

Washington University in St. Louis

## Washington University Open Scholarship

---

Arts & Sciences Electronic Theses and  
Dissertations

Arts & Sciences

---

Winter 12-15-2018

### Fast Objective Coupled Planar Illumination Microscopy

Cody Jonathan Greer

*Washington University in St. Louis*

Follow this and additional works at: [https://openscholarship.wustl.edu/art\\_sci\\_etds](https://openscholarship.wustl.edu/art_sci_etds)



Part of the [Computer Sciences Commons](#), [Neuroscience and Neurobiology Commons](#), and the [Optics Commons](#)

---

#### Recommended Citation

Greer, Cody Jonathan, "Fast Objective Coupled Planar Illumination Microscopy" (2018). *Arts & Sciences Electronic Theses and Dissertations*. 1706.

[https://openscholarship.wustl.edu/art\\_sci\\_etds/1706](https://openscholarship.wustl.edu/art_sci_etds/1706)

This Dissertation is brought to you for free and open access by the Arts & Sciences at Washington University Open Scholarship. It has been accepted for inclusion in Arts & Sciences Electronic Theses and Dissertations by an authorized administrator of Washington University Open Scholarship. For more information, please contact [digital@wumail.wustl.edu](mailto:digital@wumail.wustl.edu).

Washington University in St. Louis  
Graduate School of Arts and Sciences  
Department of Neuroscience

Dissertation Examination Committee:  
Timothy E. Holy, Chair  
Martha Bagnall  
Dennis Barbour  
Matthew Lew  
Steve Petersen

A Dissertation on Fast Objective Coupled Planar Illumination Microscopy

by

Cody Greer, B.S. Biology

A dissertation presented to  
the Graduate School  
of Washington University in  
partial fulfillment of the  
requirements for the degree of  
Doctor of Philosophy

December 2018  
Saint Louis, Missouri



# Contents

<b>List of Figures</b> . . . . .	<b>iv</b>
<b>Acknowledgments</b> . . . . .	<b>vi</b>
<b>Abstract</b> . . . . .	<b>viii</b>
<b>1 Introduction</b> . . . . .	<b>1</b>
1.1 Large-scale neuroscience . . . . .	1
1.2 Optical imaging of neural function . . . . .	5
1.2.1 Fluorescence contrast . . . . .	7
1.2.2 Fluorophores for functional imaging . . . . .	8
1.2.3 Basics of image formation . . . . .	10
1.2.4 Point-scanning versus widefield microscopy . . . . .	14
1.3 Optical sectioning strategies with widefield methods . . . . .	16
1.3.1 Planar illumination . . . . .	17
1.3.2 Computational sectioning . . . . .	23
1.4 Algorithmic neuroscience . . . . .	26
<b>2 Fast Objective Coupled Planar Illumination Microscopy</b> . . . . .	<b>28</b>
2.1 Design rationale . . . . .	29
2.2 Scanning faster with OCPI . . . . .	34
2.2.1 Optimizing the scanning control system . . . . .	38
2.2.2 Optimizing camera exposure timing during fast scanning . . . . .	39
2.2.3 Pulsed illumination . . . . .	44
2.2.4 Bidirectional imaging . . . . .	46
2.3 Acquiring images faster with MCIS . . . . .	47
2.3.1 Scalability of MCIS . . . . .	52
2.4 Identification and removal of heartbeat artifact in calcium imaging timeseries . . . . .	55
2.5 Discussion . . . . .	63
2.6 Methods . . . . .	66
2.6.1 Calculation of image quality in non-native focal planes . . . . .	66
2.6.2 Off-the-shelf components . . . . .	66
2.6.3 Custom components . . . . .	67
2.6.4 Calibration of piezo closed-loop controller . . . . .	67
2.6.5 Generation of smoothed triangle wave commands . . . . .	68

2.6.6	Scan range tuning . . . . .	68
2.6.7	Pulsed illumination . . . . .	69
2.6.8	Image timing calibration . . . . .	70
2.6.9	Interpolation of bidirectional images . . . . .	71
2.6.10	KEM alignment with focal plane for MCIS . . . . .	71
2.6.11	Camera alignment for MCIS . . . . .	72
2.6.12	Stitching MCIS images . . . . .	73
2.6.13	Zebrafish imaging . . . . .	74
2.6.14	Stripe removal . . . . .	74
2.6.15	Image registration . . . . .	75
2.6.16	Extraction of neuronal calcium traces . . . . .	75
2.6.17	Power spectral density and filtering of neuronal calcium traces . . . . .	76
2.6.18	Software . . . . .	76
<b>3</b>	<b>Analyzing data from a fast OCPI microscope . . . . .</b>	<b>77</b>
3.1	Before the experiment . . . . .	78
3.2	Computing with terabyte-scale imaging data . . . . .	81
3.2.1	Computing hardware . . . . .	81
3.2.2	Software design strategies . . . . .	86
3.3	Special considerations for light sheet microscopy data . . . . .	91
3.3.1	Point spread function . . . . .	91
3.3.2	Depth effects . . . . .	94
3.3.3	Stripe artifacts . . . . .	96
3.4	Special considerations for live imaging . . . . .	98
3.4.1	Photodamage, Photobleaching, and Phototoxicity . . . . .	99
3.4.2	Unintentional stimulation of neural systems . . . . .	100
3.5	Preprocessing pipeline for optical neurophysiology . . . . .	100
3.5.1	Image registration . . . . .	101
3.5.2	Extracting neuronal timeseries . . . . .	106
3.5.3	Spike deconvolution . . . . .	110
3.6	Future work interpreting neuronal dynamics . . . . .	112
3.6.1	Cross-correlation analysis . . . . .	113
3.6.2	Predictive modeling . . . . .	114
<b>Appendix A</b>	<b>Spherical aberrations from remote focusing . . . . .</b>	<b>121</b>
<b>Appendix B</b>	<b>Microscope components and drawings . . . . .</b>	<b>125</b>
<b>Appendix C</b>	<b>Derivation of Wiener filter for calcium deconvolution . . . . .</b>	<b>139</b>
<b>Appendix D</b>	<b>Sparsity of the weight matrix with linear regression . . . . .</b>	<b>146</b>
<b>References</b>	<b>. . . . .</b>	<b>150</b>

# List of Figures

1.1	Rendering of the <i>C. elegans</i> anatomical connectome . . . . .	2
1.2	Adjacency matrix summarizing the <i>C. elegans</i> anatomical connectome	4
1.3	Basics of image formation . . . . .	6
1.4	The fluorescence process . . . . .	8
1.5	The calcium response function . . . . .	10
1.6	Spherical aberration ray diagram . . . . .	11
1.7	Abbe Sine, Herschel, and perfect imaging systems . . . . .	12
1.8	Widefield versus point-scanning microscopy . . . . .	15
1.9	Light Sheet Fluorescence Microscopy . . . . .	17
1.10	Selective Planar Illumination Microscopy . . . . .	18
1.11	Original Objective Coupled Planar Illumination microscope . . . . .	19
1.12	Remote focusing with an electrically tunable lens . . . . .	20
1.13	Aberration-free remote focusing designs . . . . .	22
1.14	Light Field Microscopy . . . . .	25
2.1	OCPI microscope schematic (single-camera) . . . . .	29
2.2	Spherical aberrations when remote focusing . . . . .	31
2.3	Piezoelectric actuator response is distorted at high scan frequencies .	32
2.4	Volume imaging rate scales with camera frame height and sample thick- ness . . . . .	33
2.5	Current CMOS cameras are limited in their imaging rates . . . . .	34
2.6	Photo of custom light-weight scanning assembly . . . . .	35
2.7	Closed-loop control of piezoelectric actuator enables highly repeatable scanning . . . . .	37
2.8	Optimized scan system achieves 20 Hz scanning for volumes up to 700 $\mu\text{m}$ thick . . . . .	38
2.9	Timing camera exposures based on sensor output from the piezoelectric scanner . . . . .	40
2.10	Sensor-guided image acquisition is insufficient to specify the correct image plane during fast scanning . . . . .	41
2.11	Image-guided refinement of camera timing improves slice alignment during fast scanning . . . . .	41
2.12	Kinetic analysis of scan system . . . . .	43
2.13	Diagram of rolling shutter with a CMOS camera . . . . .	45
2.14	Bidirectional image acquisition doubles imaging rate . . . . .	46
2.15	OCPI with Multi-Camera Image Sharing . . . . .	49

