standard deviation of the means. For a group of size N, with statistics σ_i and μ_i for all i in N, the total mean μ_T and total standard deviation σ_T are derived in the following way:

$$\sigma_\mathrm{T} = \sqrt{\frac{\sum_{i=1}^{N}\left(\sigma_i^2 + (\mu_i - \mu_\mathrm{T})^2\right)}{N}}, \quad \mu_\mathrm{T} = \frac{\sum_{i=1}^{N}\mu_i}{N} \tag{5.5}$$

5.4 Results

5.4.1 Acquisition of training and evaluation data

The first objective of the experiment was to acquire all the data necessary to train the denoising neural networks and then to evaluate their performance. The network architectures of all of the selected deep-learning methods were self-supervised and therefore they could be trained on the same data which needs to be denoised. It was decided

to acquire data of different dimensions, in order to evaluate as many experimental scenarios as possible and to ensure the denoising algorithms are presented with rich real-world data, comprised of images with spatial and temporal correlations in their signal. The data was divided in two main categories: live time-series of planes (2D+t) and volumes (3D+t), and fixed individual 3D volumes.

Capturing live mRNA data is difficult and it is not really possible to obtain good ground truth data for it. These obstacles were overcome by imaging the live expression of *Drosophila* EB1 protein instead. It localises to the plus tips of microtubules and exhibits great dynamics during their growth and the cytoskeletal reorganisation in general. Therefore, EB1 is an excellent model which provides reliable fast dynamic movement of puncta-like objects.

The live data was acquired by alternating low and high exposures, in order to have one noisy and one relatively clean image at each time point. This resulted in two time-series with roughly the same underlying signal, with one of them being able to act as a reference (ground truth). The low exposure time was set to 10ms, which was the lowest possible value for the microscope. The high exposure time was set to 100ms, which offered a good trade-off between SNR and speed, which is an important parameter because the two exposures need to be taken as quickly as possible, to ensure the scene remains mostly the same. This kind of experiment was performed at five different laser powers in order to evaluate the denoising algorithms at various noise levels. Starting at 20%, below which there was no meaningful signal, the laser power was increased by 20% four times, for a final power of 100%. This set of five experiments was done first with single planes and then with volumes, where a stacks of multiple sections along the z-axis (depth) were acquired for each frame. Table 5.3 summarises the data acquisition settings. Figures 5.2 and 5.3 show maximum projections along the t-axis (time) for the 2D+t and 3D+t datasets, respectively.

The fixed training and evaluation data, consisting of individual 3D volumes, was acquired from a commercial cell slide. This allowed for some degree of standardisation, but most importantly it helped in two distinct ways. On one hand, it enabled the capturing of volumes of higher depth than the live data where speed was a deciding factor in how many section can be obtained. Secondly, for the same reasons, it provided the opportunity to acquire high quality ground truth data, which also happens to be

Table 5.3: Experimental settings for 2D+t and 3D+t data.

Type	Laser power	Exposure time	Frames	Sections	Channels
Low SNR	Level 1: 20% Level 2: 40% Level 3: 60% Level 4: 80% Level 5: 100%	10 ms	2D: 140 3D: 18	2D: 1 3D: 8 (150nm step)	1
High SNR	40%	100 ms			

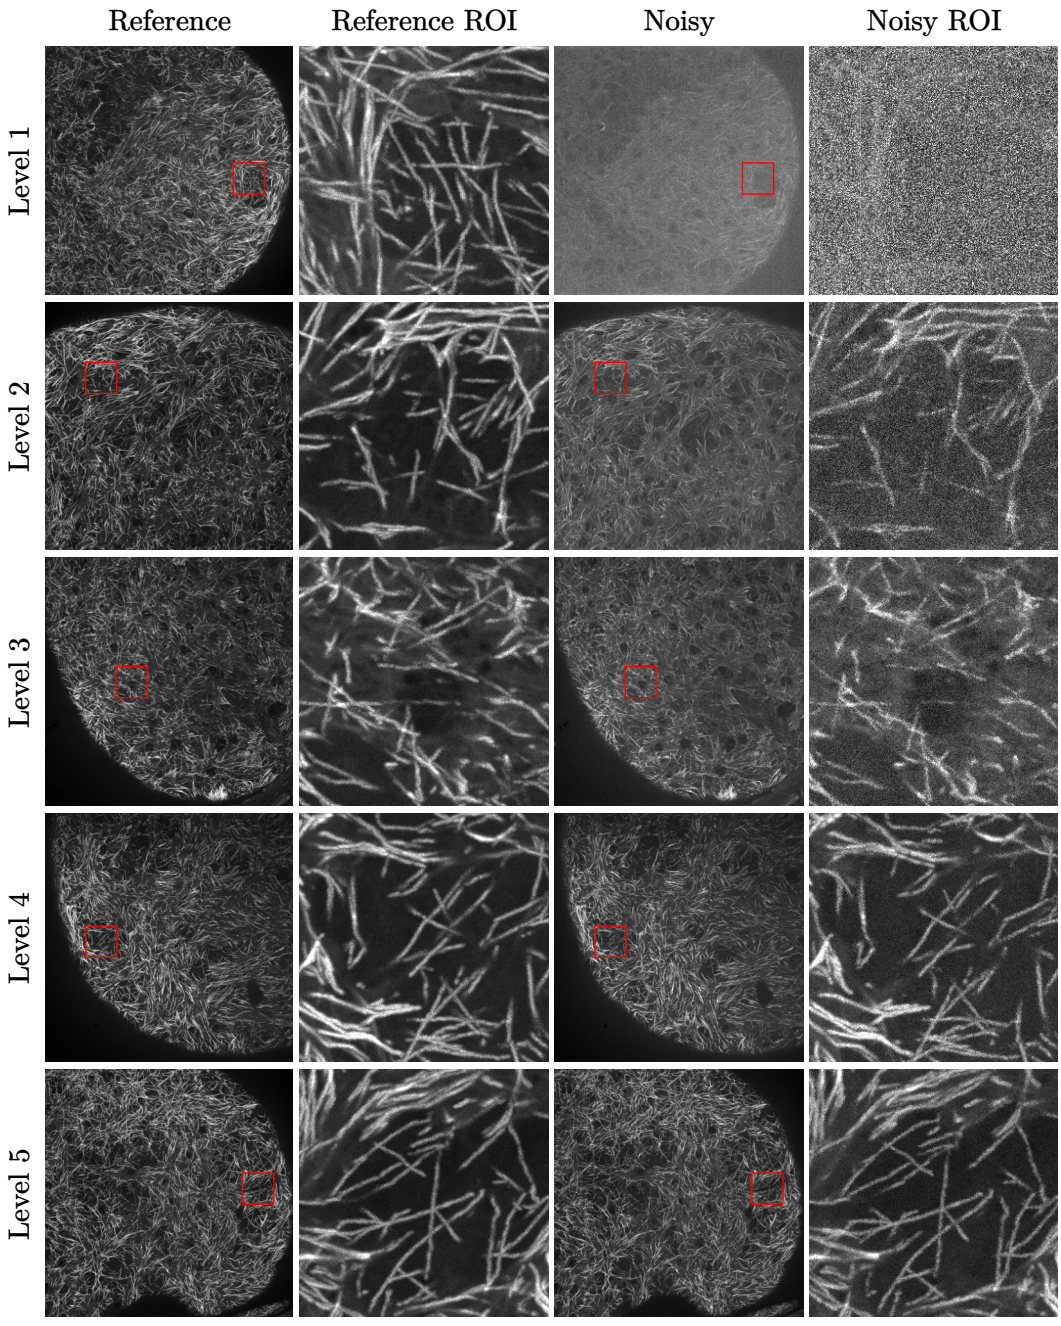


Figure 5.2: Maximum t projections of the 2D+t dataset.

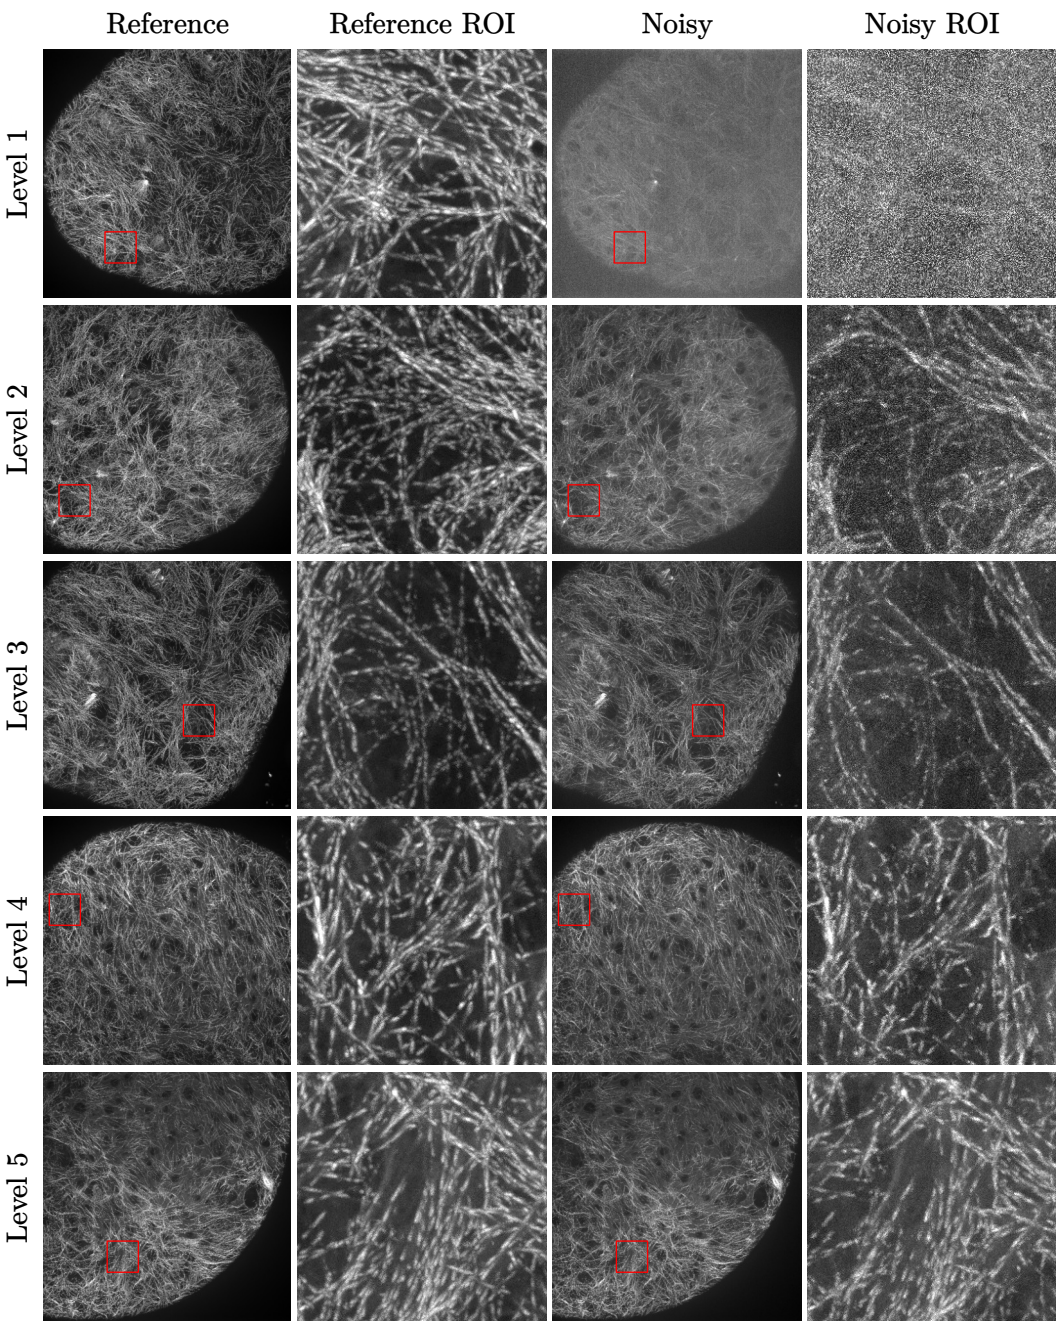


Figure 5.3: Maximum t projections of the 3D+t dataset.

Table 5.4: Experimental settings for 3D data.

Level	Laser power	Exposure time	Frames	Sections	Channels
1	488 nm: 15% 561 nm: 2%	488 nm: 10 ms 561 nm: 10 ms	1	17 (150nm step)	2
2	488 nm: 15% 561 nm: 2%	488 nm: 25 ms 561 nm: 25 ms			
3	488 nm: 17% 561 nm: 4%	488 nm: 25 ms 561 nm: 25 ms			
4	488 nm: 20% 561 nm: 7%	488 nm: 25 ms 561 nm: 25 ms			
5	488 nm: 25% 561 nm: 15%	488 nm: 25 ms 561 nm: 25 ms			
GT	488 nm: 100% 561 nm: 100%	488 nm: 500 ms 561 nm: 500 ms			

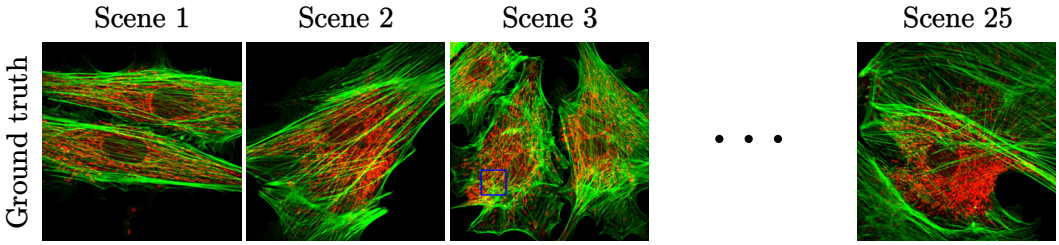
well correlated and aligned, because the sample was not dynamic. The entire dataset consisted of 25 scenes, each containing one or more cells in the field of view. Figure 5.4a shows maximum z projections of the first three and the last scenes.

Each scene was imaged with six different acquisition settings, progressively increasing the light dose, in order to match the format of the 2D+t and 3D+t data. This time, however, there was no constraint on the exposure times, which added extra flexibility in selecting the parameters of each level. The first level was configured to be barely above the noise threshold and the rest of the levels were configured in a way where each would bring a visually noticeable change to the noise statistics. For the red channel, this resulted in roughly doubling of the excitation photon count, whereas in the green channel the relationship was more complex because of non-linearity of the laser line in this power range. Table 5.4 lists all the different acquisition settings and Figure 5.4b gives a particular example, showing the gradation of the noise with each level.

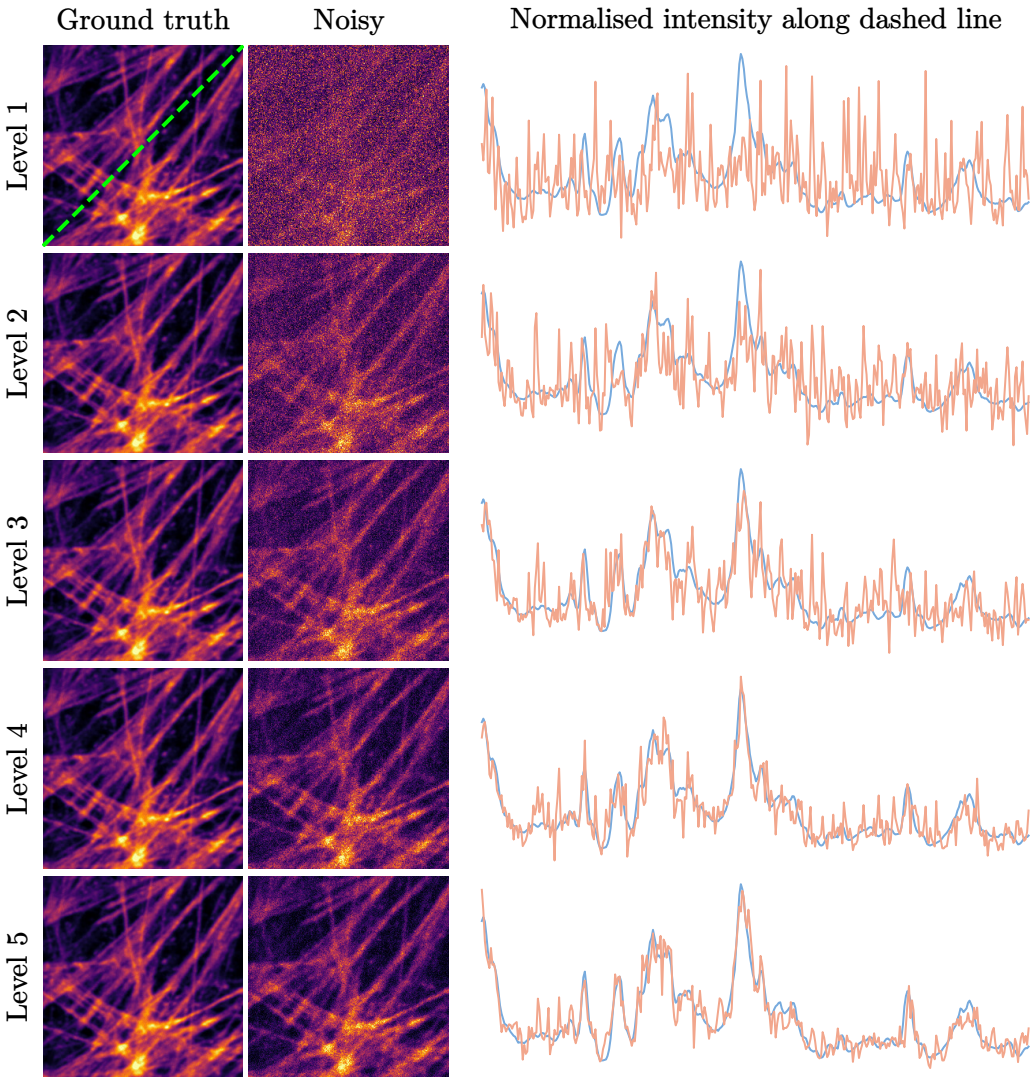
5.4.2 Denoising performance

Introduction

The second objective was to evaluate the performance of the different denoising methods on the acquired data. The PSNR and SSIM metrics were used to quantify the performance. The two metrics can be calculated in multiple ways for multidimensional data. The default (trivial) way is to calculate them for each image individually, then



(a) Maximum z projections.



(b) Noise levels in the actin channel of the region indicated in blue in scene 3.

Figure 5.4: Overview of the 3D data.

to report a statistical measure of all the results, such as the mean. Alternatively, the multidimensional data can first be projected, in the sense of linear algebra, to produce data of lower dimensionality before calculating the metric. Two commonly used projections are average and maximum, which take the average or maximum value along one of the dimensions such as depth or time. A combination of subscript and superscript notation is used to indicate when a metric has been calculated after the data has been projected first. For example, $\text{PSNR}^{\text{T}}_{\text{mean}}$ indicates that all the frames in a time-series have been averaged first, whereas just PSNR is the default case and it is the average of the metric for each frame.

This distinction between metrics calculated with or without projecting first is important because dimensionality reduction decreases the variance of the prediction – one of the two key parameters that characterise the denoising performance, together with bias. The two are analogous to accuracy and precision, and there is usually a trade-off between them. The MSE can be expressed as the sum of the variance and the squared bias, which is how these parameters can influence metrics such as the PSNR.

What this means in practice is that a denoiser, or an estimator in general, with a lower bias could appear less efficient if it has a higher variance, as illustrated in Figure 5.5a. Effectively, the variance hides the true accuracy of the denoiser, which is then revealed by averaging multiple predictions. In the same sense, looking at the performance of a denoiser on individual images does not give a full description of its capabilities.