2.16	Aligned and stitched images of fluorescent beads acquired with MCIS	50
2.17	Measuring image redundancy in an MCIS system . . . . .	51
2.18	MCIS image of a larval zebrafish brain slice . . . . .	51
2.19	Diagram of a 4-camera MCIS system . . . . .	52
2.20	Analysis of MCIS scalability . . . . .	54
2.21	Maximum intensity projection of neural activity in the larval zebrafish forebrain . . . . .	55
2.22	Extracted neuronal timeseries and correlations in zebrafish forebrain .	57
2.23	Power spectral density estimates of individual neurons show peaks corresponding to zebrafish heartbeat . . . . .	59
2.24	Anomalies in correlations at heartbeat frequency bands . . . . .	61
2.25	Changes in correlation values after removal of heartbeat frequency bands	62
3.1	Data flow with an OCPI microscope . . . . .	79
3.2	Imagine GUI window . . . . .	84
3.3	Imagine thread hierarchy . . . . .	85
3.4	Julia language benchmarks . . . . .	87
3.5	Lazy evaluation example . . . . .	90
3.6	Light sheet geometry . . . . .	93
3.7	Refractive indices of cellular components . . . . .	95
3.8	Image destriping results . . . . .	97
3.9	Preprocessing pipeline . . . . .	101
3.10	Parameterizing specimen motion . . . . .	103
3.11	Block registration . . . . .	104
3.12	Registration of zebrafish brain slice . . . . .	105
3.13	Neural activity in larval zebrafish registered to a brain atlas . . . . .	106
3.14	Cross-correlation analysis of a neural network . . . . .	113
3.15	Predictive analysis of a neural network with linear regression . . . . .	116
3.16	Fast OCPI with point-scanning . . . . .	120

# Acknowledgments

I would like to thank my advisor for his many intellectual contributions to the work described in this dissertation, but perhaps more importantly for providing me an inspiring scientific exemplar. Tim's consistent enthusiasm, ambition, and professional standards have exerted a great influence on me, and he managed to do this while at the same time allowing me a great deal of freedom to choose my own scientific path.

I thank my dissertation committee for their crucial role in bringing this work to fruition. I also thank the faculty and the students involved in the Cognitive, Computational, and Systems Neuroscience pathway for broadening the intellectual scope of my work, and also for financial support delivered through NIH T32 1T32NS073547-01. The work described in this dissertation would have been impossible without the skill and creativity of the Washington University School of Medicine machine shop: John Kreidler, John Witte, and Kevin Poenicke. Dae Woo Kim also provided (and continues to provide) invaluable software engineering expertise, developing software for instrumentation and image processing. I also thank the staff at the Washington University Center for Cellular Imaging, particularly James Fitzpatrick and Dennis

Oakley, for their continuing management of the core facility microscope described in this dissertation.

Finally I thank all members of the Holy lab for their contributions to my work, and for their companionship throughout my PhD studies. Likewise I owe much to the many graduate students and distinguished faculty within my department who have supported the research comprising this dissertation. A special thanks goes to Paul Taghert, who in collaboration with Tim obtained the grant that funded the development of the microscope described here (NIH R24 1R24NS086741-01). I also thank the NIH for additional support via R01 5R01NS068409-08.

Cody Greer

*Washington University in Saint Louis*

*December 2018*

## ABSTRACT OF THE DISSERTATION

A Dissertation on Fast Objective Coupled Planar Illumination Microscopy

by

Cody Greer

Doctor of Philosophy in Neuroscience

Washington University in St. Louis, December 2018

Research Advisor: Professor Timothy E. Holy

Among optical imaging techniques light sheet fluorescence microscopy stands out as one of the most attractive for capturing high-speed biological dynamics unfolding in three dimensions. The technique is potentially millions of times faster than point-scanning techniques such as two-photon microscopy. This potential is especially poignant for neuroscience applications due to the fact that interactions between neurons transpire over mere milliseconds within tissue volumes spanning hundreds of cubic microns. However current-generation light sheet microscopes are limited by volume scanning rate and/or camera frame rate. We begin by reviewing the optical principles underlying light sheet fluorescence microscopy and the origin of these rate bottlenecks. We present an analysis leading us to the conclusion that Objective Coupled Planar Illumination (OCPI) microscopy is a particularly promising technique for recording the activity of large populations of neurons at high sampling rate.

We then present speed-optimized OCPI microscopy, the first fast light sheet technique to avoid compromising image quality or photon efficiency. We enact two strategies to develop the fast OCPI microscope. First, we devise a set of optimizations that

increase the rate of the volume scanning system to 40 Hz for volumes up to 700  $\mu\text{m}$  thick. Second, we introduce Multi-Camera Image Sharing (MCIS), a technique to scale imaging rate by incorporating additional cameras. MCIS can be applied not only to OCPI but to any widefield imaging technique, circumventing the limitations imposed by the camera. Detailed design drawings are included to aid in dissemination to other research groups.

We also demonstrate fast calcium imaging of the larval zebrafish brain and find a heartbeat-induced motion artifact. We recommend a new preprocessing step to remove the artifact through filtering. This step requires a minimal sampling rate of 15 Hz, and we expect it to become a standard procedure in zebrafish imaging pipelines.

In the last chapter we describe essential computational considerations for controlling a fast OCPI microscope and processing the data that it generates. We introduce a new image processing pipeline developed to maximize computational efficiency when analyzing these multi-terabyte datasets, including a novel calcium imaging deconvolution algorithm. Finally we provide a demonstration of how combined innovations in microscope hardware and software enable inference of predictive relationships between neurons, a promising complement to more conventional correlation-based analyses.



# Chapter 1

## Introduction


### 1.1 Large-scale neuroscience

Nervous systems are staggeringly complex. The human brain contains on the order of  $10^{11}$  neurons and an even greater number of glial cells, though precise counts are still debated[34]. Each neuron can make 1000s of synaptic connections with other neurons, placing the total number of synapses in the trillions[86]. Other animals generally have fewer cells and synapses. Popular model organisms for neuroscience range in their numbers over many orders of magnitude, with the nematode *C. elegans* at the small end with 302 neurons. With current tools even *C. elegans* is too complex to simultaneously observe all aspects of the structure and function of its nervous system. Despite the scale of nervous systems, most scientific progress has been driven by the careful application of tools that extract data about a tiny fraction of the system, or else about the whole system at a coarse scale.

These tools continue to drive progress, but recently there has been great interest in developing tools that allow more comprehensive measurements. This interest is inspired in part by the success of the Human Genome Project and other projects with

the -omics suffix that exhaustively catalogue aspects of the biology of an organism. Neuroscientific hypotheses that implicate large numbers of cells separated by large distances remain difficult to test without more comprehensive techniques.

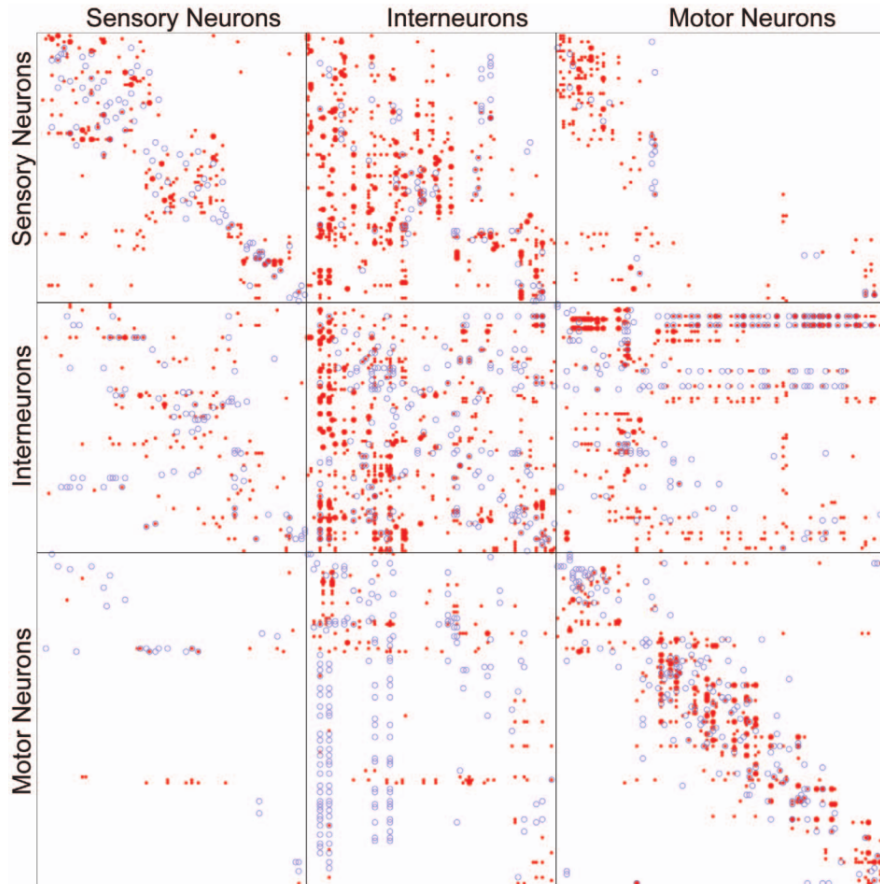
Some of the first comprehensive projects in neuroscience have been termed “connectomics” projects. These projects seek to fully characterize anatomical connections between brain regions and/or individual neurons. The first connectome was mapped in *C. elegans* in 1986 [96], and now efforts are underway in more complex organisms[3, 27, 61, 62]. While the Human Genome Project is an undisputed success, connectomics efforts in neuroscience often arouse skepticism[4]. This is because, differently from the genome, it is difficult to agree on how to define *comprehensive* in the context of a neural system. There is growing agreement that mapping anatomical connections is not comprehensive enough, and that one useful complement would be large-scale recordings of the dynamics of functioning neural systems[81]. As with many areas in biology, it is difficult to infer function from structure alone, and measurement of both aspects will provide much better constraints on interpretation.



**Figure 1.1:** 3D rendering of anatomical connections between all 302 neurons in the highly-stereotyped nematode *C. elegans* nervous system (open-worm.org). The data used to generate this rendering and Figure 1.2 were first published in 1986 [96].

Functional recording methods for neuroscience currently do not scale as well as methods for anatomical mapping. The main (perhaps obvious) reason for this is that functional recordings must be performed while the animal (or the tissue) is alive. Thus most of the physical and chemical transformations that aid anatomical mapping techniques cannot be employed when recording neural function. An additional challenge is that a comprehensive functional method must sample both in the spatial and temporal dimensions. Classically the temporal signals of most interest to neuroscientists have been the fast electrochemical signals mediating the propagation of information within and between neurons. This propagation is extremely fast in biological terms; in order to observe it the sampling rate must be high enough to keep up with synaptic transmission and local membrane potential changes within a neuron lasting less than 1 ms. Neuronal electrophysiology was developed to record these fast dynamics, and this has been a method-of-choice for decades of neurophysiology studies. While electrophysiology achieves high temporal resolution, it requires physically inserting an electrode into the animal, and currently it only scales to a few hundred neurons per insertion[40].

Recently optical imaging methods have grown in popularity due to their ability to scale to greater numbers of neurons (potentially 1000s) than electrophysiology. Optical imaging exhibits a fundamental tradeoff between the number of neurons sampled and the sampling rate; typical sampling rates are orders of magnitude slower than with electrophysiology. Researchers are quickly developing new optical techniques to improve the temporal resolution of large-scale optical imaging and bring it closer to parity with electrophysiology. Chapter 2 of this dissertation describes techniques developed in this spirit.



**Figure 1.2:** *C. elegans* Anatomical information from [96] as visualized by Varshney et al[94]. Neurons are classified by type. Gap junctions between pairs of neurons are marked with blue circles while chemical synapses are marked by red points. Only recently has it become possible to record the activity of all neurons in this network simultaneously at video rate in a live animal, an advance enabled by optical imaging[71].

In addition to describing a novel high-speed recording technique, this dissertation also addresses some of the challenges involved in interpreting fast functional imaging datasets. For example, there is not yet a consensus about how to extract functional counterparts to anatomical connectivity graphs like the one shown in Figure 1.2. This topic is rich enough to warrant an entire dissertation. However rather than exploring this topic in depth we focus on a simple and ubiquitous component of many functional analyses: *correlations* in the activity of neurons over time. Correlation (or

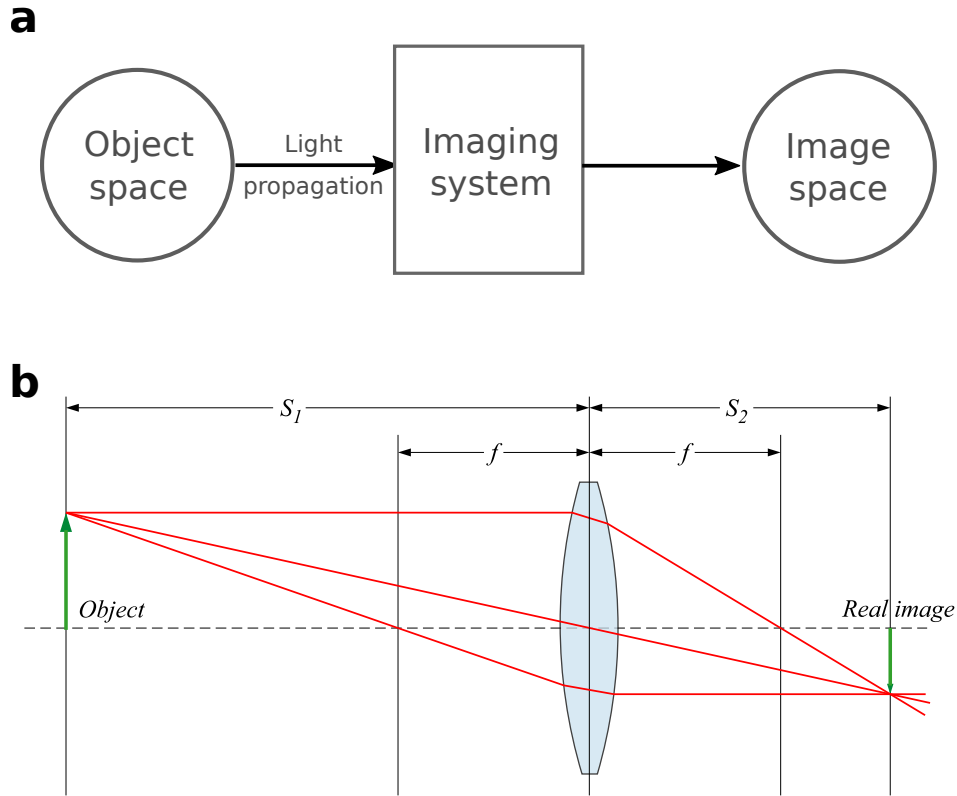
its un-normalized sibling, *covariance*) can be computed for each pair of neurons in a functional dataset and summarized in a matrix. In Section 2.4 we apply our fast imaging method to record neural activity in the larval zebrafish forebrain, and we examine the matrix of neuronal correlations. We find that the correlation measure is biased by an artifact resulting from the heartbeat of the fish, and we show that fast imaging enables removal of the artifact by filtering. Managing this and similar artifacts will be a crucial step for future studies aiming to generate comprehensive functional descriptions of neural circuits.