This phenomenon is particularly noticeable when the signal is low and close to the noise threshold, as illustrated in Figure 5.5b for level 1 of the 2D+t dataset. N2V has lower variance in its predictions, compared to PN2V, which is difficult to show with individual images, but can be noticed in the noisy single frames in the left half of the middle row. N2V appears cloudy and there is little contrast, with most of the intensities concentrated in a small range of values. On the other hand, PN2V shows more variation and contrast, with clear dark areas. If multiple frames were inspected, this variation would appear as intensity fluctuations, reflecting the uncertainty of the denoiser. The metric scores of N2V are therefore better for single frames, 25.37 and 0.342 compared to the 24.30 and 0.303 of PN2V. However, by averaging all the frames in the time-series it becomes obvious that N2V has a much higher bias, appearing as a

glow in the areas with signal. In this case PN2V is much closer to the reference image, with scores of 29.29 for PSNR and 0.669 for SSIM.

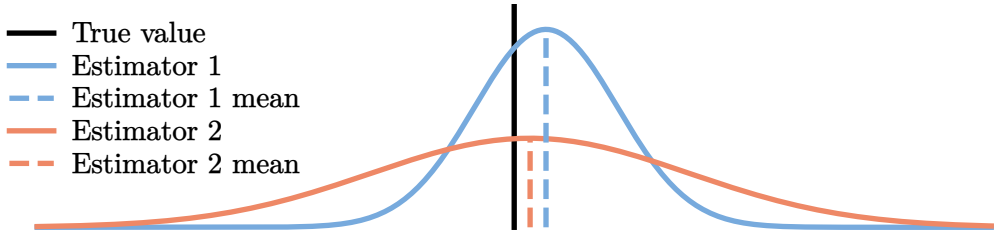
The goal of this experiment was to evaluate the performance of various denoising methods on real microscopy data, which is multidimensional most of the time, and therefore the metrics were calculated for both individual images and projected images. Full breakdown of the metric scores is given in Appendix C.

Live time-series data

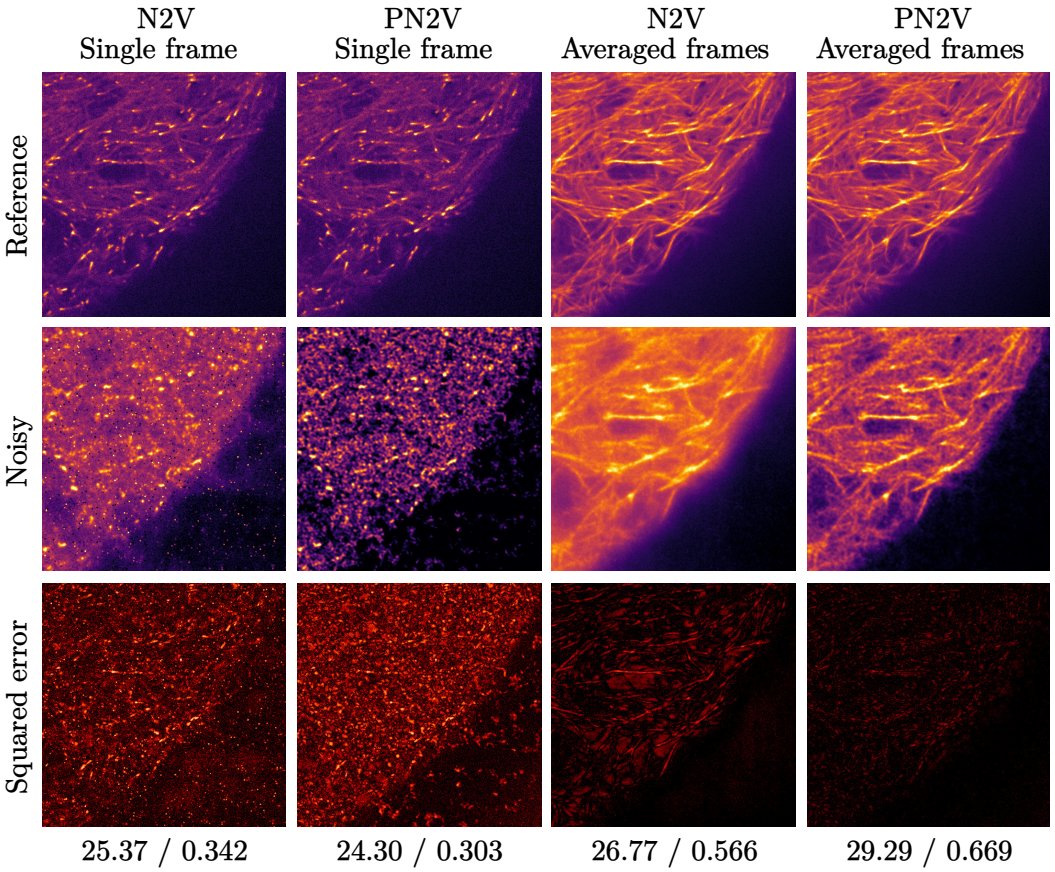
The quantitative and qualitative results for the 2D+t and 3D+t dataset are shown in Figures 5.6, 5.7, 5.8, and 5.9. Maximum projection was used for dimensionality reduction because it was more sensible for this kind of sample where it is of interest to track bright spots across space and time. The images in the qualitative results were cropped 32×32 pixel regions, in different locations for each level.

Starting with the 2D+t data, at level 1 none of the methods did particularly well. nD-SAFIR and other patch-based algorithms generally struggle with high noise because it impairs their ability to select similar patches. Likewise, the downsampling approach of N2F probably destroyed the little signal there was in the data. From the remaining three, N2V was better individually, because of its lower variance, but its bias was worse than the other two. PN2V and N2N had very similar performance, each excelling at one of the metrics. All of the methods performed well from level 2 onwards. One of the limiting aspects of these datasets was the absence of good ground truth data. As a consequence, at level 2 to 5 all of the methods produced results which appear visually better than the reference. nD-SAFIR showed the best performance after projection in this noise range, followed closely by the deep-learning methods.

Regarding the 3D+t data, N2V showed dominant performance, scoring first in 17 out of the 20 tests. It produced a meaningful estimation even at level 1, where nD-SAFIR and N2F failed again, and the other two deep-learning methods struggled. The excellent performance of N2V was probably caused by the fact that the network used 3D convolutions as part of its architecture and therefore it probably made better use of the multidimensionality of the data. nD-SAFIR performed worse at all levels after projection, creating pronounced artefacts with abrupt intensity transitions.



(a) Comparison of the distributions of two biased estimators with different means and variances. Estimator 2 has a higher MSE than estimator 1, but it is more accurate on expectation.



(b) Difference between calculating metrics for a single frame or for the average of all frames. The images show a cropped region from level 1 of the 2D+t dataset. The metrics are reported below each error image, in the format PSNR/SSIM. N2V shows better performance for single frames but worse performance on average, compared to PN2V.

Figure 5.5: Overview of the 3D data.

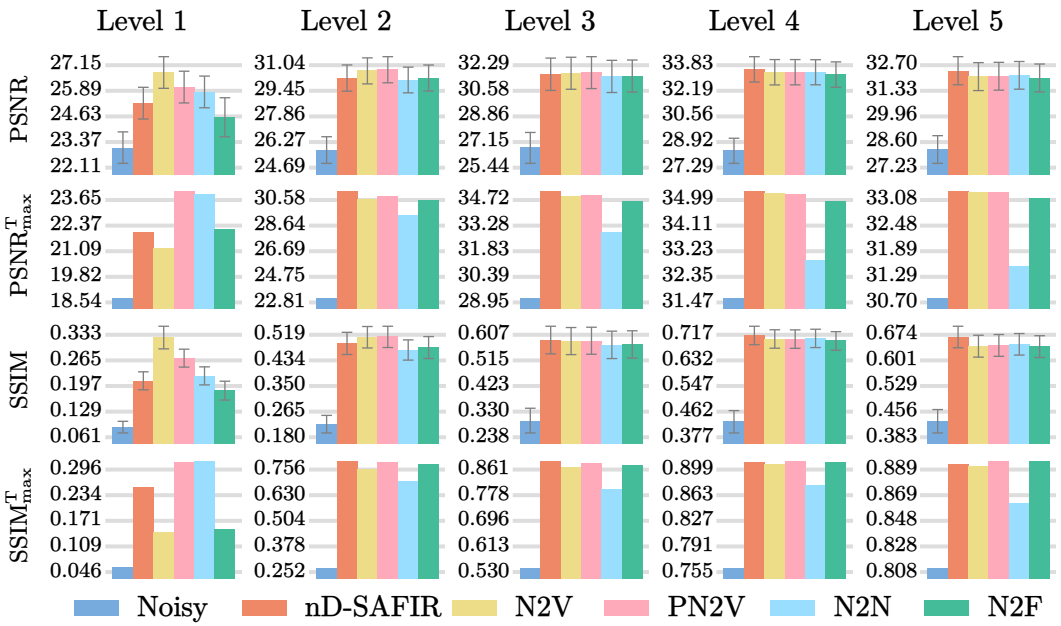


Figure 5.6: Quantitative results for the 2D+t dataset. Error bars indicate $\pm\sigma$ across all frames.

Fixed 3D data

The results are shown in Figures 5.10 and 5.11. This time, average projections were used for dimensionality reduction, because they portray better the overall structure of the cells. The metrics were calculated as an average across all 25 scenes and both of the two channels. In a similar way, the standard deviations are the combined values across all scenes and channels, calculated with equation 5.5. The qualitative results show both channels, with 32×32 px cropped regions from scene 8, with different locations for each channel.

The performance of nD-SAFIR was poor in all tests, but this was expected because the algorithm was designed to work with time-series data, where it can find correlated patches across multiple frames. The algorithm is severely limited with a single 3D volume.

All the denoisers failed to produce meaningful results for channel 1 in the region in Figure 5.11 at level 1, but the bright circular feature in channel 2 was restored well at all noise levels. PN2V did a particularly good job on it, and overall on the tests after projection, up until levels 4 and 5 where N2F outperforms it in three out of the four tests.

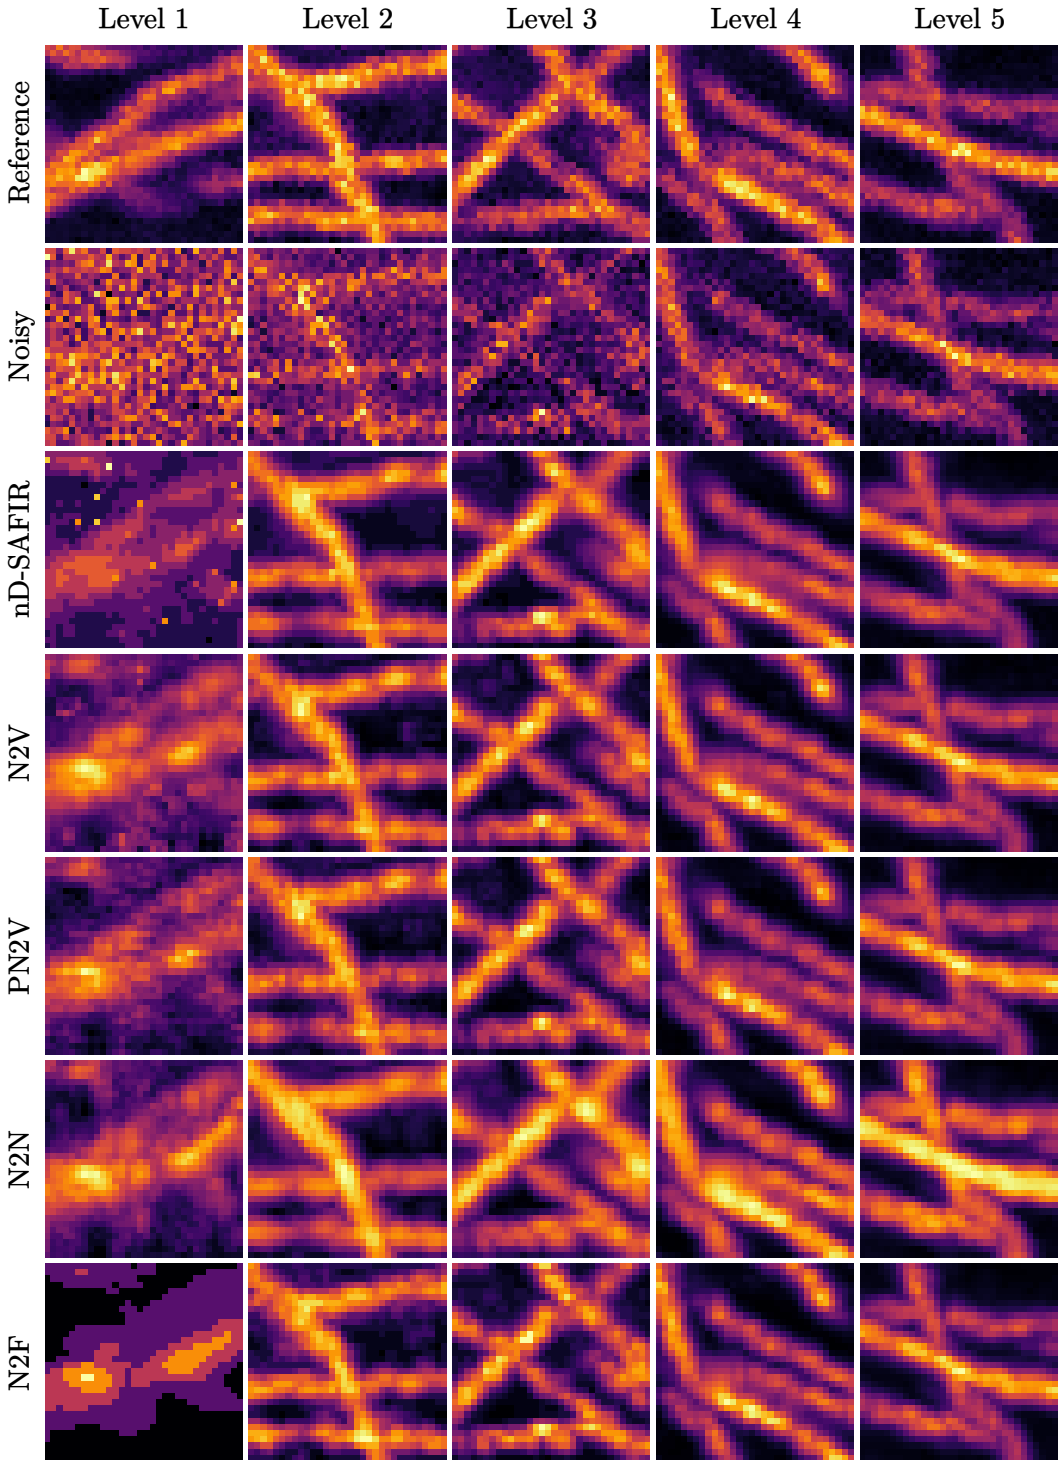


Figure 5.7: Qualitative results for the 2D+t dataset.

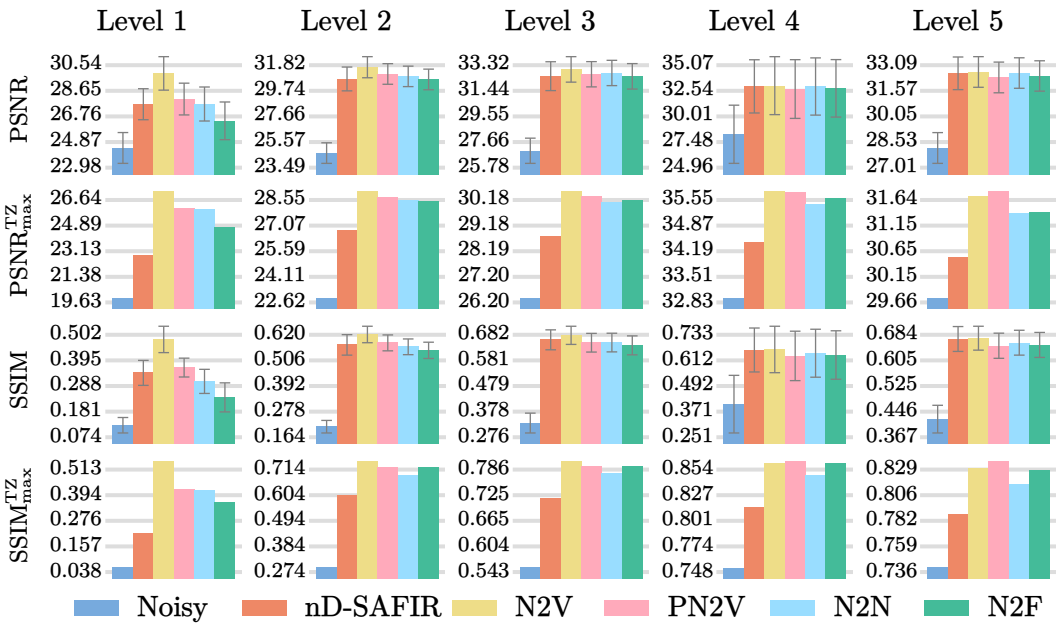


Figure 5.8: Quantitative results for the 3D+t dataset. Error bars indicate $\pm\sigma$ across all sections and frames.

Running times

The operation of the denoising methods was broadly split into two phases: training and prediction. nD-SAFIR and N2F did not require training, and those were making predictions directly, without the need of a pre-trained model. The remaining three methods did require training, with N2N being a special case where the pre-trained models were obtained from elsewhere. Figure 5.12 shows approximate training and prediction times, achieved on the hardware described in section 5.3.2. No statistics were recorded for several reasons. Many of the methods were deterministic to a large extent, and except for a little bit of overhead their processing time were fairly precise, for the same amount of training or prediction data. An exception to this was N2F, whose validation strategy created some ambiguity in the timings. Furthermore, all the software was ran on a shared computer, whose overall load and usage had an impact on the processing times.