The remainder of this chapter gives further background information required to understand our fast imaging technique and its context.

## 1.2 Optical imaging of neural function

Neuroscientists use myriad optical tools in their work. The term “optical” encompasses any technique utilizing electromagnetic radiation. In this dissertation “light” will be used interchangeably with electromagnetic radiation even though “light” refers to only a subset of the possible wavelengths of radiation, and many of the principles discussed here also apply to wavelengths outside of that subset. Light can be harnessed in many ways to enable advances in neuroscience, but here we focus on a very common application: imaging. Imaging systems exploit the fact that objects (and regions of a single object) differ in their interactions with light. Most imaging devices, including the common cell phone, are designed to sample these interactions at a collection of spatial locations, with the collection of samples comprising an image. Thus the abstract function of an imaging device is to map points in an “object” space

to points in an “image” space. In practice this is achieved through a lens system that focuses light emitted from the object space to points in the image space (Figure 1.3).



**Figure 1.3:** (a) Abstract description of an imaging system as an operator that maps points in an “object” space to points in an “image” space using light. (b) A more practical description of a typical imaging system. A lens (or a multi-lens system) is designed to refract light emitted from points in the object space such that the light is refocused to a corresponding point in the image space. The ray diagram shown illustrates this process for a single point, but a conventional imaging system is also capable of focusing other points at the same horizontal distance from the lens.  $f$ : focal length of the lens,  $S_1$ : object distance,  $S_2$ : image distance. (Adapted from commons.wikimedia.org)

Devices vary in the mechanism employed to form images; some mechanisms are further described in Section 1.2.4. A microscope is simply an imaging device specialized for imaging small things; microscopes use varying degrees of *magnification* to capture details that are too small to be seen by eye. Most imaging systems can only form a clear image of a planar spatial region, though systems differ in the rate that


image quality degrades with distance from this plane. An explanation for this and a discussion of exceptional systems is deferred until Section 1.3.

Systems can be designed to sample many types of electromagnetic interactions. In microscopy the type of interaction sampled is generally chosen to ensure that interesting details in a specimen can be distinguished from uninteresting details, thus generating *contrast*. Contrast, when defined in this way, is a joint property of the sampling mechanism, the specimen, and the intentions of the microscope user. Some common ways to generate contrast are to sample the absorption, scattering, phase-altering, or fluorescence properties of the specimen. Neuroscience has long benefitted from contrast mechanisms that expose the structure of neural systems[14], but imaging of function became possible only relatively recently[33]. Fluorescence is currently the most direct contrast mechanism available for sampling the electrochemical activity of individual neurons. It is the mechanism employed by the novel microscope described in Chapter 2. Fluorescence microscopy is further described in the following section.

### 1.2.1 Fluorescence contrast

Fluorescence is the process by which a molecule can emit light after being “excited” by light of a shorter wavelength (and thus a higher energy), illustrated in Figure 1.4. Since emitted light is always of a longer wavelength than excitation light, it is possible to isolate only the emission light using filters. This ability to isolate light emitted by a particular molecule, or *fluorophore*, affords fluorescence microscopy its utility. Many biological molecules exhibit weak fluorescence; i.e. they emit a small amount of light when excited. Neurons and other components of tissue are naturally weakly fluorescent, but typically this “autofluorescence” does not provide sufficient contrast to observe interesting structure. Instead experimenters introduce

stronger fluorescent molecules into the specimen in a way that confines the molecules to particular structures of interest. After introducing fluorophores in this way one can image structures of interest by imaging the emitted fluorescence.



**Figure 1.4:** Fluorescence explained through a Jablonski diagram (image from [19]). Light of a certain wavelength is absorbed by the fluorophore, raising it from a low-energy ground state  $S_0$  to an excited state. The molecule may visit multiple intermediate excited states, eventually falling back to the ground state and emitting a photon of a longer wavelength (and thus lower energy) than the excitation light. The average time taken to emit the photon after the molecule is excited is termed the fluorescence lifetime. Note that the time axis in the diagram is somewhat misleading; typically the fluorescence lifetime is almost entirely dominated by the time spent in  $S_1$  before decaying to  $S_0$ , whereas the diagram implies that this transition is instantaneous.

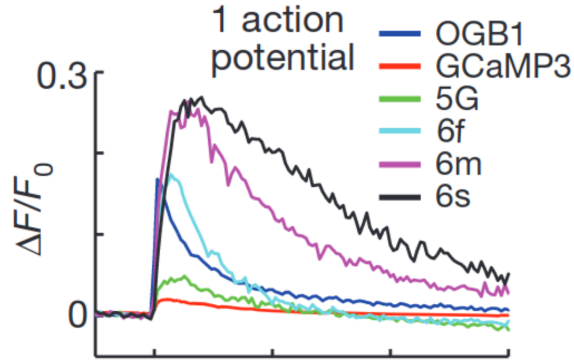
### 1.2.2 Fluorophores for functional imaging

Imaging dynamic processes in neurons involves additional considerations. Autofluorescence does not change with the moment-to-moment activity of the neuron and is thus an unsuitable contrast mechanism. Specialized fluorophores have been developed



to track dynamic processes. The earliest such molecules used in neuroscience were synthetic[33], but fluorescent genetically-encoded proteins are now more commonly used.

The imaging demonstrated in this dissertation utilizes a very popular family of genetically-encoded fluorescent proteins called GCaMP[17]. GCaMP proteins exhibit increasing fluorescence emissions with increasing ambient calcium ion concentration. This is suitable for tracking neural activity because intracellular calcium concentration increases sharply when a neuron generates an action potential. Neither GCaMP fluorescence nor calcium concentration is a direct measure of the electrical state of a neuron. Relative to the action potential, the fluorescence response is delayed and spread over a longer time period (Figure 1.5). Several computational methods have been developed in order to infer the timing of the action potential from calcium measurements. This deconvolution process is further described in Chapter 3 and a novel approach is described in appendix C. Very recently new proteins have been developed that provide fluorescence readouts of neurotransmitter release[56] and voltage[32, 84]. If these proteins can be sufficiently optimized they may replace calcium indicators for many use cases, and the imaging speed increases pioneered in this dissertation will only increase in relevance.