5.5 Discussion and conclusions

In this chapter I described an experiment which compared the performance of several denoising methods, both traditional and deep-learning ones, with the aim of identifying

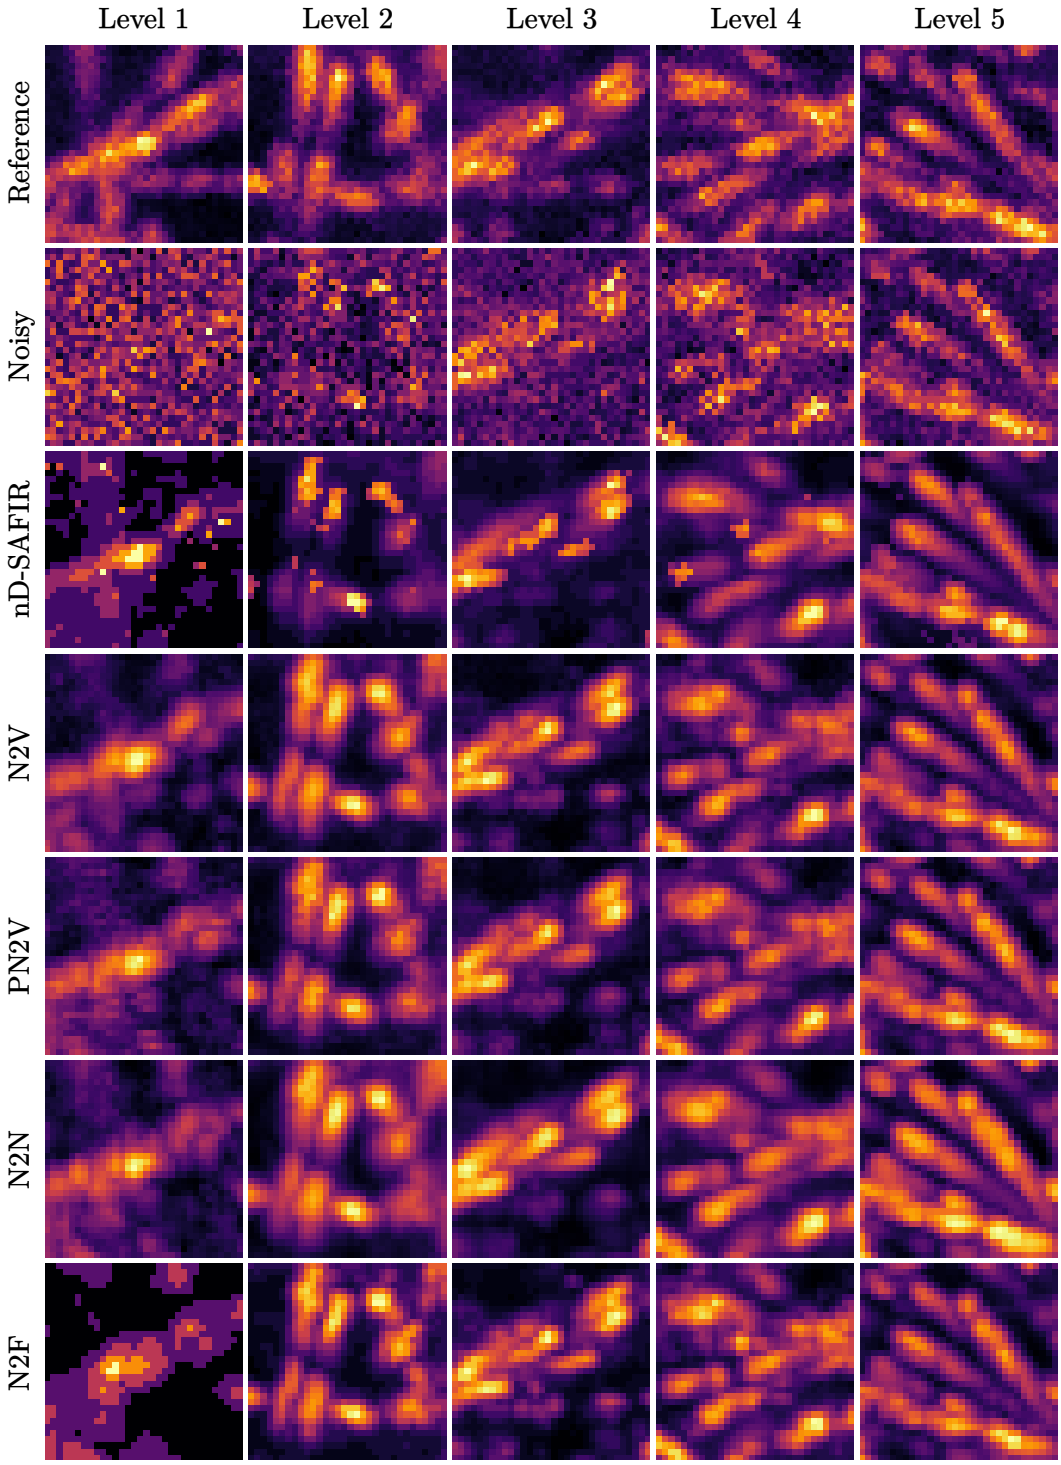


Figure 5.9: Qualitative results for the 3D+t dataset.

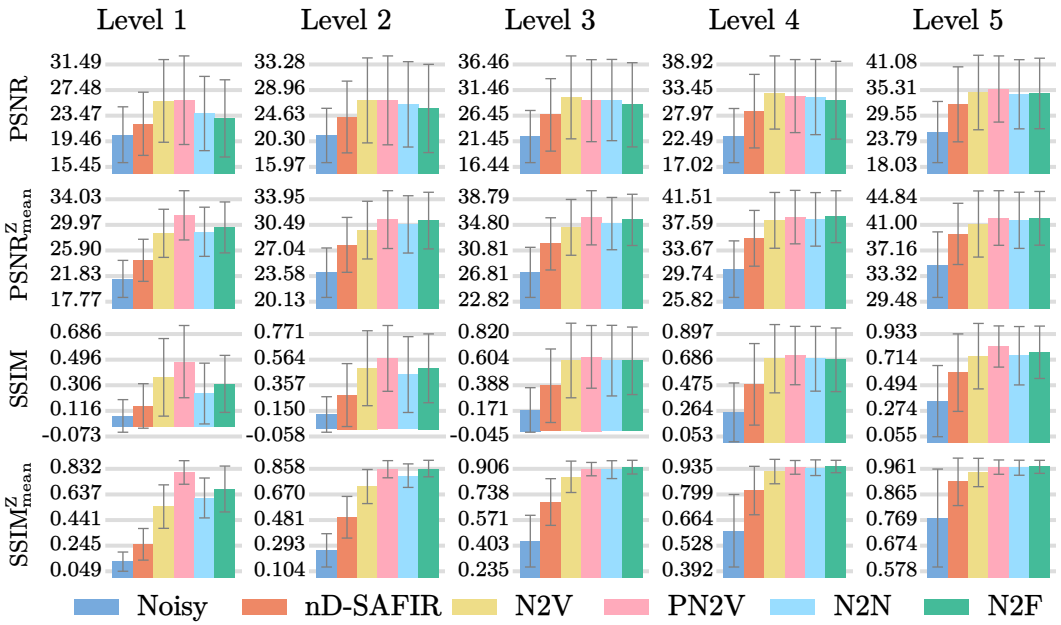


Figure 5.10: Quantitative results for the 3D dataset. The metrics were calculated as usual separately for each scene and channel. They were then combined into a single statistic using equation 5.5, which is why standard deviations were included for $\text{PSNR}^{\text{Z}}_{\text{mean}}$ and $\text{SSIM}^{\text{Z}}_{\text{mean}}$. Error bars indicate $\pm\sigma$.

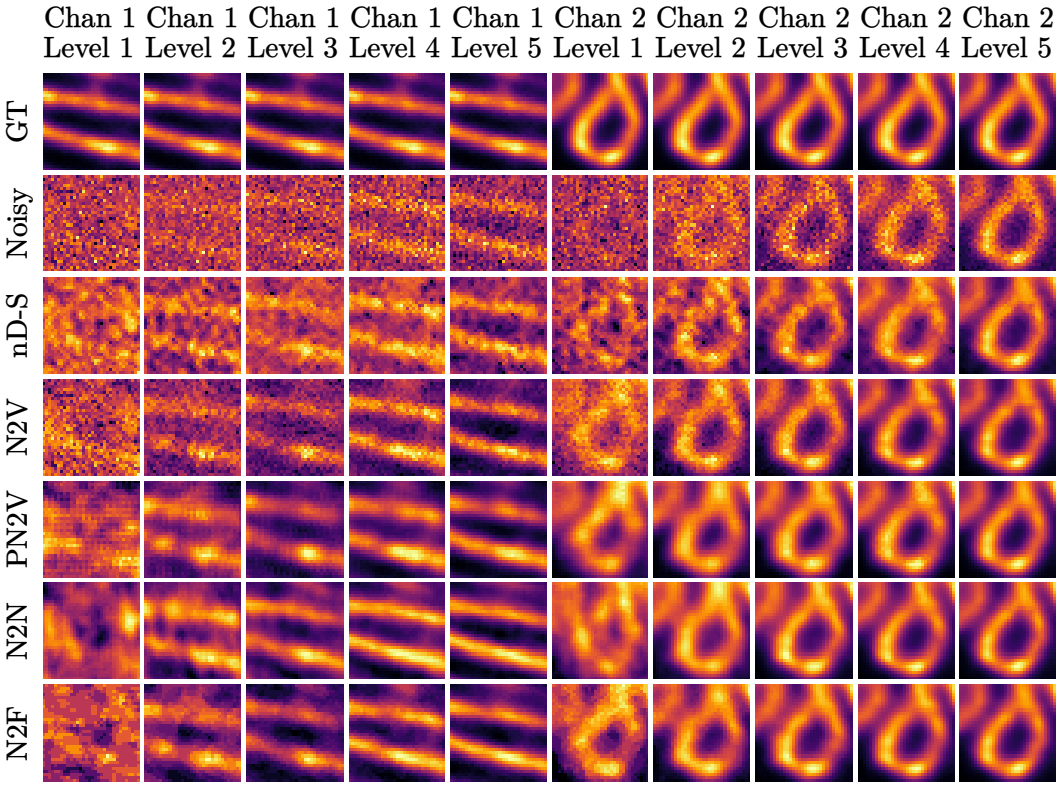


Figure 5.11: Qualitative results for scene 8 of the 3D dataset. Chan: channel, GT: ground truth, nD-S: nD-SAFIR.

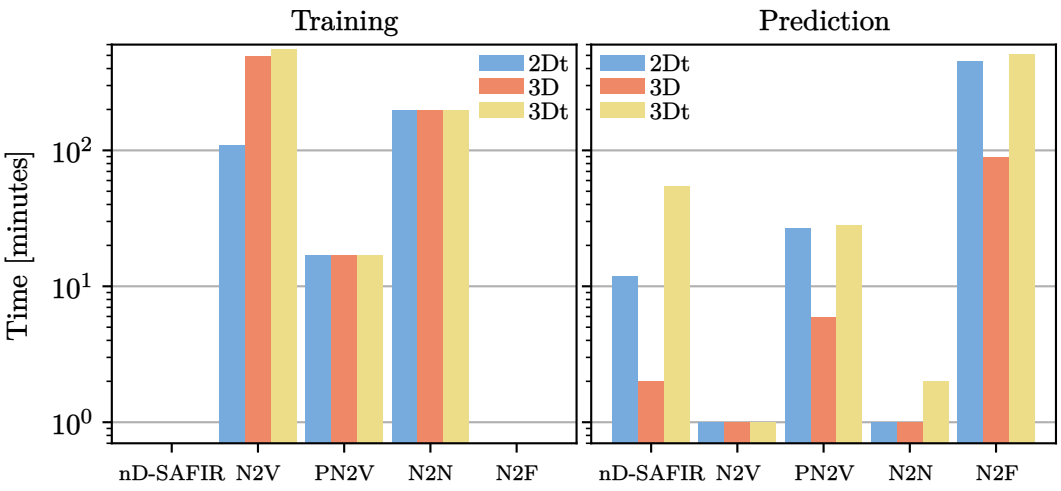


Figure 5.12: Approximate training and prediction times for all methods, rounded to a minute. The training time for PN2V included the training of the noise model. The training time for N2N was taken from [112].

better ways of processing fluorescence image data, with an emphasis on live microscopy. I set out three objectives at the start of the experiment and so far I have described the successful completion of the first two, namely the acquisition of real microscopy training and evaluation data, and the assessment of a wide range of denoising methods. In this section I want to discuss and highlight the key results outlined in section 5.4, completing the third and final objective of the experiment.

The most important message is that the deep-learning methods perform better. They only challenged on the 2D+t dataset and even there the difference seems mostly technical, as all of the methods seem to restore the structures well, from level 2 onwards. However, the same cannot be said for the other datasets, where nD-SAFIR outright fails to produce satisfactory results. In terms of processing times, some deep-learning methods, such as the tested PN2V, are quicker than traditional methods, even if training is required for every dataset. If models are reused, the deep-learning methods can be considered orders of magnitude faster.

On the subject of reusing of models, another important result is the excellent performance of the pre-trained N2N network. It was trained on an entirely different data, microscope, and noise statistics, and yet it came second on many of the tests. This suggests that deep-learning networks can generalise really well and that they could be prepared in advance for rapid live denoising. It also poses the question whether training is best done for a specific noise level or whether the training data should be inclusive

of all noise levels.

The superior performance of the deep-learning methods was achieved without the need to acquire special or voluminous training data. They were all trained on the same data that they then denoised, and even in the case of the 3D data, where multiple scenes were captured, this was more of a safety measure and it may not have been necessary. Therefore, contrary to popular belief, neural networks do not require dedicated training data to perform well, and this false requirement should not be a deterrent to their adoption. Some of the networks, like N2V and N2F, also provide software frameworks that streamline their use and make their integration fairly straightforward.

The excellent performance of N2V on the 3D+t dataset hints at the potential importance of the architecture of the neural network, especially in terms of dimensionality of the operations. Using 3D convolutional blocks and working on volumes instead of images is likely better for data which has rich spatial information, such as *Drosophila* larval brains. N2V performance was less impressive on the 3D dataset, but this could potentially be explained by the fact that the dataset was collected from cell culture sample, which is thin and the majority of its signal is contained within just a handful of sections.

The multidimensionality of the data is beneficial not just for the training, where networks such as N2V can make use of the extra information and all networks equally benefit from the sheer increased size of available images, but it is also beneficial for the prediction stage where it reduces the variance of the estimate and emphasises the true accuracy of the network. This is particularly important for low-light imaging where the uncertainty of the prediction can be removed by dimensionality reduction.

Finally, it is not optimal to recommend a single denoising method. This decision needs to be informed by the level of the noise and the type of the data, amongst other factors. Perhaps the only consistent trend was that deep-learning methods were better. For multidimensional data with rich spatial information N2V performed the best. For other types of data, PN2V excelled when the noise was the strongest and it did really well otherwise. For 2D time-series with weaker noise, traditional methods such as nD-SAFIR are still a viable option.

Chapter 6

Developing a motility assay for mRNA dynamics

6.1 Introduction

The nucleus is viewed as the core of the cell, but the structure and the function of the cell is not radially symmetrical around the nucleus. Cells are not perfectly round and compact. They usually have asymmetric shape, often exhibiting strong compartmentalisation and specialisation. One of the most striking examples are the cells of the nervous system: neurons and glia. Their long cytoplasmic projections have more than just structural and mechanical function. They participate actively in the interaction with other cells and in the transmission of signals.

Remote sites of the cytoplasmic projections, far away from the nucleus, have unique requirements for molecular signals and resources, most prominently protein. However, it is not completely understood how such remote protein gradients and patterns are generated. Emerging evidence hints at the importance of mRNA transport and local translation, and post-transcriptional regulation in general. Indeed, the complex intracellular transport machinery can traffick mRNA molecules, as well as various organelles and vesicles. Amongst them, however, mRNA plays a much more fundamental role, and it is connected to many forms of disease, which makes it a prime candidate for study.

Early investigations of mRNA transport used dyes which were not specific, for example labelling all nucleic acids [129]. This limited severely the scope of the studies. The field was greatly transformed by the development of the MS2-MCP system. Since

then it has been used in various organisms and with various samples, significantly improving our understanding of mRNA transport and mRNA biology in general [130]. Similar systems and entirely different techniques for live visualisation of mRNA have been developed, too, but they are generally less mature and less optimised [131, 132].

Despite having the means to label individual mRNA molecules, in vivo live imaging has remained a technically challenging task [133]. One of the main problems is that the tissues of living organisms are optically scattering and prevent light from penetrating through. This is usually addressed by using organisms whose nervous system can be reached with microscopic methods that also possess the necessary spatial and temporal resolution.

Drosophila stands out as a particularly good candidate for the study of mRNA transport. Historically, many landmark discoveries in the field were made precisely in this organism [130]. The MS2-MCP system has been used previous to study mRNA transport in dendritic arborization neurons, part of the CNS, close to the surface of intact *Drosophila* larvae [134]. However, this type of neurons form complex networks where it becomes difficult to disentangle the different interactions.

In this chapter I describe the development of an improved motility assay that uses the *Drosophila* PNS instead. The advantage of this approach is that it provides a simplified system where the interaction of a single motor neuron and its glial cells can be studied in the context of the synapses it forms with a muscle fibre. These kind of synapses are glutamatergic, which makes them a good model for neuronal processes in the mammalian central brain. The assay uses the MS2-MCP system to label single mRNA molecules. Together with the powerful genetic toolkit of *Drosophila* and its ease of use, it offers an unmatched potential for probing the post-transcriptional regulation of genes.

I demonstrate the potential of the assay by looking at the mRNA dynamics of *shot* in glia. This choice was based on the idea that local translation in the peripheral cytoplasm is used to regulate the function of not only the glial cell itself but also that of the adjacent motorneuron synapse. The Davis Lab recently completed comparative analysis of the spatial distribution and expression of 200 mRNA and the proteins they encode [135]. A large proportion of those transcripts were revealed to be present at the NMJ, in the distal ends of the axonal projections, with 18 of them present at

the glial periphery. Francesca Robertson from the lab had tagged one of these 18 transcripts endogenously with 24 MS2 loops (unpublished work) and this one was *shot*. It was later found out that a UAS-RNAi knockdown, driven by the pan-glial repo-Gal4, caused morphological defects to the glia at the NMJ and it affected the synaptic morphology of the adjacent otherwise wild-type motoneuron axon [136]. Interestingly, the RNAi knockdown of *shot* did not affect the plasticity of the synapses in a potassium stimulation assay.

shot encodes a member of the spectraplakin family and it is the only such representative in *Drosophila*. These are proteins which can connect all three types of filaments of the cytoskeleton and as such play an important role in its regulation. A patient study discovered that breakpoints which interrupt the human *dystonin* gene, an orthologue of *shot*, lead to neurodevelopmental disorders [137]. This makes it a particularly exciting candidate for investigation.

6.2 Aims and objectives

The overall aim was to build the fundamentals of a methodology for the live observation of mRNA dynamics, both in neurons and in glia. The idea was to develop a motility assay where any arbitrary gene, expressed in the *Drosophila* PNS, can be assessed in a robust and quantifiable way by deriving particle tracks, which can then be used to answer questions about mRNA transport and post-transcriptional regulation in general. This followed up work started by Francesca Robertson, in the Davis Lab, who had already created MS2 lines for several genes of interest.

As discussed in the Statement of Mitigation, the microscope with adaptive optics described in Chapter 4 had to be decommissioned before it could be used to acquire live mRNA image data. Therefore, all the motility assay data presented in this chapter was acquired on different instruments. As a result, there is no aberration correction or remote focusing. To address this shortcoming, a secondary aim was to create simulated synthetic data, closely mimicking the mRNA behaviour observed *in vivo*. Different versions of this data can then be generated, with and without aberrations and noise, to demonstrate the necessity for aberration correction and image denoising.

The two aims were approached with the following set of objectives:

- Drive expression of MCP in the glial cells in the PNS, using the Gal4/UAS system, and acquire live image data of moving shot mRNA particles
- Tune and train an image denoising neural network, and develop an image processing pipeline for extracting mRNA particle tracks from fluorescence image data
- Investigate if the acquired *shot* track data can answer questions about the transport and post-transcriptional regulation of the gene
- Demonstrate the benefits of adaptive optics and image denoising to the motility assay, by using simulated data which mimics closely real mRNA dynamics and the image corruption caused by the optical system

6.3 Methods and materials

6.3.1 Fly strains and sample preparation

The samples were prepared from a cross of the following two lines:

$$; \frac{shotMS2v5x24}{CyOGFP}; \frac{repo-Gal4}{Tm6B,tb}$$

$$; \frac{UAS-mcd8-mCherry, UAS-MCP-GFP}{CyOGFP}; \frac{tubgal80ts}{Tm6B,tb}$$

Both of them were part of the Davis Lab's stocks and they were made by Francesca Robertson. Flies from the two lines were crossed and kept at 25 °C for 4 days, then shifted to 30 °C for 24 h, and then back to 25 °C for another 24 h. L3 larvae were collected and dissected in Schneider's medium at room temperature. Brains were carefully removed to keep their axonal projections intact, and mounted in Schneider's medium on a glass slide under a #1.5 glass coverslip. The samples were imaged for about 30 minutes before being discarded.

6.3.2 Microscopy

The imaging was done on a custom widefield microscope, powered by the Cockpit software described in chapter 3. The system consisted of an Olympus IX83 inverted

microscope, equipped with Photometrics Kinetix camera. The light was delivered by Lumencor Spectra III and a standard GFP filter cube was used to filter the fluorescent emission. The synchronisation of the experiments was done with a RedPitaya 125-14 v1.0.

The image data was acquired with Olympus UPLSAPO100XS objective lens for a voxel size of $0.065 \times 0.065 \times 0.3\,\mu m$. The total height of the volumes was $5.4\,\mu m$. The acquisition speed was one volume per second, for a total duration of 120 s. The exposure time was 48 ms and the GFP channel of the Spectra III was operated at 25% power, equivalent to about 125 mW of input power.

6.3.3 Image processing

The images were denoised with Noise2Void version 1 with MSE loss and `uniform_withCP` manipulator. The network was configured in 3D mode with patch size of $16 \times 32 \times 32$ pixels. The kernel size was $3 \times 3 \times 3$ pixels and 2% of the pixels were masked. The training was done for 200 epochs with a rate of 0.0004.