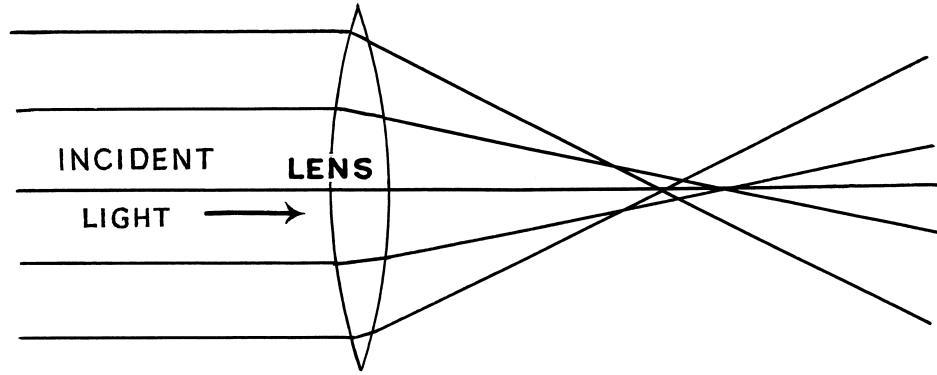
**Figure 1.5:** Shown are the fluorescence responses of several calcium-sensitive fluorophores due to single action potentials in dissociated rat hippocampal neurons (image from [17]). Fluorophores vary in their kinetics but all exhibit faster on-transients than off-transients. Note that all timecourses are much longer in duration than the action potential itself ( $\approx 1$  ms), motivating algorithms that extract spike timing from the calcium signal.

### 1.2.3 Basics of image formation

Before comparing fluorescence microscopy methods it will be helpful to describe the image formation process more concretely. As introduced by Figure 1.3, the task of any imaging system is to relay light originating from a source (the “object”) to another location (the “image”) *while preserving spatial information about the source*. The image may occupy a larger or smaller space than the object if the system *magnification* is less than or greater than 1, respectively. Magnifications of greater than 1 are generally more useful for imaging the tiny structures studied in neuroscience.

With or without magnification, imaging systems strive for 1:1 correspondences between spatial locations in the object and in the image. This is accomplished by a set of lenses between the object and the image. Lenses change the angle of incident light rays (through refraction) so that light emitted from a point in the object space gets refocused to another point in the image space. Typical lens systems are optimized to simultaneously map all points from a single 2D object plane to a single 2D image

plane. Light emanating from a plane outside of the object plane will generally not be focused to a single image plane. These imperfections in focus are termed *aberrations*. One type of aberration, *spherical*, is diagrammed in Figure 1.6.



**Figure 1.6:** A lens exhibiting spherical aberration is incapable of focusing incident light to a single point. Instead rays of light converge at various distances from the lens, as shown. A critical aspect of modern lens design is to adjust the curvature of the lens such that aberrations are minimized. (image source: [physicssimplifiedforyou.blogspot.com](http://physicssimplifiedforyou.blogspot.com))

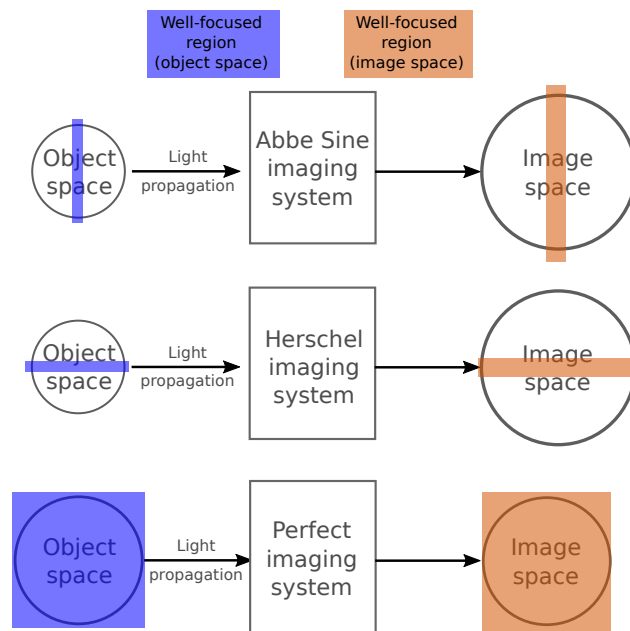
Lens systems that minimize aberrations in a single plane orthogonal to the optical axis satisfy the Abbe Sine Condition. Most lens systems for microscopy fall into this category. Lens systems can also be designed to minimize aberrations in a plane orthogonal to the optical axis. Such systems satisfy the Herschel condition[35]. If a system satisfies both the Sine and the Herschel conditions then it can achieve “perfect” aberration-free imaging of a 3D volume. It was shown by Maxwell[57] that this is only possible at a specific magnification

$$M = \frac{n_o}{n_i} \tag{1.1}$$

where  $n_o$  and  $n_i$  are the refractive indices of the media in object space and image space, respectively. Thus for a system working entirely in air or entirely in water

$M = 1.0$ . A common case in biological microscopy is that the object is immersed in water while the image is formed in air, implying that perfect imaging is possible only when  $M = 1.33$ .

Rather than having a single object plane and a corresponding image plane, it might be said that a perfect imaging system has an object *volume* and an image *volume*. The difference between Abbe/Herschel (magnifying) systems and a perfect imaging system is illustrated in Figure 1.7.



**Figure 1.7:** The conceptual diagram in Figure 1.3 has been modified to illustrate that the well-focused region (exhibiting minimal aberrations) differs depending on the optical condition satisfied by the lens system. A system meeting the Abbe Sine condition (the most common type of system) can focus light from a single plane orthogonal to the optical axis. A system meeting the Herschel condition can focus light from a single plane aligned with the optical axis. In order to avoid aberrations throughout an entire volume and achieve *perfect* imaging it is necessary to meet both conditions. Maxwell[57] showed that this only occurs when magnification  $\approx 1$  (The precise condition is given in equation 1.1).

Until recently it was widely believed that perfect imaging systems would not be useful for neuroscience because they provide insufficient magnification, but this was proven

incorrect by several recent microscope designs[26, 72, 25, 63, 28, 82, 100, 45]. These systems and the tradeoffs that they incur are discussed further in Section 1.3.1.

Even in an aberration-free system the correspondence between object and image points is never 1:1 due to the phenomenon of diffraction. The wave nature of light results in spreading of light from a single object point over a region of space in the image. The size of this region of spread is termed the *diffraction limit*. Usually this limit is quantified as first done by Abbe[1] in 1873 for points in an object plane orthogonal to the optical axis:

$$\frac{\lambda}{2NA} \tag{1.2}$$

where  $\lambda$  is the wavelength of light emitted from the object and  $NA$  is the numerical aperture of the system. Numerical aperture is defined as  $n \sin(\theta)$ , where  $n$  is the refractive index of the medium and  $\theta$  the half-angle of the cone of light collected by the lens system. Thus lens systems with higher NAs exhibit better *resolution* and also collect a greater fraction of the light originating from the object.

The above equation describes the *lateral* resolution of the imaging system. A full description of the system resolution must also consider the *axial* spread of light, which is the spread of light in the direction parallel to the optical axis. The axial resolution is described by a different equation:

$$\frac{2\lambda}{(NA)^2} \tag{1.3}$$


In order to acquire an image of a 3D sample often multiple 2D images are acquired at different axial offsets within the sample. Many strategies are available to achieve this *optical sectioning*. The following section will focus only on how a single 2D fluorescence image may be acquired, with 3D acquisition deferred until Section 1.3.

#### 1.2.4 Point-scanning versus widefield microscopy

Microscopes differ in the way that they capture a 2D fluorescence image of a specimen. Microscopes can be broadly classified as *point-scanning* or *widefield* (Figure 1.8). Point-scanning techniques acquire a single pixel of the image at a time, with each pixel corresponding to a spatial location in the specimen. Popular fluorescence point-scanning techniques either excite only a single point in the specimen at a time (2-photon microscopy) or excite a cone-shaped volume within the sample and block emission light emanating from all but one point (confocal microscopy). Thus collecting a 2D image requires *scanning* the optics laterally over multiple points in the specimen. Typically this is done with a set of rotating mirrors (galvanometers).

The fundamental difference between point-scanning and widefield is that lateral scanning is not required with widefield techniques. Scanning might be expected to result in a slower rate of image acquisition versus widefield imaging due to constraints on how fast the mirrors can move. However modern hardware such as the resonant galvanometer scanner already allows impressive scan rates, and further improvements might be considered a solvable engineering problem.

Yet there is a separate rate-limiting factor specific to *fluorescence* point-scanning methods: there is an important time constant associated with fluorescence emission. This constant, the fluorescence *lifetime*, describes the average amount of time it takes



**Figure 1.8:** Widefield microscopes (left) illuminate and image an entire plane of the sample onto a camera sensor. Point-scanning methods (right) such as confocal and 2-photon microscopy scan a focused laser across the sample, acquiring images one point/pixel at a time. Fluorescence-based point-scanning methods are fundamentally slower in their imaging rates, with pixel/voxel rates subject to a minimum dwell time[66]. (image from [9])

for an excited fluorophore to fluoresce (emit light). Thus in order to collect sufficient emission light a point scanning technique is subject to a minimum *dwell time*, the amount of time spent exciting and detecting at a single point in the specimen[66]. Lifetime is fluorophore-specific, typically on the order of a few nanoseconds. While this may seem brief, it results in a stringent limit on imaging rate when multiplied by the number of pixels (potentially millions) comprising a 2D image. Moreover it is common that only a modest number of fluorophores are present and able to contribute to a pixel, implying that it may be necessary to wait several fluorophore lifetimes to accumulate sufficient signal.

In a widefield microscope, on the other hand, fluorescence from many points in the object is detected simultaneously. Thus by parallelizing pixel acquisition widefield techniques exchange a minimum dwell time per-pixel for a minimum dwell time *per-plane*. This effectively eliminates dwell time as a consideration, and instead the

imaging rate of a widefield system is limited simply by the framerate of a modern CMOS camera. Section 2.3 introduces a new method for circumventing this camera rate bottleneck.

If widefield is potentially millions of times faster, why not use it exclusively for neuroimaging? With widefield it is generally more difficult to ensure that excitation/emission light is confined to a single plane in the sample. Until recently widefield techniques suffered crippling disadvantages in this ability to *optically section* samples. Details and recent innovations are discussed in the next section.


### **1.3 Optical sectioning strategies with widefield methods**

Perhaps the simplest and most widely used widefield fluorescence method is epifluorescence microscopy. Unfortunately this method provides very poor optical sectioning; excitation light is applied to the entire specimen, and the resulting emission light is sampled indiscriminately by the detector. This means that the image formed contains light from all planes of the sample. Recall that with a magnifying optical system only a single plane of the sample can be relayed in-focus to the detector. Images of other planes are blurred; and moreover they vary in their magnification (Appendix A describes this in detail). Thus epifluorescence images contain the superposition of a single well-focused object plane and many other poorly-resolved planes.




### 1.3.1 Planar illumination

Recently Light Sheet Fluorescence Microscopy (LSFM) has gained recognition as an alternative approach with much better sectioning capability. The core idea of LSFM is to use a *planar* illumination pattern so that only fluorophores within a single plane/sheet of the specimen are excited (Figure 1.9). Light sheet microscopes are generally designed so that this plane of illumination corresponds with the best-resolved plane in the object space. A 3D image of the specimen can then be obtained by assembling multiple well-focused images taken sequentially at different planes in the sample. This approach to optical sectioning has been implemented in several ways and is a major factor distinguishing variants of LSFM. Key LSFM variants are discussed in the following sections.



**Figure 1.9:** Generic light sheet fluorescence microscope design: Illumination light is confined to a single plane within the sample (actually a thin volume due to diffraction). An objective lens collects light emitted from the illuminated plane and forms a widefield image. In most LSFM designs the plane of the light sheet is orthogonal to the principle axis of the objective lens as shown, but LSFM also includes techniques that employ an oblique light sheet[25, 11]. (image from [37])

**Selective Planar Illumination Microscopy** Selective Planar Illumination Microscopy (SPIM) was one of the earliest-developed LSFM methods, and it continues to be one of the most popular (Figure 1.10). Excitation light is delivered through a lens oriented orthogonally to the lens that collects the emission light. Collimated excitation light is focused to a thin sheet by a cylindrical lens. Optical sectioning is accomplished by translating the sample along the axis orthogonal to the plane of the light sheet while capturing multiple 2D image slices. Precise translation is usually achieved with a piezoelectric positioning system.




**Figure 1.10:** Selective Planar Illumination Microscopy was one of the first LSFM designs and remains popular. 3D imaging is achieved by acquiring multiple planar optical sections while translating and/or rotating the sample, shown mounted in a cylindrical tube. cyl: cylindrical lens, ch: imaging chamber (filled with water), s: sample, det obj: detection objective. (image from [37])

**Objective Coupled Planar Illumination Microscopy** Objective Coupled Planar Illumination Microscopy (OCPI) was developed concurrently and independently from SPIM[36]. The primary difference between SPIM and OCPI is that the OCPI microscope accomplished optical sectioning by translating the *optics* rather than the

sample (Figure 1.11). Compared with SPIM this offers greater flexibility in specimen mounting and the potential for faster scanning.

Fast scanning with SPIM would require translating the specimen at high rates that can induce problematic side effects. These fast vibrations are an unwanted mechanical stimulus to the neural system under study that can bias neural activity or deform the tissue in complex ways that are difficult to correct with image registration. Moreover in order to move a load at high rates the piezoelectric positioner of the scan system requires a load-dependent calibration (discussed further in Chapter 2). With an OCPI microscope this calibration can be performed once (per imaging objective) whereas with a SPIM microscope the load is dependent on the specimen, requiring specimen-specific calibrations.




**Figure 1.11:** Objective Coupled Planar Illumination microscopy, developed concurrently and independently from SPIM, differs in that optical sectioning is achieved by translating optical elements rather than translating the sample. An adjustable coupler maintains the light sheet within the focal plane of the objective lens while the lens is translated. (image from [36])

An additional innovation of OCPI is that the excitation optics and imaging objective are both rotated  $30^\circ$  so that all optics are positioned above the sample. This allows imaging of a wide variety of specimen types and sizes ranging from zebrafish to

mouse. These advantages motivated the choice of OCPI over SPIM for the fast volumetric imaging microscope described in Chapter 2. In addition to SPIM and OCPI several other LSFM methods have been developed. The following sections describe alternative LSFM methods and their tradeoffs in fast imaging applications.

**Remote focusing methods** Scanning with OCPI involves translating components using a linear actuator (usually a piezoelectric device). Since modern objectives are infinity-corrected, this translation does not affect the quality of the image, which is formed using a conventional tube lens and acquired by a camera. This contrasts with “remote focusing” methods that image outside of the objective’s native focal plane, avoiding mechanical translation in order to scan faster. Remote focusing methods dynamically reposition the focal plane by using fast tunable lenses, scanning mirrors, or exploiting refractive index mismatches. Most of these methods have a critical disadvantage relative to SPIM or OCPI: they trade away image quality in favor of scan speed.




**Figure 1.12:** Shown is a LSFM method that achieves fast volume scanning using a lens with electrically tunable focal length. The lens (labeled ETL) is dynamically tuned to maintain focus on the light sheet while it is scanned through the sample[28]. This method and several others[20, 26, 72, 63, 28, 77] suffer from spherical aberrations when imaging outside of the native focal plane of the objective lens. (Image modified from [28]).

The nature of the tradeoff depends primarily on whether the microscope satisfies the condition for perfect imaging (equation 1.1). Any microscope that does not satisfy this condition suffers from spherical aberrations that worsen with increasing system NA, system magnification, and distance from the native focal plane. We plot in Figure 2.2 the theoretical imaging performance of a system with  $20\times$  magnification for various NA values. The aberrations are not a technical limitation of existing lenses, as we show in Appendix A that they follow directly from the Abbe sine condition and apply to any magnifying system that collects images away from the classical focal plane. Due to these aberrations such methods are unable to maintain cellular resolution when imaging specimens of moderate-to-large size [20, 26, 72, 63, 28, 77]. An example of such a system is shown in Figure 1.12.