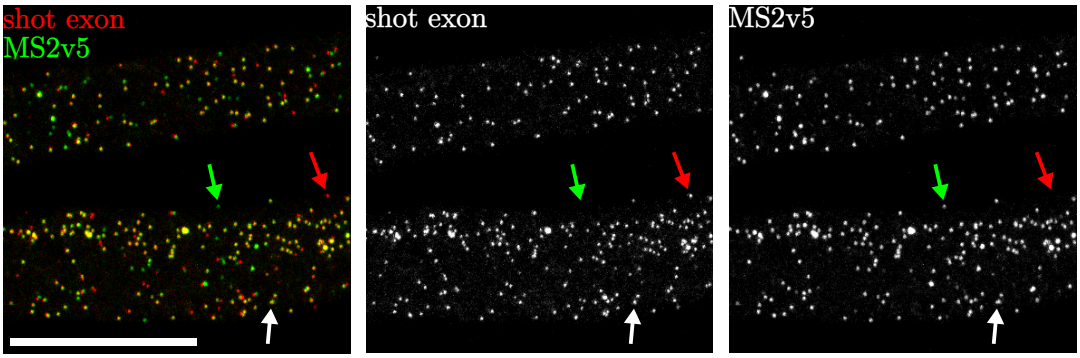
After denoising, background subtraction was performed in Python, using the scikit-image library [138]. The images were then loaded in ImageJ [139] and the tracking of particles was done in the same software with the TrackMate plugin [140].

6.4 Results

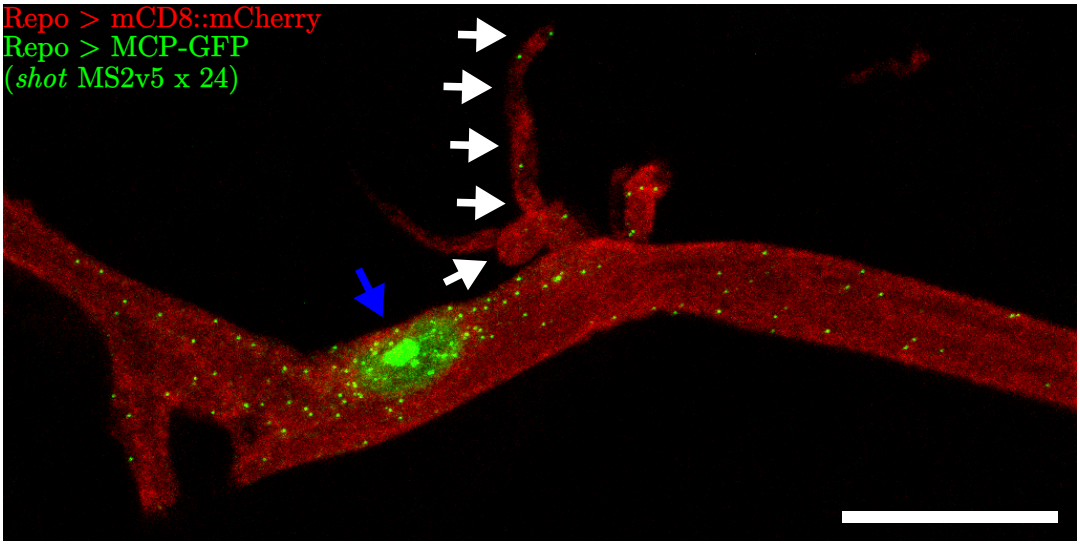
6.4.1 Labelling *in vivo* live single mRNA molecules in a tissue-specific manner

The first task was to establish the labelling technique for the motility assay, which was based on the MS2-MCP system. smFISH was performed by simultaneously using a probe that targets the largest exon in *shot*, which to my knowledge is present in all isoforms and therefore should label all *shot* mRNA transcripts, and a probe that targets just the MS2 loops (sequences are listed in Appendix D). This was done in fixed L3 *Drosophila* larval fillets and image data was acquired along axonal projections. A region of two nearly parallel such axon bundles in shown in Figure 6.1a.

The smFISH spots for each of the probes were counted in 3D and then matched to determine the overlap between the two sets. There were 223 co-localised spots, with



(a) Co-localisation of two smFISH probes in L3 larval axons. Images show maximum projections. Green and red arrows indicate spots which are unique for each probe and do not co-localise. The white arrow points to a spot which is incorrectly detected only for the MS2v5 probe. Scale bar is 20 µm.



(b) Labelling of *shot* mRNA in glial cells. The image shows maximum projection of axons in fixed L3 larva. The glial ensheathing around the axon bundle is labelled in red, while the *shot* transcripts are labelled in green. The blue arrow indicates a nucleus and the white arrows indicate a branching projection which goes towards a neuromuscular junction. Scale bar is 20 µm.

Figure 6.1: Validation of the *shot* MS2-MCP line.

57 and 40 unique spots for the exon and MS2 probes, respectively. This is equal to a co-localisation rate of 69.7%. Example of unique spots are indicated with green and red arrows in Figure 6.1a, while the white arrow indicates an example of a spot which is falsely detected in only one of the channels.

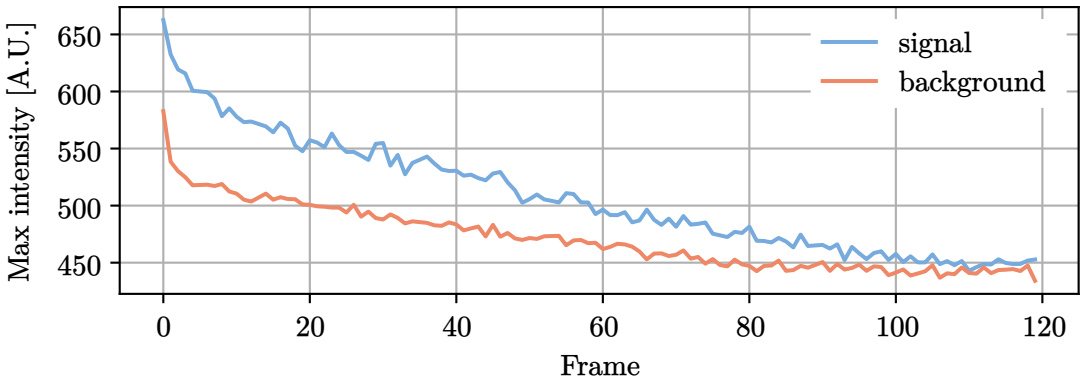
The high co-localisation rate showed that the MS2 loops were correctly inserted and that *shot* transcripts are both present and abundant in the PNS, meaning that there were no obvious problems caused by the insertion. To further validate the mutant fly line, expression of the MCP was driven in glial cells by using the *repo* panglial driver. Again, L3 larval fillets were fixed and images were taken along axonal projections. Figure 6.1b shows one such projection which branches off to make contact with the muscles below and to form synapses. This branch is indicated with white arrows.

There is strong fluorescence signal, showing *shot* expression both in the nucleus (indicated with a blue arrow) and in the cytoplasm. The presence of *shot* transcripts in the branch is a clear indication that such molecules reach the far ends of the cytoplasm, near the synapses, validating the suitability of both the gene and the labelling system for the motility assay.

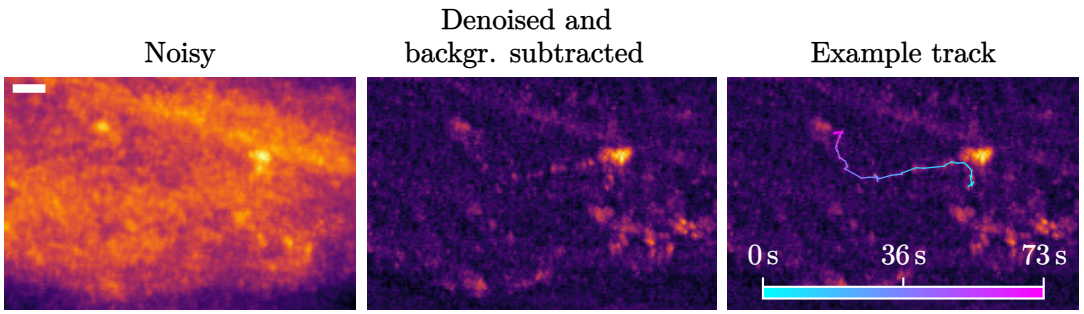
6.4.2 Image processing pipeline

After establishing that the MS2-MCP system can be successfully driven in the glial cells of axonal projections, the next step was to create an imaging pipeline that can produce mRNA particle tracks. It was decided to use a widefield microscope for data acquisition, because of its speed and gentleness. The diameter of an axon bundles varies, but it is usually in the range 5 µm to 10 µm. Therefore multiple Z sections needed to be taken for each time point, in order to cover the entire volume of the bundle. As a result, it was difficult to capture more than 1 frame every second, even with conservative exposure settings of 48 ms and 25% light power (around 48 mW of input power). Photo-bleaching was a significant problem, as shown in Figure 6.2a. The signal became increasingly difficult to detect after about 60 frames, corresponding to 60 seconds, and it was almost completely gone by frame 100.

The acquired image data was denoised with Noise2Void, because it was the best performing 3D+t method in chapter 5. The neural network was trained on a subset of the data. After denoising, background subtraction was used to improve the contrast,



(a) Photobleaching curves for MCP-GFP. The signal intensities were taken over the area of one of the detected spots, while for the background an arbitrary nearby area of the same size was used.



(b) Demonstration of the three main image processing steps. Raw image data was first denoised, then underwent background subtraction, and then particle tracks were derived. The three panels show maximum intensity projections along the Z and time axes. Scale bar is 1 µm.

Figure 6.2: Image acquisition and processing pipeline for live *shot* mRNA data.

and then particle tracking software was used to automatically obtain *shot* mRNA tracks. These three steps are shown in Figure 6.2b.

6.4.3 Particle track analysis

Once a working pipeline was established, five different *shot* datasets were processed and analysed. For ease of use, the data was acquired in explanted brains, where the axonal projections were cut and loosely hanging from the brain, instead of being tethered to the body wall muscles of the larva. This made it possible to image the sample on an inverted microscope, through a coverslip, instead of requiring an upright microscope with a dipping objective lens, commonly employed with larval fillet samples.

Each of the five datasets corresponded to an image volume that contained multiple particles. The total number of tracks derived from all five datasets was 701. However, a large proportion of this amount were short noisy tracks. To address this problem, tracks with less than 10 spots in them were filtered out, leaving 149 tracks. The tracking software, TrackMate, provided many features for each track, and the following list was used:

- Number of spots
- Track displacement – the distance between the first spot and the last spot in the track
- Track speed – mean, standard deviation, minimum, and maximum of the track's speed (4 features in total)
- Maximum distance travelled – the maximum distance between any two spots in the track
- Mean directional change rate – mean angular velocity of the track

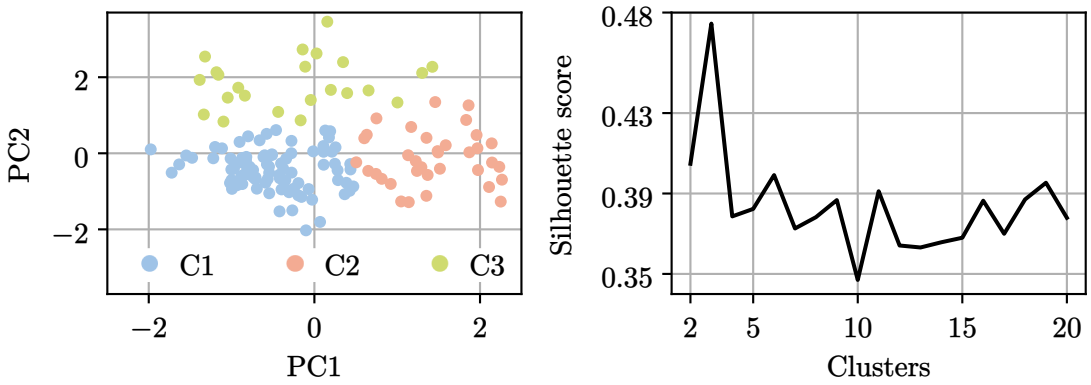
The eight features were reduced to two, using principal component analysis (PCA). The projected data is shown in Figure 6.3a. K-means clustering was used to partition the data, and effectively to determine different classes of particle motion. The number of clusters was determined by using calculating the silhouette scores for different amounts and then selecting the one with highest score, which in this case was three clusters, labelled C1 to C3.

The tracks for all the five datasets, coloured according to their cluster, are shown in Figure 6.3b, together with a selection of 8 tracks, chosen for their interesting character. There were mostly three types of motion: confined diffusion, more loose but still restricted diffusion, and directed motion. The tightly confined motion was really well captured by cluster C2, which also included a couple of tracks that had bursts of directed motion, indicated in the selected tracks (Figure 6.3b) with the pink arrows and the bottom grey arrow. Because they spent the majority of their time in confined motion, the features offered by TrackMate were not very suitable to make them stand out, and therefore they were difficult to separate into their own cluster.

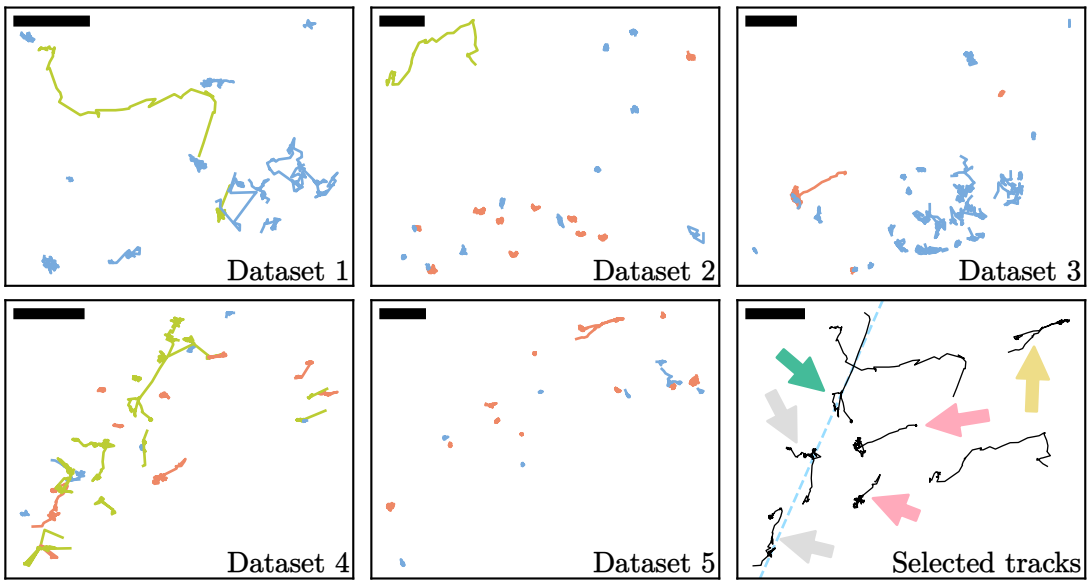
The particles that diffused more loosely and showed more oscillation, yet remained restricted, were captured well by cluster C1. The two areas where all the C1 tracks were concentrated, the bottom right corners of datasets 1 and 3, were actually nuclei. Spots in these areas appeared and disappeared so fast between frames, that it was difficult to say if this type of motion was genuine or just an artefact of insufficient temporal resolution and noise.

There were also three tracks that exhibited long directed motion. Two of them came from datasets 1 and 2, and are shown amongst the selected tracks without any arrow indication. The third one came from dataset 4 and it is indicated with a green arrow. The first two are part of cluster 3, which also included several other tracks that happened to have bursts of linear motion, characterised by low angular velocity, high max speed, and moderate to long track displacement.

Finally, it is interesting to note that there were two instances of directed motion with opposite directions. The first instance was three tracks from dataset 4 that were positioned on a common axis, yet one of them moved in the opposite direction from the other two. They are shown in the selected tracks with grey (travelled downwards) and green (travelled upwards) arrows, and the common axis is drawn with a dashed line. A second example of similar behaviour was the track indicated with a yellow arrow, which came from dataset 5. This track first started on the left side with a long period of tightly confined motion, then it travelled directionally and quickly to the right. There, it experienced a second period of long tightly confined motion, followed by another burst of directed motion in the opposite direction, bringing the final position a little beyond where the track had originally started.



(a) Dimensionality reduction and clustering. Left: projection of the data along the principal components. Right: plot of silhouette scores for different numbers of clusters.



(b) Particle tracks. Colours correspond to clusters. The bottom right panel shows interesting tracks. The grey and green arrows indicate tracks from the same dataset, positioned along common axis (dashed line), yet with opposite direction of motion (grey downwards, green upwards). The pink arrows indicate tracks that switch between directed and restricted motion. The track indicated with the yellow arrow does a round trip, exhibiting directed motion in opposite directions. The remaining two tracks just show long direction motion. Scale bar is 2 µm.

Figure 6.3: Particle track analysis for five datasets of live *shot* mRNA.

6.4.4 Simulation of optical aberration and noise

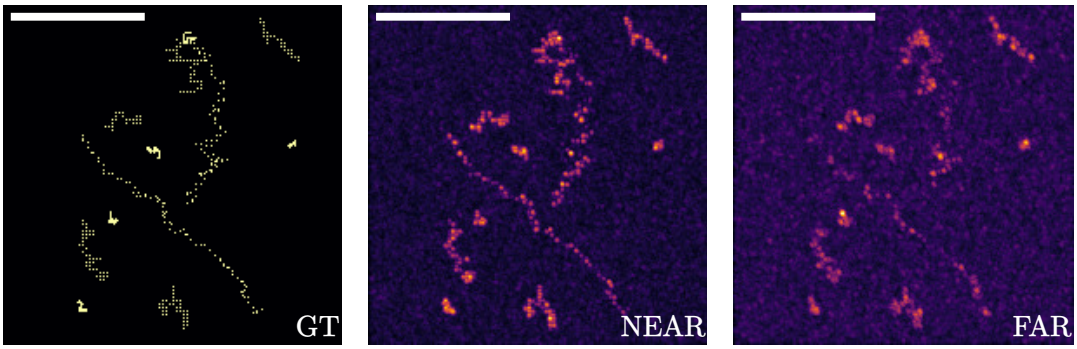
Synthetic data, generated by simulation, was used to demonstrate the necessity for aberration correction and image denoising. The data was made by first creating 3D particle tracks in a volume, as shown in Figure 6.4a. Twelve different tracks were created, split between the three main categories outlined in subsection 6.4.3: 5 tracks of tightly confined diffusion, 5 tracks of restricted diffusion, and 2 tracks of long directed motion. This constituted the ground truth (GT) data.

Next, two widefield PSFs were used to simulate the image formation of a real microscope. They were acquired on the microscope with adaptive optics described in Chapter 4, with a sample similar to the one illustrated in Figure 4.24a. One of them was taken right on the surface of the coverslip, where the sample aberrations are negligible, while the other was taken under an entire *Drosophila* larval brain, and thus suffered from strong distortion, dominated by spherical aberrations. In addition to the blurring with the PSFs, the GT data was also further corrupted by a combination of Gaussian ($\mu = 0, \sigma = 0.03$) and Poissonian noise ($\alpha = 5$, $x \in [0; 1]$), following the model in Eq. 2.16. The results of the blurring and noise are shown in the middle (near PSF) and right (far PSF) panels of Figure 6.4a.