Botcherby et al[10] introduced two system designs that avoid aberrations when using remote focusing (Figure 1.13). These designs can both be described as multi-stage microscopes. A first-stage microscope images a volume of the specimen with  $1\times$  magnification (or  $1.33\times$  when imaging from water to air) so that the perfect imaging condition is satisfied. Thus the entire image volume is relayed from the object space to an intermediate image space. A second-stage microscope then images a single plane of the relayed volume. Optical sectioning is accomplished by a scanning system that sequentially images object planes while maintaining each plane within the native focal plane of the second-stage imaging objective.

One of the designs by Botcherby et al accomplishes scanning via translation of an objective lens while the other requires translating only a mirror. The former design maintains high photon efficiency but fails to reduce the inertia of the scan system. The latter design has much less inertia, translating only a tiny mirror, but by design it loses 50% of the light collected by the objective lens due to the use of a beamsplitter.



**Figure 1.13:** Botcherby et al[10] introduced two optical designs that allow changing the plane of focus by translating optical elements located distant from the objective that images the specimen (remote focusing). Unlike previous (and many subsequent) remote focusing designs, these two designs satisfy the perfect imaging condition by choosing lenses so that the magnification from the object space to the intermediate image space respects equation 1.1. Additional magnification is provided in both cases by the final objective and a tube lens (tube lenses marked by red arrows). 3D imaging using the first design (top) requires translating a heavy objective lens (either L2 or L3). 3D imaging using the second design (bottom) requires translating only a small mirror, but half of the emission light is lost. (image adapted from [10])

Neither design was originally adapted for LSFM, but LSFM adaptations of these ideas emerged quickly[26, 72, 25, 63, 28, 82, 100, 45]. In all of these LSFM designs, the light sheet is generated by the same objective lens that collects emission light, incurring an additional cost in photon efficiency due to the extreme angles of light emanating from the intermediate image volume. A representative example[82] loses 79% of the light collected by the objective lens. There exists no theory suggesting that it is *impossible* to improve this efficiency and reach image parity with OCPI

and SPIM while maintaining fast scan speeds, but no method to do so has yet been discovered.

### 1.3.2 Computational sectioning

Several other widefield imaging techniques have been developed to perform optical sectioning without the use of a light sheet. Sectioning is performed by computationally unmixing light emanating from multiple image planes. Multiple methods can be employed to perform this unmixing. The following paragraphs briefly describe two popular groups of methods, Structured Illumination Microscopy (SIM) and Light Field Microscopy, and their tradeoffs for fast volumetric imaging.

**Structured Illumination Microscopy** SIM is a broad classification encompassing many different techniques that apply patterned (rather than uniform) illumination to the specimen. In SIM microscopy multiple snapshots of the specimen are taken at the same focus while varying the pattern of illumination. By recording images of the sample in multiple known illumination contexts it is possible to extract more information than would be possible with a single illumination pattern. SIM imaging systems can be specialized for many purposes[79] including optical sectioning [60].


SIM can perform optical sectioning without a light sheet by simple algebraic combination of corresponding pixels from a set of images taken under structured illumination[60, 79]. Unfortunately since multiple images are required, optical sectioning via SIM is inherently slower than (unstructured) LSFM by a factor equal to the number of illumination contexts (at least two[49]). Moreover SIM requires mechanical scanning similar to LSFM, so it has no advantage in this regard. While it has no speed

advantage, SIM may be combined with LSFM to improve image quality; structured light sheets have been used to improve LSFM resolution beyond the diffraction limit [42].

**Light Field Microscopy** Other computational widefield methods avoid the imaging speed penalty of SIM by attempting to solve a more difficult computational problem: *3D deconvolution*. Deconvolution methods attempt to redistribute light intensity across multiple pixels in multiple planes based on physical models of light propagation. Some degree of optical sectioning can be achieved by deconvolving images from an epifluorescence microscope taken at various depths of focus[52]. However since deconvolution is an ill-posed inverse problem[75], better results can be achieved by specialized microscopy methods designed to constrain the problem. Light Field Microscopy (LFM) is one such method[48].

LFM is capable of extracting many optical sections from a single camera image under uniform (volumetric) illumination *without scanning*. LFM constrains the deconvolution problem by operating on the *light field* of the sample. The light field includes not only spatial information but also angular information about light emitted by the specimen. Like conventional microscopy LFM captures information about *where* in the sample is fluorescing, but it also captures information along two angular dimensions: for each point in space an LFM records the intensity of light traveling along a set of angles within the collection cone of the objective lens. LFM achieves this by utilizing an array of microlenses to separate out angular information at a coarsely-sampled set of locations in the image plane. A set of computational steps, including 3D deconvolution, has been developed to extract a sectioned 3D volume from the 4D light field image. Further implementation details of LFM are omitted here in favor of discussing factors that are implementation-independent.





**Figure 1.14:** Diagram of a Light Field Microscope. Rather than recording a 2D image of the sample, such a microscope records a 4D *light field* that also includes angular information. The additional information provided by the light field enables optical sectioning via computation after the image is recorded. The light field is imaged by placing an array of microlenses in the typical image plane of the microscope. Each microlens projects its incident rays onto the camera sensor behind so that different incident ray angles arrive at different camera pixels, thus trading spatial resolution for angular resolution. The green rays trace an image of the objective’s back focal plane. (image adapted from [48])

When compared with LSFM, LFM has one major advantage and two major drawbacks. The advantage is that no scanning is required in order to acquire 3D information, so LFM is limited only by camera speed. One disadvantage is that LFM faces a difficult computational problem. The deconvolution problem is slow to solve (algorithmic complexity of  $O(n^2 \log n)$ ), and no current algorithm is guaranteed to find the best solution. The light field image provides better conditioning for the deconvolution problem than conventional images, but nevertheless optimization often converges to a suboptimal solution and, thus, to suboptimal images (the optimization

problem is nonconvex). Moreover unavoidable artifacts result from incomplete spatial and angular sampling of the specimen, especially near the native object plane[12].

The second major disadvantage of LFM is that lateral spatial resolution must be traded for angular resolution. The significance of this tradeoff is not obvious until one considers the relationship between angular resolution and optical sectioning with LFM. Improving angular resolution results in better-resolved optical sections. Critically the relationship between angular resolution (sectioning ability) and lateral resolution is nonlinear: an  $n$ -fold increase in the number of optical sections requires an  $n^2$ -fold deterioration in lateral resolution[48]. This means that the only way to improve sectioning ability without compromising lateral resolution is to use an image sensor with a higher pixel density, but unfortunately the size of the required sensor grows quadratically with volume thickness. These factors suggest that LFM, unlike LSFM, scales poorly to larger image volumes.

## 1.4 Algorithmic neuroscience

Advances in the rate of data acquisition, enabled by instruments like the microscope described in Chapter 2, are driving neuroscience research into what might be described as a new era in which *algorithms* are central to scientific advancement. While qualitative descriptions of data are still valuable, scientists rely increasingly on software to transform raw data into something more understandable. Within the imaging domain, for example, many techniques now generate more images than can possibly be examined by eye. The microscope described in Chapter 2 can generate images at a constant rate of up to 2 GB/s, with a single imaging experiment lasting up to an hour (7.2 TB). For perspective one can compare this to the Human Genome Project,

which produced a 3.3 GB human genome sequence over a period of 13 years. Thus we can now generate the equivalent of 2100 genomes per hour. This rate is increasing rapidly; we describe a method in Section 2.3 that can easily scale the microscope’s imaging rate by another order of magnitude.