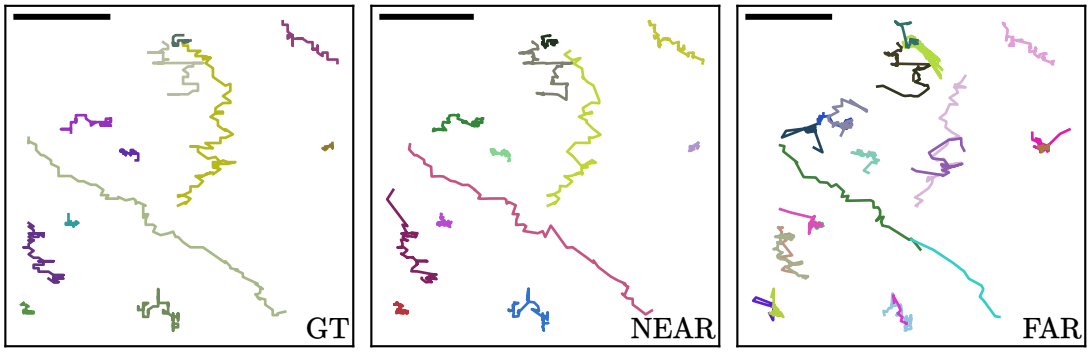
The three variants of the synthetic data were used to perform particle tracking. The results are shown in Figure 6.4b. With the near dataset the particles were tracked very well, with small deviations from the ground truth data. However, the quality of the far dataset was not good enough to reproduce properly the particle tracks. Most notably, the long directed motion tracks were incomplete or broken into multiple fragments, some of the restricted diffusion tracks were malformed and aberrant, and there were several tracks that were caused by the noise and that resembled restricted diffusion.

6.5 Discussion and conclusions

The improved motility assay described in this chapter presents an exciting opportunity for studying mRNA transport and post-transcriptional regulation in the *Drosophila* nervous system. In theory, any gene expressed in neurons or glia can be investigated in great detail. In practice, the two main pre-requisites are that it is possible to insert MS2 stem loops into the gene (ideally 24 of them for good signal) and that the insertion



(a) Synthetic particle data. Images show maximum projections along the Z and time axes. Left: ground truth. Middle and right: simulating aberrations near and far from the cover slip. Scale bar is 5 µm.



(b) Particle tracks. Each track is drawn with an arbitrary colour. Scale bar is 5 µm.

Figure 6.4: Simulation of mRNA track data with aberrations and noise.

does not lead to disruption of the normal processing of the mRNA.

Extra consideration should be given to the microscopy, and particularly the data acquisition, part of the assay. The results clearly show the importance of good data for the tracking of spots and the subsequent analysis. Spatial resolution does not seem to be important, the particles are large enough to be seen well on any type of modern microscope with high NA objective lens. Temporal resolution, however, is crucial and 1 fps was not enough to tell what is happening inside the nucleus. Gentleness is also important, otherwise photobleaching limits the experiment duration and phototoxicity can lead to non-physiological effects. For all these reasons, data is probably best acquired on a widefield microscope optimised for fast live fluorescence imaging. MS2-MCP sample often suffer from poor contrast due to strong background signal, but this can be overcome with deconvolution or other computational methods.

Tracking with the TrackMate software may be good for some applications, but I question its suitability for the task at hand. The spot detection part of it does not seem sophisticated enough, for example not making full use of the available information, such as PSF of the imaging system and other parameters. The particle tracker is not entirely adequate either. Most of the mRNA particles that exhibit directed and deliberate motility are switching between different modes of motion, such as confined diffusion and directed motion. Most of the provided track features become invalid in such cases, heavily biased towards the more dominant modes, and there is no easy way to decompose such complex trajectories for individual consideration.

When it comes to *shot* in particular, the results from this chapter confirm that it is being transported in glial cells over large distances of several micrometres. The data is insufficient to dissect the exact molecular mechanism, but the assay could be repeated under different conditions to answer specific questions. Two instances were shown where motion happens in opposite directions, once for a group of three particles and another time for a single particle. For the group, the lack of reliable cell markers makes it difficult to claim if all three particles were in the same cell. However, the fact that they move along a common line makes it likely that it is the same cell. For the single particle that travels in opposite directions, it is conceivable that it is a false merge of two separate tracks, one ending where the other begins. Even then, it is a concrete proof that two different modes of transport are possible for *shot* mRNA molecules.

In any case, this type of bidirectional transport is very characteristic of the trafficking along microtubules. The speeds observed for the directed motion tracks were up to $1\,\mu m\,s^{-1}$, which is consistent with the speeds previously reported in the literature for trafficking of mRNA along microtubules in neurons [141]. The paper also shows how only 14% of the mRNA particles exhibit directed transport, which is also consistent with the saltatory behaviour of *shot* particles observed in this chapter.

Chapter 7

Conclusions and future work

7.1 Software for microscopy

Chapter 3 outlined the development of an open-source Python-based software for control of bespoke imaging systems. The software was then used to drive the microscope described in Chapter 4. Many improvements were made to the GUI and to the core functionality, with an emphasis on the adaptive optics side of it.

Microscope-Cockpit is a mature platform with a long history, but it is difficult to use. It is more suitable for engineers and even then only those with some fluency in Python, which challenges some of the core values of the project, such as openness and collaboration.

It is a very powerful tool, with a wide scope of operation and full of advanced features. It is not only extensive but very extensible and users are, indeed, encouraged to make modification to the code. However, documentation is absent. What is available out there is not easy to access, things such as code comments, obscure text files, etc. It is very commendable, though, that the core developers organise monthly meetings where people can ask questions and get help. They are also quite active online and have presence on popular imaging discussion forums.

The software is multiplatform but it is also clunky and features a very plain GUI, with a poor choice of GUI toolkit in general. It does the job but it is not intuitive and it lacks consistency and cohesion. More importantly, it does not seem like a very friendly or welcoming environment for new or even regular users. The configuration setup and the way the software is launched definitely needs to change. The config files

and daemons need to be replaced. The overarching workflow of the software is great, and it has been inspired by and it encourages best practices. However, it is not intuitive and not robust enough, sometimes leading to crashes, memory hogging, or something just stops working suddenly.

Yet there are not many suitable alternatives out there. They are all much less complete and mature. So despite all the problems, Microscope-Cockpit still finds use and it is arguably the best open-source software at what it does. It would be great if it finds more enthusiasm in the broader microscopy community, and if it receives more attention to things like automation, scripting, and especially interoperation with other software. It can also benefit from a more strategic and organisational approach, creating and documenting specifications of the software, including unit tests, and looking into the overall architectural design.

7.2 Adaptive optics for fluorescence microscopy

The development of adaptive optics methods, described in Chapter 4, was quite comprehensive and it produced great results. One thing that could have been done better is the treatment of the multi-plane aberration correction. It was done for just two planes, one on top and one on bottom of the volume, but maybe more planes would have given better interpolation.

A lot of scope for improvement is on the software side, especially for the remote focusing and the overall device control during image acquisition. For example, it would be very interesting to define an acquisition mode where corrections can be recalculated between frames, instead of using a single correction throughout an entire acquisition session. The advantage is obvious for time-lapse experiments and the live imaging of dynamic samples where it is conceivable for the sample-induced aberrations to change significantly over short periods of time. In general, a lot of the adaptive optics functionality is clunky and tedious to carry out, and there are many opportunities for streamlining and automation.

In terms of improvements, they would come almost naturally from better and more modern or even entirely different hardware devices. For example, replacing the DM could improve the range and quality of the remote focusing, replacing the CCD cam-

eras with sCMOS ones could improve the acquisition speed and the field of view, etc. Another research direction would be more complex arrangements of adaptive optics, such as multiconjugate adaptive optics or sensor-based closed-loop approaches. AI and machine-learning methods are an active area of research that shows really promising performance [142].

Finally, there is always the consideration of optimising and simplifying the optical system, and putting it in a box, suitable to be used as an add-on with commercial microscopes.

7.3 Denoising of fluorescence images

Chapter 5 demonstrated that deep-learning methods lead the current state-of-the-art in the denoising of fluorescence image data. The emphasis there was mostly on self-supervised neural networks because of the difficulty, and even impossibility, of acquiring good ground truth images under the relevant experimental conditions. Supervised methods, however, offer a lot better performance and their application to microscopy deserves consideration.

The task of denoising is defined almost entirely by the imaging system, and little by the sample. It is important how many photons reach the detector, but not really where they came from. Of course, knowledge of the sample could help with the constraining of the solution space via regularisation, but the most popular prior models, such as smoothness and NSS, are already generic enough to apply to all natural images, let alone microscopy ones. Therefore, an ideal denoiser should be a function of the imaging system, and by extension the noise model, and it should be agnostic to the imaged sample.

This line of reasoning suggests that a neural network could be trained on any kind of data, as long as the noise statistics are similar to those of the data on which the network will later make predictions. A small caveat to this idea is that some neural networks, such as CNNs, include contextual information into their prediction, but this can be counteracted by data diversification and augmentation. Still, the results in chapter 5 demonstrated that even a self-supervised network such as Noise2Noise can perform well on different data than it was trained on, and under different imaging conditions.

Having a flexibility in the choice of imaging sample, for the collection of training data, could potentially address the issue of obtaining good ground truth data in biomedical applications. Theoretically, there is no need to capture dynamic time-series or other challenging datasets. All training could be done on simple static samples. The only concern then becomes the intensity of the emission light, so that it matches the intensity range of the noisy data. The reason is that the noise statistics are determined by the intensity levels and the overall imaging conditions, and it is not known if a neural network can learn to remove well more than one type of noise.

An interesting idea and a potential new avenue of research would be the development of a calibration microscope slide specifically designed for the acquisition of standardised training data. It needs to be equipped with a light source which can provide constant steady emission of known precise intensity levels. A device like this would be impervious to photobleaching and it would allow the capturing of multiple noise realisations of the same underlying signal.

Regardless of the existence of a such specialised device, it is still important to create and share high quality training data. Unlike other image restoration tasks such as classification, there are not many public datasets for denoising, let alone ones specifically for fluorescence microscopy. The FMD dataset [143] was recently developed to fill this gap and although it is a step in the right direction it is lacking in many ways, mainly in the choice of noise levels and statistics, the limited selection of detectors, and the absence of multidimensional data. Perhaps a good way to obtain training data is via generative methods, such as in [144], where a clean image and a noise model can be used to create many clean-noisy training pairs.

The point about multidimensionality is particularly important because nowadays fluorescence image data usually comes in the form of a multidimensional set of images, whether that would be a series of frames over time, sections over depth, or something else. Yet vast majority of denoising neural networks are designed to process single 2D images, and these networks consist of 2D operations as part of their design. On one hand the images in a microscopy dataset should have independent noise realisations because noise is a consequence of image formation and microscopes capture data one image at a time. But on the other hand, some neural networks, such as blind spot ones, do use contextual information to some degree. The experiments in chapter 5 suggest

that 3D architectures can offer superior performance in some situations when dealing with data of many dimensions.

The transition from 2D to 3D network blocks is fairly straightforward and it should not impact the performance on 2D data. From a computational perspective, however, 3D networks are slower and more demanding. Blocks of fourth order or higher are poorly supported in popular neural network software frameworks, and it is conceivable that they would be even more computationally intensive.

Besides an increase of dimensionality, more concrete benefits can come from overall network optimisation. Many networks are initially developed around almost arbitrary design decisions and then left unchanged after the proof of concept. Exceptions to this trend, which have revised the original architecture and tuned its parameters, have demonstrated significant improvements in processing times and memory footprint [111].

Another important parameter, specifically for CNNs, is the receptive field of the network. This plays a key role in blind spot networks, where the engineering of the receptive field has been one of the main driving forces in recent research. Considerable efforts have been spent on modern such network to adjust the shape and the size of the receptive field [145, 146]. Although their design is motivated by the removal of spatially correlated noise in colour photographs, it would be interesting to see if they can bring any improvement to fluorescence microscopy.

Finally, another interesting recent architectural development has been the adoption of transformers. They are an alternative to CNNs and they have shown superior performance in domains such as image recognition [147]. They bring global contextual information to the attention of the network, and this has been shown to offer advantages in image denoising [146, 148]. Their application to multidimensional data with rich spatio-temporal content may be especially beneficial, because of the capturing of global features.

7.4 Live mRNA motility assay

The improved motility assay presented in Chapter 6 provided an easy-to-use yet effective method of studying mRNA dynamics in the cells of the nervous system in *Drosophila*. It combined work and ideas from all the previous technical chapters, and

served as a demonstration of their interoperability.

Besides the pre-requisite to have MS2 stem loops inserted in a gene of interest, one the first important step of the assay is the data acquisition. The results from Chapter 6 highlighted several problems in this regard. The time resolution is important, and the rate of 1 fps was not enough to capture the trajectories of some of the particles. At the same time, axonal projections have diameters of at least 5 µm, so it is difficult to acquire an entire volume at the required rate. The solution is likely to focus on the surface of the axon bundle, if the observed gene is expressed there, or maybe to look at the tips of glia at the NMJ. Imaging fillet preparations is not easy either, so yet another option is to look at the peripheral nerve of the wing which may be more optically tractable [149].

Another challenge with the image acquisition is the poor contrast. Using a confocal microscope for its ability to reject background light could help, but it will likely come at the expense of reduced temporal resolution and increased light dose. A widefield system with subsequent deconvolution would probably perform much better. However, neither can help if the signal of mRNA particles is close to or lower than the background signal. This can sometimes be caused by unbound MCP expression which can be reduced by using heterozygous flies [150]. Alternatively, a more drastic measure is to omit the Gal4 driver entirely and to rely solely on weak passive ubiquitous expression of MCP. Other things that could help are increasing the number of stem loops and using a nuclear localisation signal to keep most of the MCP within the nucleus.

The third problem with the image acquisition was photobleaching, which also hints at concerning levels of photo-toxicity. The only way to address this problem is to reduce the light dose, but this will also reduce the noise and reduce the contrast. This further accentuates the need for good SNR and efficient collection of light. The next step of the improved motility assay was the detection of spots and the generation of tracks. The popular software TrackMate was used, but more sophisticated methods could be employed, reminiscent of the quantification procedures for smFISH [151].

Once the tracks were generated, three key modes of transportation were identified among 149 tracks: confined, restricted, and directed motion. This was done by first using PCA to project the features of the data to two components and then by k-means clustering, all unsupervised. It may be worthwhile to capture more data initially,

label some of it by hand, and try supervised learning. Regardless, more suitable ways to analyse such tracks exist, and two examples are the calculation of mean-square displacements and the use of Hidden Markov Models with Bayesian model selection [152].

Regarding the particular gene used to test the improved motility assay – *shot* – the results show that it is likely transported along microtubules, exhibited in some cases long directed motion. Trajectories in opposite directions were observed, but there was no reliable way to confirm it. This highlights the usefulness of nuclear and membrane markers. In order to further narrow down the type of motor, first the directionality of microtubules in PNS glia needs to be elucidated. For comparison, it is known that other glial cells, such as astrocytes have dense microtubule networks with the minus end pointing towards the nucleus, similar as in neurons [153].

The assay can be further enhanced with the use of more specialised glial drivers, instead of the pan-glial Repo, which can help with determining the tissue-specific expression pattern of the gene being tested with the assay. It is also possible to look for localised translation events, which are particularly interesting for mRNA molecules that find themselves at the periphery of the cell [154].

Appendices

Appendix A

Details about the Alpao DM69-15

A.1 Firmware bug

While integrating the Alpao DM69-15 into the imaging system, it was noticed that stepping through the queue of actuator patterns worked well for shorter sequences but failed for longer ones. To investigate this behaviour further, an experiment was performed where a set of three distinct actuator patterns were repeated 150 times. The three patterns poked actuators in the following shapes: four dots in each corner, vertical line, and horizontal line. The device was then triggered enough times to go through the entire queue and an interferometer was used to detect the corresponding wavefront phase at each step. The results are shown in Figure A.1.

It was observed that every 131st pattern is skipped, as indicated by the three magenta lines which are 131 images apart. As a result the queue terminated prematurely with three less images, shown with a green line. To address this issue every 131st pattern was duplicated.

A.2 Step response

The step response of the Alpao DM69-15 was assessed to determine how quickly the device can switch between patterns. An interferometer was used to monitor the transition in wavefront phase when some of the actuators were set from their neutral normalised voltage of 0.5 to a normalised voltage of 0.8. An oscilloscope was used to keep track of the different events, as shown in Figure A.2. Ten interferograms were captured in total, spaced 500 µs apart. The camera could not go any faster because even with a

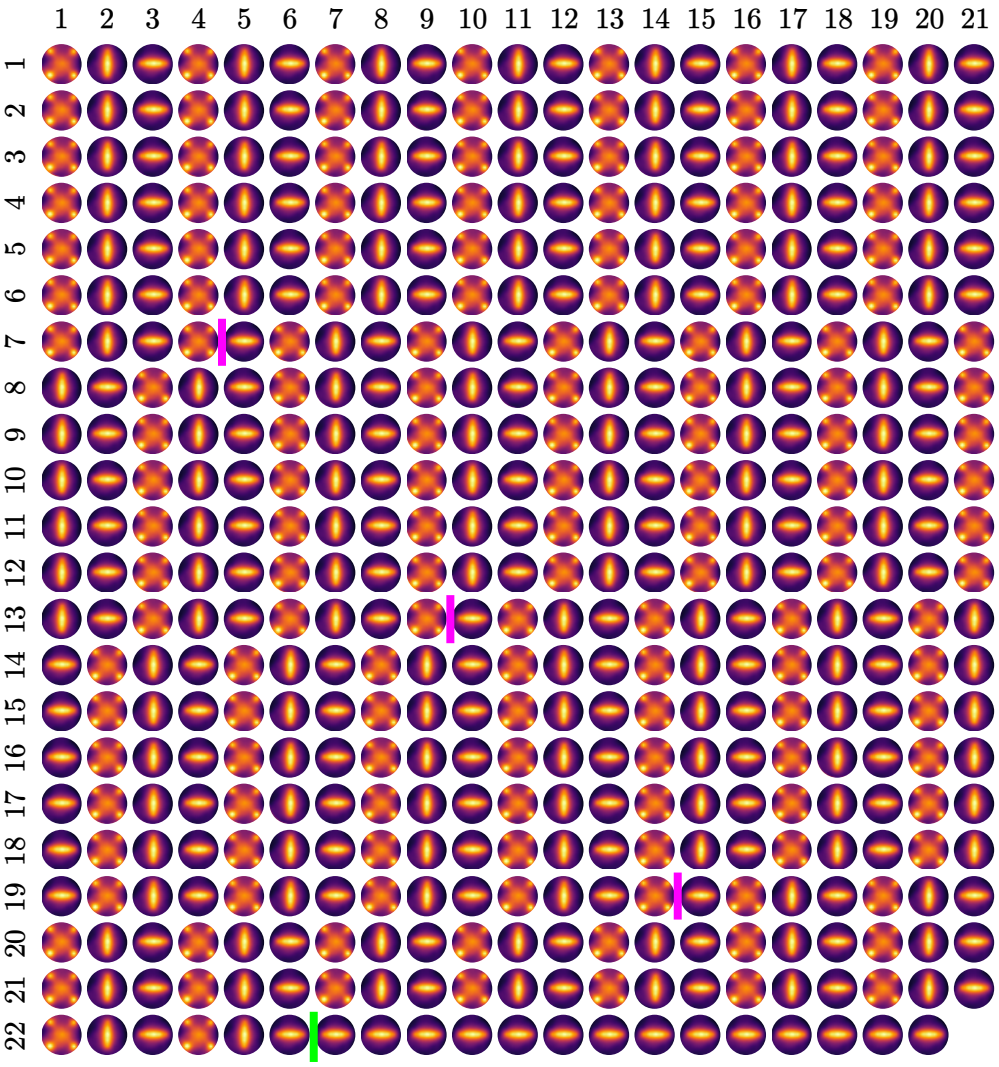


Figure A.1: Alpao DM69-15 triggering behaviour. Three patterns were repeated 150 times and queued, then they were iterated via triggering and the shapes of the corresponding wavefronts were measured. The three magenta lines indicate the positions where a pattern is missing. The green line indicates the end of the queue.

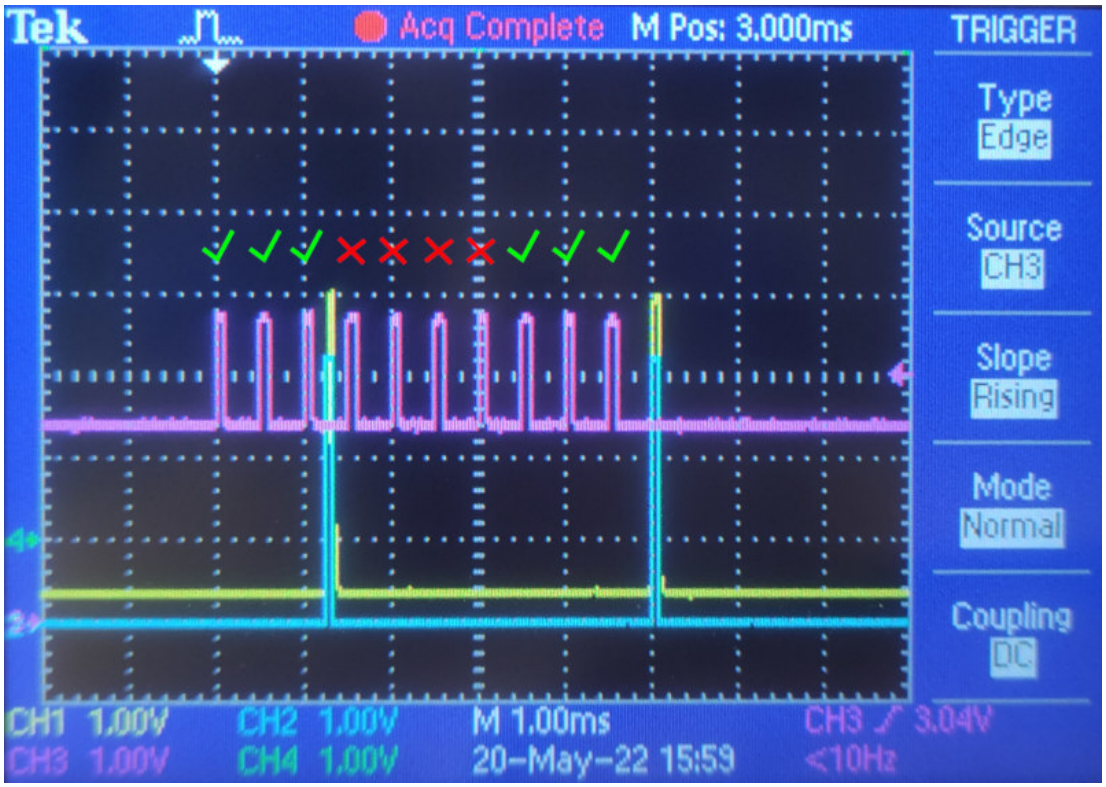


Figure A.2: Timing diagram of one of the experiments to test the speed of the Alpao DM69-15. The magenta waveform shows the triggers for the interferometer camera and effectively the sampling rate of 2000 Hz. The yellow and cyan waveforms show the same thing – triggers for the DM – just on different sides of the connection. The green checkmarks and the red crosses indicate whether the wavefront shape was stable or not, respectively, as judged by the interferometer.

tiny ROI of 64×64 pixels at the lowest exposure time of 103 µs the readout time was 305.5 µs. The step change was introduced at 1250 µs. The stability of the surface of the DM was determined with the interferometer, where unstable surface would fail to produce fringe pattern. This is indicated in Figure A.2 with green checkmarks and red crosses. The device took up to 2000 µs to assume its final steady shape. This kind of experiment was repeated several time and with different step sizes, not just 0.8, and the results were the same.

The performance of the DM was compared to that of the Z piezo stage (PI P-736 driven by PI E-709). The step responses of the stage were recorded in two different ways: digital and analogue. The digital responses used the serial connection (USB) and the manufacturer software (PIMikroMove) to change the position setpoint and then to read out the actual position for a period of time. The maximum sampling frequency of 10 kHz was used. The results of the digital tests are shown in Figure A.3.

Alternatively, the stage provides an analogue output signal which was read with

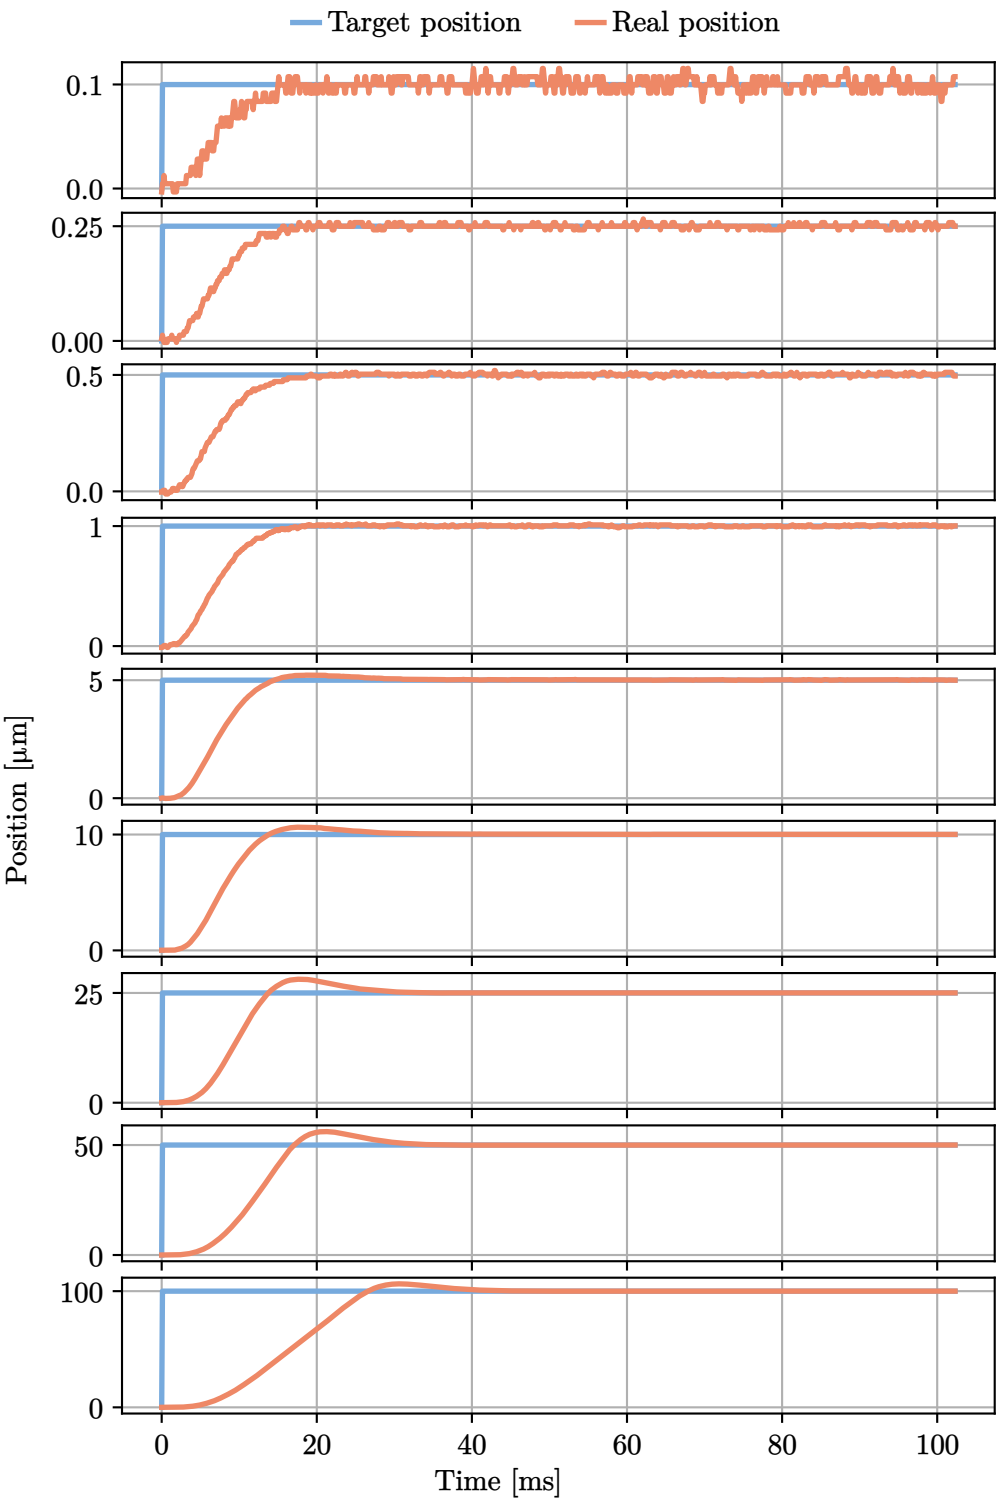
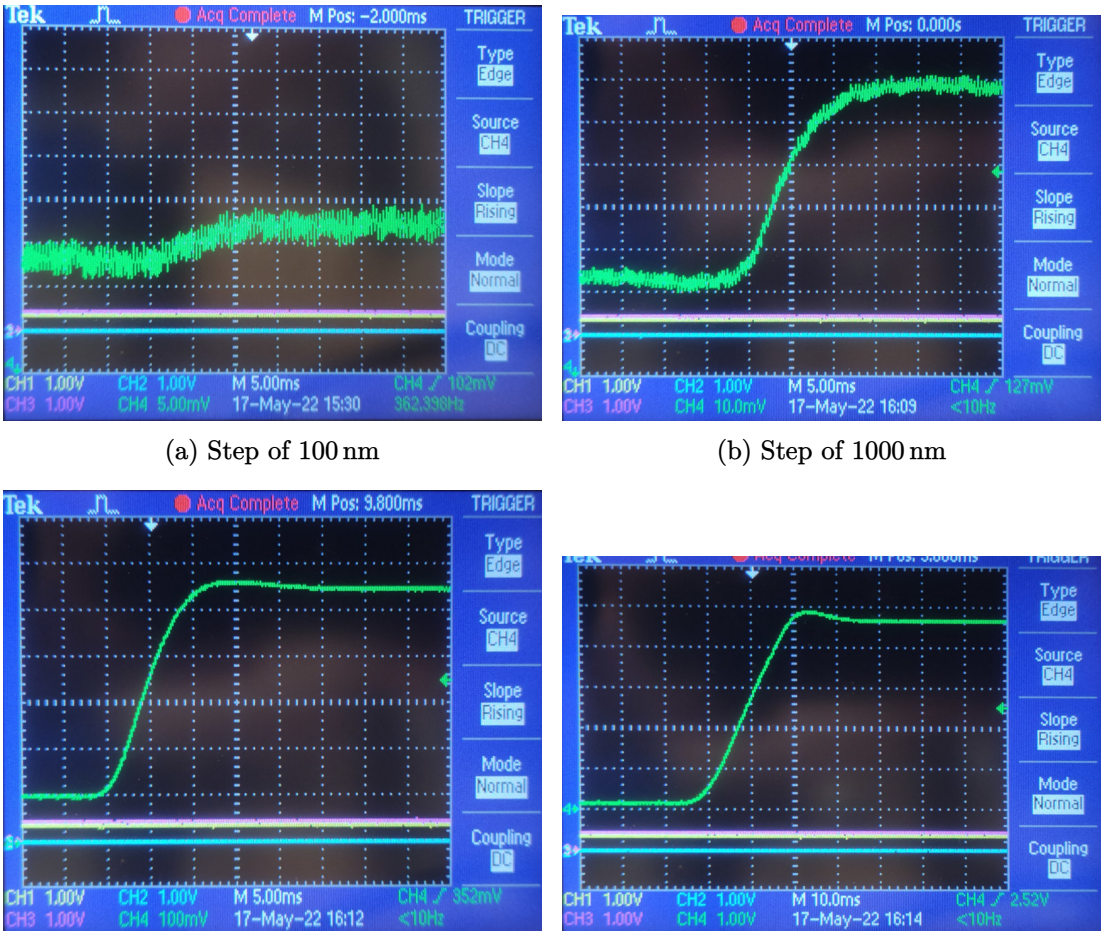


Figure A.3: Digital step responses of the PI P-736 stage driven by PI E-709. The data was acquired with PIMikroMove with a sampling frequency of 10 kHz.



(a) Step of 100 nm

(b) Step of 1000 nm

(c) Step of 10 µm

(d) Step of 100 µm

Figure A.4: Analogue step responses of the PI P-736 stage driven by PI E-709.

an oscilloscope. The results are shown in Figure A.4 and they are consistent with the previously shown digital ones. Overall, the stage was significantly slower than the DM. It may be possible to increase the speed the stage by tuning its controller differently, but this would likely result in other compromises such as oscillations or steady-state errors.

Appendix B

Mathematical proofs

B.1 Minimising the MSE

Given a neural network f_θ and two random variables x and y, training with MSE loss can be expressed in the following way:

$$\begin{aligned} f_\theta^\star &= \underset{f_\theta}{\operatorname{argmin}}\, \mathbb{E}\left[(f_\theta(y) - x)^2\right] \\ &= \underset{f_\theta}{\operatorname{argmin}}\, \mathbb{E}\left[(f_\theta(y) - \mathbb{E}\left[x \mid y\right] + \mathbb{E}\left[x \mid y\right] - x)^2\right] \\ &= \underset{f_\theta}{\operatorname{argmin}} \Big\{ \mathbb{E}\left[(f_\theta(y) - \mathbb{E}\left[x \mid y\right])^2\right] + 2\mathbb{E}\Big[(f_\theta(y) - \mathbb{E}\left[x \mid y\right])(\mathbb{E}\left[x \mid y\right] - x)\Big] \\ &\qquad + \mathbb{E}\left[(\mathbb{E}\left[x \mid y\right] - x)^2\right] \Big\} \end{aligned} \tag{B.1}$$

The cross term can be shown to be zero:

$$\mathbb{E}\Big[(f_\theta(y) - \mathbb{E}\left[x \mid y\right])(\mathbb{E}\left[x \mid y\right] - x)\Big] = \mathbb{E}\Big[\mathbb{E}\Big[(f_\theta(y) - \mathbb{E}\left[x \mid y\right])(\mathbb{E}\left[x \mid y\right] - x) \;\Big|\; y\Big]\Big] \tag{B.2}$$

$$= (f_\theta(y) - \mathbb{E}\left[x \mid y\right])\mathbb{E}\Big[\mathbb{E}\Big[\mathbb{E}\left[x \mid y\right] - x \;\Big|\; y\Big]\Big] \tag{B.3}$$

$$= (f_\theta(y) - \mathbb{E}\left[x \mid y\right])\mathbb{E}\Big[\mathbb{E}\Big[\mathbb{E}\left[x \mid y\right] \;\Big|\; y\Big] - \mathbb{E}\left[x \mid y\right]\Big]$$

$$= (f_\theta(y) - \mathbb{E}\left[x \mid y\right])\mathbb{E}\Big[\mathbb{E}\left[x \mid y\right] - \mathbb{E}\left[x \mid y\right]\Big] \tag{B.4}$$

$$= (f_\theta(y) - \mathbb{E}\left[x \mid y\right]) \cdot 0$$

$$= 0$$

Line B.2 uses the law of total expectation. Lines B.3 and B.4 use the property of conditional expectation $\mathbb{E}[f(x) \mid x] = f(x)$ and the fact that conditional expectation is itself a function of the condition. Equation B.1 now simplifies to:

$$f_\theta^\star = \underset{f_\theta}{\operatorname{argmin}} \left\{ \mathbb{E}\left[(f_\theta(y) - \mathbb{E}\left[x \mid y\right])^2 \right] + \mathbb{E}\left[(\mathbb{E}\left[x \mid y\right] - x)^2 \right] \right\} \tag{B.5}$$

The second term in equation B.5 is not a function of f_θ and therefore can be ignored during the optimisation. The optimum is then found at $f_\theta^\star = \mathbb{E}[x \mid y]$.

B.2 Comparison between noisy and clean training targets

The learning of a neural network f_θ depends on the data it is presented with during training. Using MSE loss, it can be shown that supervised learning with pairs of clean and noisy data (x, y) is comparable to that of pairs of only noisy data (y, z):

$$\begin{aligned}
f_\theta^\star &= \underset{f_\theta}{\operatorname{argmin}}\, \mathbb{E}\left[(f_\theta(y) - z)^2 \right] \\
&= \underset{f_\theta}{\operatorname{argmin}}\, \mathbb{E}\left[(f_\theta(y) - x + x - z)^2 \right] \\
&= \underset{f_\theta}{\operatorname{argmin}}\, \mathbb{E}\left[((f_\theta(y) - x) - (z - x))^2 \right] \\
&= \underset{f_\theta}{\operatorname{argmin}}\, \mathbb{E}\left[(f_\theta(y) - x)^2 - 2(f_\theta(y) - x)(z - x) + (z - x)^2 \right] \\
&= \underset{f_\theta}{\operatorname{argmin}}\, \mathbb{E}\left[(f_\theta(y) - x)^2 - 2f_\theta(y)(z - x) + 2x(z - x) + (z - x)^2 \right] \\
&= \underset{f_\theta}{\operatorname{argmin}} \left\{ \mathbb{E}\left[(f_\theta(y) - x)^2 \right] - 2\mathbb{E}\left[f_\theta(y)(z - x) \right] + 2\mathbb{E}\left[x(z - x) \right] + \mathbb{E}\left[(z - x)^2 \right] \right\} \\
&= \underset{f_\theta}{\operatorname{argmin}} \left\{ \mathbb{E}\left[(f_\theta(y) - x)^2 \right] - 2\mathbb{E}\left[f_\theta(y)(z - x) \right] \right\} + \text{const.}
\end{aligned} \tag{B.6}$$

The last line of equation B.6 uses the fact that any terms that are not function of f_θ are effectively constant and can be ignored during the optimisation.

Appendix C

Complete denoising metrics

The metrics in Figure C.1 and C.2 were calculated as described in section 5.3.7. The metrics in Figure C.3 were calculated in the same way, but in addition the statistics (mean and standard deviation) were combined across scenes and channels by using equation 5.5.

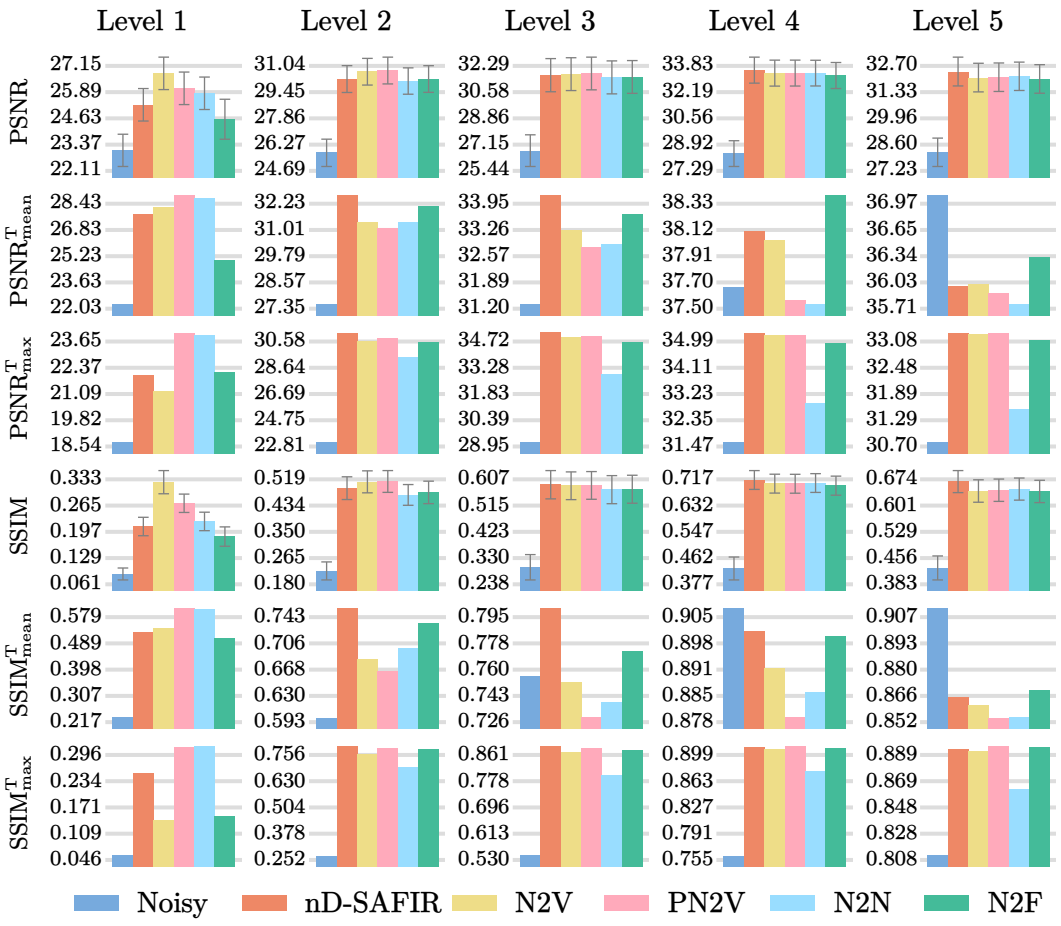


Figure C.1: Complete quantitative results for the 2D+t dataset. Error bars indicate $\pm\sigma$.

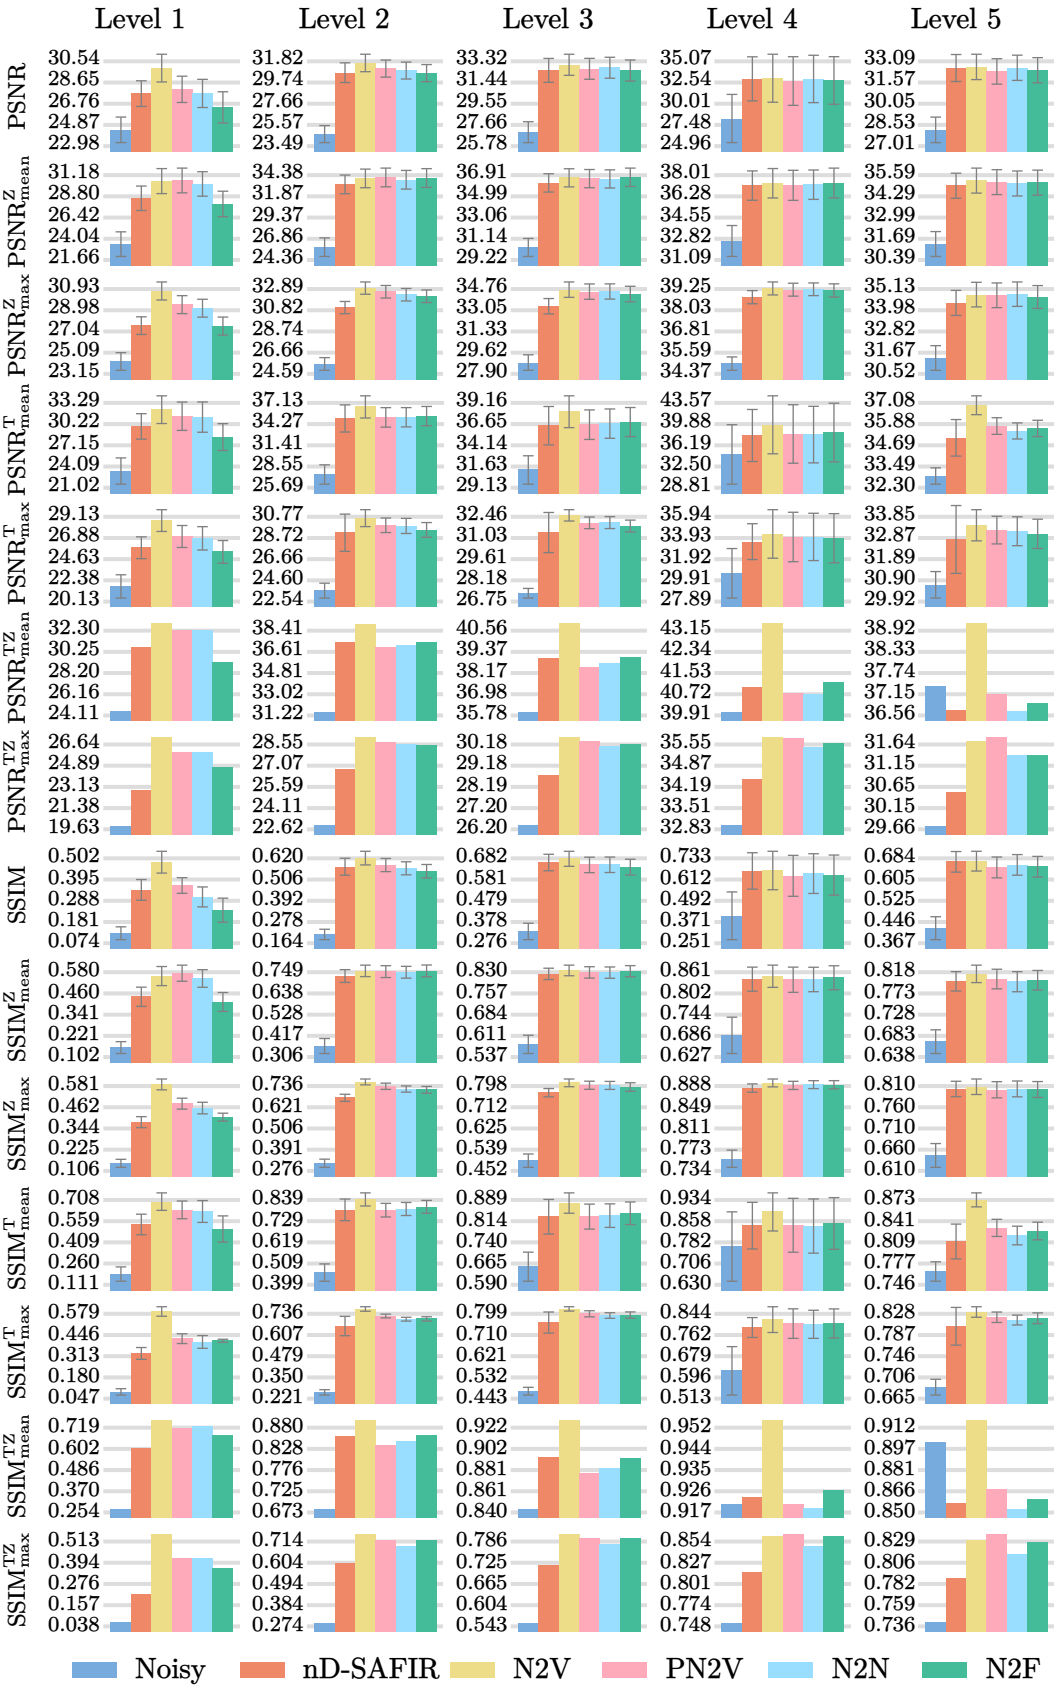


Figure C.2: Complete quantitative results for the 3D+t dataset. Error bars indicate $\pm\sigma$.

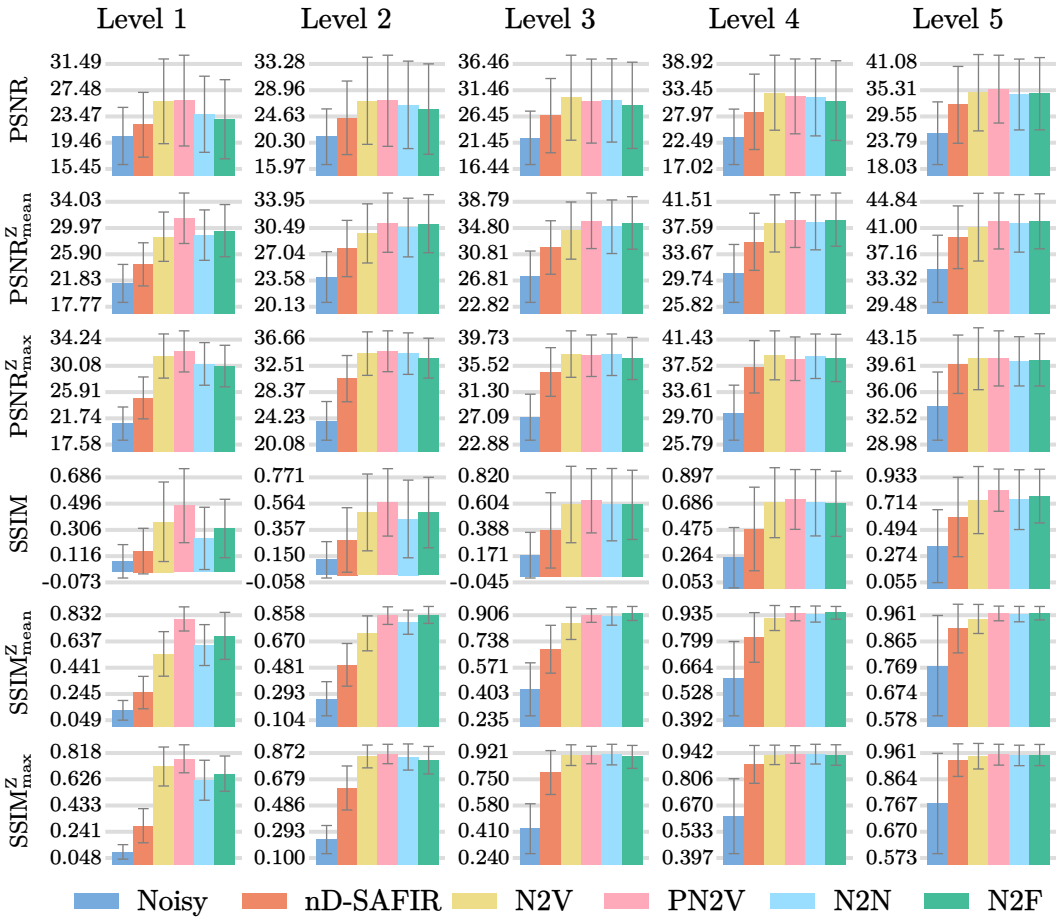


Figure C.3: Complete quantitative results for the 3D dataset. Error bars indicate $\pm\sigma$.

Appendix D

smFISH probe sequences

Table D.1: smFISH probes for the largest exon of the *shot* gene in *Drosophila*.

Probe #	Sequence (5'-3')	Probe #	Sequence (5'-3')
1	tctgacggatttgctcca	25	cccgtagcgattgttgat
2	tccagagtggtcacttcg	26	agctagtcccttcagatt
3	gggcgacttgcaaccaat	27	tgcttaatgcagcgctcg
4	tccttgtcctccaaagtg	28	tcttcgaggatcttgcgc
5	gatgaaccgcagatcggc	29	cggaggatatccttcacc
6	acaaacttctgtgcggcc	30	atgctatcgttctcctgg
7	ccaaaccggagacgcgat	31	acgttccttactgtccag
8	ttaggatcgccggcaatc	32	tggaccatacggggatcg
9	ccttgataagcgaggcgg	33	gcggaattcctcgcgaag
10	cgctggcaagcagatgtt	34	atgtccaagctcctggaa
11	ttcttggcaatgcgtgca	35	agtttgatttcgaccgcc
12	ttcaccacgtttgttggc	36	caaccagatcatcccact
13	taagctggccatacgctg	37	caaactattggcccggtc
14	gcgacagttcgttcatgc	38	tccacatcggctaacagc
15	tcggtgacatcacgcagg	39	cgggatcacccagaatgt
16	tgttgcgccactgagatt	40	tgtcggtcaggttggaga

Continued on next page

Table D.1: smFISH probes for the largest exon of the *shot* gene in *Drosophila*. (Continued)

Probe #	Sequence (5'-3')	Probe #	Sequence (5'-3')
17	gtaggccagttgttgttc	41	tgttggcttgtgtgcaga
18	taactcgcgcttcagtgc	42	tggcgatcggttagatcc
19	atccttgcatatgggctt	43	cgccgggtgtaagatcag
20	agttggtgcatagtcgcg	44	gtctgcagatcacgcaga
21	tctggacggattagcgca	45	tcgcagctgctgaggaat
22	ggactggtccgcttaatc	46	atgtcttgaagttccgcc
23	ccactggaaccagtcgat	47	gagtctctcccaacgatc
24	ggaaaccatcctgggagg	48	tggctgcagattgcagtt

Table D.2: smFISH probes for the MS2v5 stem loops.

Probe #	Sequence (5'-3')	Probe #	Sequence (5'-3')
1	tcgtgctttcttggcaataa	7	tattccagggttcatcagat
2	aatactggagcgacgcgtga	8	aaacgacgacgatcacgcgt
3	accgtaggatctgatgaacc	9	agaaagcacgagcatcagcc
4	cgtttgaagattcgacctgg	10	gatcctacggtacttattgc
5	gtgatcgtcgtcgtttgaag	11	tccagtattccagggttcat
6	tgaaccctggaatactggag		

References

1. Maday, S., Twelvetrees, A. E., Moughamian, A. J. & Holzbaur, E. L. Axonal transport: cargo-specific mechanisms of motility and regulation. *Neuron* **84,** 292–309 (2014).

2. Holt, C. E., Martin, K. C. & Schuman, E. M. Local translation in neurons: visualization and function. *Nature structural & molecular biology* **26,** 557–566 (2019).

3. Blanco-Urrejola, M. *et al.* RNA localization and local translation in glia in neurological and neurodegenerative diseases: lessons from neurons. *Cells* **10,** 632 (2021).

4. Fernandopulle, M. S., Lippincott-Schwartz, J. & Ward, M. E. RNA transport and local translation in neurodevelopmental and neurodegenerative disease. *Nature neuroscience* **24,** 622–632 (2021).

5. Liao, Y.-C. *et al.* RNA granules hitchhike on lysosomes for long-distance transport, using annexin A11 as a molecular tether. *Cell* **179,** 147–164 (2019).

6. Buxbaum, A. R., Haimovich, G. & Singer, R. H. In the right place at the right time: visualizing and understanding mRNA localization. *Nature reviews Molecular cell biology* **16,** 95–109 (2015).

7. Pilling, A. D., Horiuchi, D., Lively, C. M. & Saxton, W. M. Kinesin-1 and Dynein are the primary motors for fast transport of mitochondria in Drosophila motor axons. *Molecular biology of the cell* **17,** 2057–2068 (2006).

8. Fischer, J. A., Giniger, E., Maniatis, T. & Ptashne, M. GAL4 activates transcription in Drosophila. *Nature* **332,** 853–856 (1988).

9. Bertrand, E. *et al.* Localization of ASH1 mRNA particles in living yeast. *Molecular cell* **2,** 437–445 (1998).

10. Shimomura, O., Johnson, F. H. & Saiga, Y. Extraction, purification and properties of aequorin, a bioluminescent protein from the luminous hydromedusan, Aequorea. *Journal of cellular and comparative physiology* **59,** 223–239 (1962).

11. Gauss, C. F. *Dioptrische Untersuchungen* (Druck und Verlag der Dieterichschen Buchhandlung, 1841).

12. Seidel, L. Zur Dioptrik. Über die Entwicklung der Glieder 3ter Ordnung welche den Weg eines ausserhalb der Ebene der Axe gelegene Lichtstrahles durch ein System brechender Medien bestimmen, vo Herrn Dr. L. Seidel. *Astronomische Nachrichten, volume 43, Issue 19, p. 289* **43,** 289 (1856).

13. Schwarzschild, K. *Untersuchungen zur geometrischen Optik: Einleitung in die Fehlertheorie optischer Instrumente auf Grund des Eikonalbegriffs. I* (Weidmann, 1905).

14. Herzberger, M. Modern geometrical optics. *Huntington* (1980).

15. Buchdahl, H. *Optical Aberration Coefficients* ISBN: 9780486620107 (Dover Publications, 1968).

16. Hopkins, H. *Wave Theory of Aberrations* ISBN: 9780598888518 (Clarendon Press, 1950).

17. Von F, Z. Beugungstheorie des schneidenver-fahrens und seiner verbesserten form, der phasenkontrastmethode. *physica* **1,** 689–704 (1934).

18. Bhatia, A. & Wolf, E. *On the circle polynomials of Zernike and related orthogonal sets* in *Mathematical Proceedings of the Cambridge Philosophical Society* **50** (1954), 40–48.

19. Niu, K. & Tian, C. Zernike polynomials and their applications. *Journal of Optics* (2022).

20. Noll, R. J. Zernike polynomials and atmospheric turbulence. *JOsA* **66,** 207–211 (1976).

21. Liu, F., Robinson, B. M., Reardon, P. J. & Geary, J. M. Analyzing optics test data on rectangular apertures using 2-D Chebyshev polynomials. *Optical Engineering* **50,** 043609–043609 (2011).

22. Hu, Q., Xiao, Y., Cui, J., Turcotte, R. & Booth, M. J. *The Lattice Geometry of Walsh-Function-Based Adaptive Optics* in *Photonics* **9** (2022), 547.

23. Shack, R. V. *Production and use of a lenticular Hartmann screen* in *Spring Meeting of Optical Society of America, 1971* **656** (1971).

24. Malacara, D. Twyman–green interferometer. *Optical shop testing,* 46–96 (2007).

25. Malacara, Z. & Servin, M. *Interferogram analysis for optical testing* (CRC press, 2018).

26. Takeda, M., Ina, H. & Kobayashi, S. Fourier-transform method of fringe-pattern analysis for computer-based topography and interferometry. *JosA* **72,** 156–160 (1982).

27. Herráez, M. A., Burton, D. R., Lalor, M. J. & Gdeisat, M. A. Fast two-dimensional phase-unwrapping algorithm based on sorting by reliability following a noncontinuous path. *Applied optics* **41,** 7437–7444 (2002).

28. Loktev, M., Soloviev, O., Savenko, S. & Vdovin, G. Speckle imaging through turbulent atmosphere based on adaptable pupil segmentation. *Optics letters* **36,** 2656–2658 (2011).

29. Milkie, D. E., Betzig, E. & Ji, N. Pupil-segmentation-based adaptive optical microscopy with full-pupil illumination. *Optics letters* **36,** 4206–4208 (2011).

30. Hu, Q. *et al.* A universal framework for microscope sensorless adaptive optics: Generalized aberration representations. *APL Photonics* **5** (2020).

31. Arimoto, R. & Murray, J. A common aberration with water-immersion objective lenses. *Journal of microscopy* **216,** 49–51 (2004).

32. Mohanan, S. & Corbett, A. D. Understanding the limits of remote focusing. *Optics Express* **31,** 16281–16294 (2023).

33. Pozzi, P. *et al.* Anisoplanatic adaptive optics in parallelized laser scanning microscopy. *Optics Express* **28,** 14222–14236 (2020).

34. Fienup, J. R. Reconstruction of an object from the modulus of its Fourier transform. *Optics letters* **3,** 27–29 (1978).

35. Gonsalves, R. A. Phase retrieval and diversity in adaptive optics. *Optical Engineering* **21,** 829–832 (1982).

36. Aviles-Espinosa, R. *et al.* Measurement and correction of in vivo sample aberrations employing a nonlinear guide-star in two-photon excited fluorescence microscopy. *Biomedical optics express* **2,** 3135–3149 (2011).

37. Tao, X. *et al.* Adaptive optics confocal microscopy using direct wavefront sensing. *Optics letters* **36,** 1062–1064 (2011).

38. Žurauskas, M. *et al.* IsoSense: frequency enhanced sensorless adaptive optics through structured illumination. *Optica* **6,** 370–379 (2019).

39. Madec, P.-Y. *Overview of deformable mirror technologies for adaptive optics and astronomy* in *Adaptive Optics Systems III* **8447** (2012), 22–39.

40. Love, G. D. Wave-front correction and production of Zernike modes with a liquid-crystal spatial light modulator. *Applied optics* **36,** 1517–1524 (1997).

41. Banerjee, K., Rajaeipour, P., Ataman, Ç. & Zappe, H. Optofluidic adaptive optics. *Applied optics* **57,** 6338–6344 (2018).

42. Hornbeck, L. J. 128× 128 deformable mirror device. *IEEE Transactions on Electron Devices* **30,** 539–545 (1983).

43. Botcherby, E. J., Juškaitis, R., Booth, M. J. & Wilson, T. An optical technique for remote focusing in microscopy. *Optics Communications* **281,** 880–887 (2008).

44. Mohanan, S. & Corbett, A. D. Sensitivity of remote focusing microscopes to magnification mismatch. *Journal of Microscopy* **288,** 95–105 (2022).

45. Booth, M. J., Neil, M. A. & Wilson, T. Aberration correction for confocal imaging in refractive-index-mismatched media. *Journal of microscopy* **192,** 90–98 (1998).

46. Cui, J., Antonello, J., Kirkpatrick, A. R., Salter, P. S. & Booth, M. J. Generalised adaptive optics method for high-NA aberration-free refocusing in refractive-index-mismatched media. *Optics Express* **30,** 11809–11824 (2022).

47. Foi, A., Trimeche, M., Katkovnik, V. & Egiazarian, K. Practical Poissonian-Gaussian noise modeling and fitting for single-image raw-data. *IEEE transactions on image processing* **17,** 1737–1754 (2008).

48. Hadamard, J. Sur les problèmes aux dérivées partielles et leur signification physique. *Princeton university bulletin,* 49–52 (1902).

49. Tikhonov, A. N. & Arsenin, V. Solutions of ill-posed problems. *(No Title)* (1977).

50. Hastings, W. K. Monte Carlo sampling methods using Markov chains and their applications (1970).

51. Jordan, M. I., Ghahramani, Z., Jaakkola, T. S. & Saul, L. K. An introduction to variational methods for graphical models. *Machine learning* **37,** 183–233 (1999).

52. Lucas, A., Iliadis, M., Molina, R. & Katsaggelos, A. K. Using deep neural networks for inverse problems in imaging: beyond analytical methods. *IEEE Signal Processing Magazine* **35,** 20–36 (2018).

53. Klann, E. & Ramlau, R. Regularization by fractional filter methods and data smoothing. *Inverse Problems* **24,** 025018 (2008).

54. Rudin, L. I., Osher, S. & Fatemi, E. Nonlinear total variation based noise removal algorithms. *Physica D: nonlinear phenomena* **60,** 259–268 (1992).

55. Strong, D. & Chan, T. Edge-preserving and scale-dependent properties of total variation regularization. *Inverse problems* **19,** S165 (2003).

56. Buades, A., Coll, B. & Morel, J.-M. *A non-local algorithm for image denoising* in *2005 IEEE computer society conference on computer vision and pattern recognition (CVPR'05)* **2** (2005), 60–65.

57. Dabov, K., Foi, A., Katkovnik, V. & Egiazarian, K. Image denoising by sparse 3-D transform-domain collaborative filtering. *IEEE Transactions on image processing* **16,** 2080–2095 (2007).

58. Boulanger, J. *et al.* Patch-based nonlocal functional for denoising fluorescence microscopy image sequences. *IEEE transactions on medical imaging* **29,** 442–454 (2009).

59. Russakovsky, O. *et al.* Imagenet large scale visual recognition challenge. *International journal of computer vision* **115,** 211–252 (2015).

60. He, K., Zhang, X., Ren, S. & Sun, J. *Delving deep into rectifiers: Surpassing human-level performance on imagenet classification* in *Proceedings of the IEEE international conference on computer vision* (2015), 1026–1034.

61. Fukushima, K. Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biological cybernetics* **36,** 193–202 (1980).

62. Krizhevsky, A., Sutskever, I. & Hinton, G. E. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems* **25** (2012).

63. Springenberg, J. T., Dosovitskiy, A., Brox, T. & Riedmiller, M. Striving for simplicity: The all convolutional net. *arXiv preprint arXiv:1412.6806* (2014).

64. Ronneberger, O., Fischer, P. & Brox, T. *U-net: Convolutional networks for biomedical image segmentation* in *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18* (2015), 234–241.

65. Prechelt, L. in *Neural Networks: Tricks of the trade* 55–69 (Springer, 2002).

66. Von Chamier, L. *et al.* Democratising deep learning for microscopy with ZeroCostDL4Mic. *Nature communications* **12,** 2276 (2021).

67. Chen, H., Sedat, J. W. & Agard, D. A. Manipulation, display, and analysis of three-dimensional biological images. *Handbook of Biological Confocal Microscopy,* 141–150 (1990).

68. Chen, H., Clyborne, W. K., Sedat, J. W. & Agard, D. A. *PRIISM: an integrated system for display and analysis of 3D microscopeimages* in *Biomedical Image Processing and Three-Dimensional Microscopy* **1660** (1992), 784–790.

69. Chhetri, R., Preibisch, S. & Stuurman, N. Software for microscopy workshop white paper. *arXiv preprint arXiv:2005.00082* (2020).

70. Pinkard, H. *et al.* Pycro-Manager: open-source software for customized and reproducible microscope control. *Nature methods* **18,** 226–228 (2021).

71. Dobbie, I. M. *et al.* OMX: a new platform for multimodal, multichannel wide-field imaging. *Cold Spring Harbor protocols* **2011,** pdb–top121 (2011).

72. Susano Pinto, D. M. *et al.* Python-Microscope–a new open-source Python library for the control of microscopes. *Journal of Cell Science* **134,** jcs258955 (2021).

73. Phillips, M. A. *et al.* Microscope-Cockpit: Python-based bespoke microscopy for bio-medical science. *Wellcome Open Research* **6** (2021).

74. Hall, N., Titlow, J., Booth, M. J. & Dobbie, I. M. Microscope-AOtools: a generalised adaptive optics implementation. *Optics Express* **28,** 28987–29003 (2020).

75. Moreno, X. C., Al-Kadhimi, S., Alvelid, J., Bodén, A. & Testa, I. ImSwitch: generalizing microscope control in Python. *Journal of Open Source Software* **6,** 3394 (2021).

76. Hardy, J. W. *Adaptive optics for astronomical telescopes* (Oxford Optical and Imaging Sci, 1998).

77. Booth, M. J. Adaptive optics in microscopy. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **365,** 2829–2843 (2007).

78. Ji, N. Adaptive optical fluorescence microscopy. *Nature methods* **14,** 374–380 (2017).

79. Débarre, D., Booth, M. J. & Wilson, T. Image based adaptive optics through optimisation of low spatial frequencies. *Optics Express* **15,** 8176–8190 (2007).

80. Booth, M. J., Neil, M. A., Juškaitis, R. & Wilson, T. Adaptive aberration correction in a confocal microscope. *Proceedings of the National Academy of Sciences* **99,** 5788–5792 (2002).

81. Sherman, L., Ye, J. Y., Albert, O. & Norris, T. B. Adaptive correction of depth-induced aberrations in multiphoton scanning microscopy using a deformable mirror. *Journal of microscopy* **206,** 65–71 (2002).

82. Gould, T. J., Burke, D., Bewersdorf, J. & Booth, M. J. Adaptive optics enables 3D STED microscopy in aberrating specimens. *Optics express* **20,** 20998–21009 (2012).

83. Kner, P., Sedat, J. W., Agard, D. A. & Kam, Z. High-resolution wide-field microscopy with adaptive optics for spherical aberration correction and motionless focusing. *Journal of microscopy* **237,** 136–147 (2010).

84. Jabbour, J. M. *et al.* Optical axial scanning in confocal microscopy using an electrically tunable lens. *Biomedical optics express* **5,** 645–652 (2014).

85. Fahrbach, F. O., Voigt, F. F., Schmid, B., Helmchen, F. & Huisken, J. Rapid 3D light-sheet microscopy with a tunable lens. *Optics express* **21,** 21010–21026 (2013).

86. Dal Maschio, M., De Stasi, A. M., Benfenati, F. & Fellin, T. Three-dimensional in vivo scanning microscopy with inertia-free focus control. *Optics letters* **36,** 3503–3505 (2011).

87. Juette, M. F., Rivera-Molina, F. E., Toomre, D. K. & Bewersdorf, J. Adaptive optics enables three-dimensional single particle tracking at the sub-millisecond scale. *Applied Physics Letters* **102** (2013).

88. Žurauskas, M., Barnstedt, O., Frade-Rodriguez, M., Waddell, S. & Booth, M. J. Rapid adaptive remote focusing microscope for sensing of volumetric neural activity. *Biomedical optics express* **8,** 4369–4379 (2017).

89. Yang, Y., Chen, W., Fan, J. L. & Ji, N. Adaptive optics enables aberration-free single-objective remote focusing for two-photon fluorescence microscopy. *Biomedical Optics Express* **12,** 354–366 (2021).

90. Marx, V. Microscopy: hello, adaptive optics. *Nature Methods* **14,** 1133–1136 (2017).

91. Marcos, S. *et al.* Vision science and adaptive optics, the state of the field. *Vision research* **132,** 3–33 (2017).

92. *Accessible adaptive optics and super-resolution microscopy to enable improved imaging* PhD thesis (University of Oxford, 2020).

93. Lin, R., Kipreos, E. T., Zhu, J., Khang, C. H. & Kner, P. Subcellular three-dimensional imaging deep through multicellular thick samples by structured illumination microscopy and adaptive optics. *Nature communications* **12,** 3148 (2021).

94. Turcotte, R., Liang, Y. & Ji, N. Adaptive optical versus spherical aberration corrections for in vivo brain imaging. *Biomedical optics express* **8,** 3891–3902 (2017).

95. Bitenc, U., Bharmal, N. A., Morris, T. J. & Myers, R. M. Assessing the stability of an ALPAO deformable mirror for feed-forward operation. *Optics express* **22,** 12438–12451 (2014).

96. Bitenc, U. Software compensation method for achieving high stability of Alpao deformable mirrors. *Optics Express* **25,** 4368–4381 (2017).

97. Hampson, K. & Booth, M. *Calibration and closed-loop control of deformable mirrors using direct sensing* 2020.

98. Antonello, J., Wang, J., He, C., Phillips, M. & Booth, M. Interferometric calibration of a deformable mirror. **10** (2020).

99. Hailstone, M. *et al.* CytoCensus, mapping cell identity and division in tissues and organs using machine learning. *Elife* **9,** e51085 (2020).

100. Brent, J. R., Werner, K. M. & McCabe, B. D. Drosophila larval NMJ dissection. *JoVE (Journal of Visualized Experiments),* e1107 (2009).

101. Parton, R. M., Vallés, A. M., Dobbie, I. M. & Davis, I. Live cell imaging in Drosophila melanogaster. *Cold Spring Harbor Protocols* **2010,** pdb–top75 (2010).

102. Ebrahimi, V. *et al.* Deep learning enables fast, gentle STED microscopy. *bioRxiv,* 2023–01 (2023).

103. Oh, K.-S. & Jung, K. GPU implementation of neural networks. *Pattern Recognition* **37,** 1311–1314 (2004).

104. Chellapilla, K., Puri, S. & Simard, P. *High performance convolutional neural networks for document processing* in *Tenth international workshop on frontiers in handwriting recognition* (2006).

105. Jain, V. & Seung, S. Natural image denoising with convolutional networks. *Advances in neural information processing systems* **21** (2008).

106. Zhang, K., Zuo, W., Chen, Y., Meng, D. & Zhang, L. Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE transactions on image processing* **26,** 3142–3155 (2017).

107. Burger, H. C., Schuler, C. J. & Harmeling, S. *Image denoising: Can plain neural networks compete with BM3D?* in *2012 IEEE conference on computer vision and pattern recognition* (2012), 2392–2399.

108. Xie, J., Xu, L. & Chen, E. Image denoising and inpainting with deep neural networks. *Advances in neural information processing systems* **25** (2012).

109. Mao, X., Shen, C. & Yang, Y.-B. Image restoration using very deep convolutional encoder-decoder networks with symmetric skip connections. *Advances in neural information processing systems* **29** (2016).

110. Lehtinen, J. *et al.* Noise2Noise: Learning image restoration without clean data. *arXiv preprint arXiv:1803.04189* (2018).

111. Li, X. *et al.* Real-time denoising enables high-sensitivity fluorescence time-lapse imaging beyond the shot-noise limit. *Nature Biotechnology* **41,** 282–292 (2023).

112. Lequyer, J., Philip, R., Sharma, A., Hsu, W.-H. & Pelletier, L. A fast blind zero-shot denoiser. *Nature Machine Intelligence* **4,** 953–963 (2022).

113. Huang, T., Li, S., Jia, X., Lu, H. & Liu, J. *Neighbor2neighbor: Self-supervised denoising from single noisy images* in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (2021), 14781–14790.

114. Moran, N., Schmidt, D., Zhong, Y. & Coady, P. *Noisier2noise: Learning to denoise from unpaired noisy data* in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2020), 12064–12072.

115. Pang, T., Zheng, H., Quan, Y. & Ji, H. *Recorrupted-to-recorrupted: Unsupervised deep learning for image denoising* in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (2021), 2043–2052.

116. Zhang, Y. *et al. Idr: Self-supervised image denoising via iterative data refinement* in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2022), 2098–2107.

117. Krull, A., Buchholz, T.-O. & Jug, F. *Noise2void-learning denoising from single noisy images* in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (2019), 2129–2137.

118. Höck, E., Buchholz, T.-O., Brachmann, A., Jug, F. & Freytag, A. *N2v2-fixing noise2void checkerboard artifacts with modified sampling strategies and a tweaked network architecture* in *European Conference on Computer Vision* (2022), 503–518.

119. Batson, J. & Royer, L. *Noise2self: Blind denoising by self-supervision* in *International Conference on Machine Learning* (2019), 524–533.

120. Laine, S., Karras, T., Lehtinen, J. & Aila, T. High-quality self-supervised deep image denoising. *Advances in Neural Information Processing Systems* **32** (2019).

121. Krull, A., Vičar, T., Prakash, M., Lalit, M. & Jug, F. Probabilistic noise2void: Unsupervised content-aware denoising. *Frontiers in Computer Science* **2,** 5 (2020).

122. Prakash, M., Lalit, M., Tomancak, P., Krul, A. & Jug, F. *Fully unsupervised probabilistic noise2void* in *2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI)* (2020), 154–158.

123. Franzen, R. Kodak lossless true color image suite. *source: http://r0k. us/graphics/kodak* **4,** 9 (1999).

124. Martin, D., Fowlkes, C., Tal, D. & Malik, J. *A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics* in *Proceedings Eighth IEEE International Conference on Computer Vision. ICCV 2001* **2** (2001), 416–423.

125. Zeyde, R., Elad, M. & Protter, M. *On single image scale-up using sparse-representations* in *Curves and Surfaces: 7th International Conference, Avignon, France, June 24-30, 2010, Revised Selected Papers 7* (2012), 711–730.

126. Mannam, V. *et al.* Real-time image denoising of mixed Poisson–Gaussian noise in fluorescence microscopy images using ImageJ. *Optica* **9,** 335–345 (2022).

127. Wang, Z., Bovik, A. C., Sheikh, H. R. & Simoncelli, E. P. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing* **13,** 600–612 (2004).

128. Hore, A. & Ziou, D. *Image quality metrics: PSNR vs. SSIM* in *2010 20th international conference on pattern recognition* (2010), 2366–2369.

129. Knowles, R. B. *et al.* Translocation of RNA granules in living neurons. *Journal of Neuroscience* **16,** 7812–7820 (1996).

130. Das, S., Vera, M., Gandin, V., Singer, R. H. & Tutucci, E. Intracellular mRNA transport and localized translation. *Nature Reviews Molecular Cell Biology* **22,** 483–504 (2021).

131. Braselmann, E., Rathbun, C., Richards, E. M. & Palmer, A. E. Illuminating RNA biology: tools for imaging RNA in live mammalian cells. *Cell chemical biology* **27,** 891–903 (2020).

132. Pichon, X., Lagha, M., Mueller, F. & Bertrand, E. A growing toolbox to image gene expression in single cells: sensitive approaches for demanding challenges. *Molecular cell* **71,** 468–480 (2018).

133. Jalihal, A. P., Lund, P. E. & Walter, N. G. Coming together: RNAs and proteins assemble under the single-molecule fluorescence microscope. *Cold Spring Harbor Perspectives in Biology* **11,** a032441 (2019).

134. Brechbiel, J. L. & Gavis, E. R. Spatial regulation of nanos is required for its function in dendrite morphogenesis. *Current biology* **18,** 745–750 (2008).

135. Titlow, J. S. *et al.* Systematic analysis of YFP traps reveals common mRNA/protein discordance in neural tissues. *Journal of Cell Biology* **222,** e202205129 (2023).

136. Gala, D. S. *et al.* Mammalian glial protrusion transcriptomes predict localization of Drosophila glial transcripts required for synaptic plasticity. *bioRxiv,* 2022–11 (2022).

137. Vincent, J. B. *et al.* A translocation t (6; 7)(p11–p12; q22) associated with autism and mental retardation: localization and identification of candidate genes at the breakpoints. *Psychiatric genetics* **18,** 101–109 (2008).

138. Van der Walt, S. *et al.* scikit-image: image processing in Python. *PeerJ* **2,** e453. ISSN: 2167-8359. https://doi.org/10.7717/peerj.453 (June 2014).

139. Schneider, C. A., Rasband, W. S. & Eliceiri, K. W. NIH Image to ImageJ: 25 years of image analysis. *Nature methods* **9,** 671–675 (2012).

140. Tinevez, J.-Y. *et al.* TrackMate: An open and extensible platform for single-particle tracking. *Methods* **115,** 80–90 (2017).

141. Turner-Bridger, B. *et al.* Single-molecule analysis of endogenous β-actin mRNA trafficking reveals a mechanism for compartmentalized mRNA localization in axons. *Proceedings of the National Academy of Sciences* **115,** E9697–E9706 (2018).

142. Hu, Q. *et al.* Universal adaptive optics for microscopy through embedded neural network control. *Light: Science & Applications* **12,** 270 (2023).

143. Zhang, Y. *et al. A poisson-gaussian denoising dataset with real fluorescence microscopy images* in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2019), 11710–11718.

144. Chen, J., Chen, J., Chao, H. & Yang, M. *Image blind denoising with generative adversarial network based noise modeling* in *Proceedings of the IEEE conference on computer vision and pattern recognition* (2018), 3155–3164.

145. Zhang, D., Zhou, F., Jiang, Y. & Fu, Z. *Mm-bsn: Self-supervised image denoising for real-world with multi-mask based on blind-spot network* in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2023), 4188–4197.

146. Wang, Z., Fu, Y., Liu, J. & Zhang, Y. *LG-BPN: Local and Global Blind-Patch Network for Self-Supervised Real-World Denoising* in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2023), 18156–18165.

147. Dosovitskiy, A. *et al.* An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929* (2020).

148. Zhang, D. & Zhou, F. Self-supervised image denoising for real-world images with context-aware transformer. *IEEE Access* **11,** 14340–14349 (2023).

149. Neukomm, L. J., Burdett, T. C., Gonzalez, M. A., Züchner, S. & Freeman, M. R. Rapid in vivo forward genetic approach for identifying axon death genes in Drosophila. *Proceedings of the National Academy of Sciences* **111,** 9965–9970 (2014).

150. Hoppe, C. *et al.* Modulation of the promoter activation rate dictates the transcriptional response to graded BMP signaling levels in the Drosophila embryo. *Developmental cell* **54,** 727–741 (2020).

151. Mueller, F. *et al.* FISH-quant: automatic counting of transcripts in 3D FISH images. *Nature methods* **10,** 277–278 (2013).

152. Monnier, N. *et al.* Inferring transient particle transport dynamics in live cells. *Nature methods* **12,** 838–840 (2015).

153. Stork, T., Sheehan, A., Tasdemir-Yilmaz, O. E. & Freeman, M. R. Neuron-glia interactions through the Heartless FGF receptor signaling pathway mediate morphogenesis of Drosophila astrocytes. *Neuron* **83,** 388–403 (2014).

154. Katz, Z. B. *et al.* Mapping translation'hot-spots' in live cells by tracking single molecules of mRNA and ribosomes. *Elife* **5,** e10415 (2016).