Leveraging data of this scale will require careful choices in computational hardware and software. In Chapter 3 we overview considerations informing these choices when analyzing data from a fast OCPI microscope. We describe general software design strategies as well as a set of algorithms that we have developed for our analysis pipeline. The chapter gives particular attention to the tradeoff between computational efficiency and accuracy, an increasingly important tradeoff as we scale to larger datasets.

# Chapter 2

## Fast Objective Coupled Planar Illumination Microscopy


This chapter introduces a novel LSFM microscope designed for fast volumetric imaging. OCPI was chosen as the optical sectioning paradigm due to its advantages in image quality and scalability (introduced in Section 1.3 and discussed further in this chapter). Several innovations are introduced in order to increase the speed and robustness of volume scanning, which is typically a rate-limiting factor for OCPI and SPIM microscopes. Section 2.3 of this chapter introduces another innovation, Multi Camera Image Sharing (MCIS), that incorporates multiple cameras in order to increase the imaging speed of the system. MCIS is useful for speeding up OCPI as well as any other widefield imaging technique (see Section 1.2.4 for an introduction to widefield imaging).

This chapter also demonstrates the utility of the microscope for neuroscience, documenting imaging of neural activity in the whole brain of a larval zebrafish at 10 Hz with 0.65  $\mu\text{m}$  resolution as well as imaging just the forebrain at 20 Hz. The forebrain recording is analyzed to show that neuronal traces are contaminated with an artifact

arising from the heartbeat of the fish. An analysis concludes that this artifact introduces spurious correlations between neurons, and that a sampling rate of at least 15 Hz is required in order to remove the artifact. Expunging this artifact will be an important processing step for future large-scale studies of zebrafish brain activity.

This microscope has successfully served more than a dozen users in a centralized imaging facility, demonstrating its flexibility and robustness. Detailed hardware schematics are provided in Appendix B in order to facilitate adoption by other research groups.

## 2.1 Design rationale

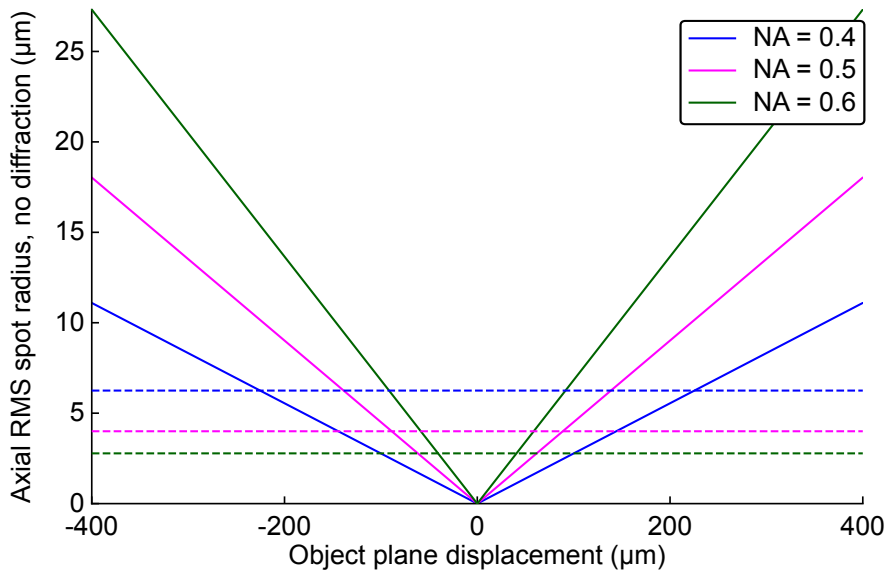


**Figure 2.1:** Schematic of an OCPI microscope with minor modifications to the design described in previous work[93, 36]. Optical sectioning is achieved by translating the optics for generating the light sheet together with the detection objective (boxed components). Inertia of these components limits the rate that a volume can be scanned. Abbreviations: PC: pigtailed collimator, LLP: light sheet lens pair, OBJ: objective, BM: broadband mirror, EF: emission filter, TL: tube lens.

This section conveys our rationale for optimizing volume imaging rate with OCPI microscopy. OCPI is similar to SPIM[39], but was the first implementation of LSFM to achieve volumetric imaging by translating the optics instead of the sample[36]. OCPI also introduced a 30–45° tilt in the optics to facilitate observation of extended horizontal samples, such as neuronal tissue slices or *in vivo* preparations, while minimizing the path length of both the illumination and emission light through the sample. Finally, from the outset OCPI reduced weight and the geometric hindrances that would arise from having two objective lenses by generating the light sheet using custom optics. Figure 2.1 illustrates the basic OCPI microscope design and shows that both the imaging objective and the lightsheet optics are translated together in the axial direction relative to the objective.

Scanning with OCPI involves translating components using a linear actuator (usually a piezoelectric device). Since modern objectives are infinity-corrected, this translation does not affect the quality of the image, which is formed using a conventional tube lens and acquired by a camera. This contrasts with “remote focusing” methods that image outside of the objective’s native focal plane, avoiding mechanical translation in order to scan faster. These methods have a critical disadvantage relative to SPIM or OCPI: they trade away image quality in favor of scan speed (discussed in Section 1.3.1). We include in Appendix A an analysis of spherical aberrations in remote scanning systems, and we plot in Figure 2.2 the theoretical performance at 20× magnification for various NA values.

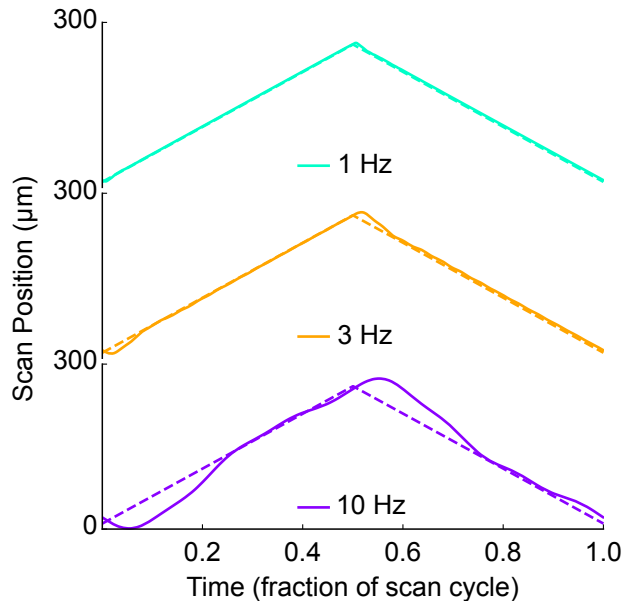
At a fairly low NA of 0.4, diffraction-limited axial resolution is limited to an area within 150 μm from the focal plane, and this degrades rapidly to only 25 μm for an objective with NA 0.6. In many fields, especially neuroscience, there is great interest in imaging axial spans of hundreds of microns efficiently at high spatial and temporal



**Figure 2.2:** Theoretically-estimated axial spherical aberrations as a function of displacement from the native focal plane when using remote focusing at  $20\times$  magnification (see supplement for derivation). Plotted are spot sizes (in object space) at various NAs. Unlike remote focusing methods, OCPI microscopy does not suffer these aberrations and is limited only by the diffraction limit (dashed lines, calculated with Abbe’s formula  $2\lambda/(NA)^2$  at wavelength  $\lambda = 500$  nm). We sought to preserve this optical advantage while bringing the scan speed of OCPI closer to these remote focusing methods.

resolution. Thus we sought to bring the scan rate of OCPI closer to the state-of-the-art while maintaining OCPI’s advantages in photon efficiency and resolution.

Volume imaging with SPIM and OCPI is usually rate-limited by the inertia of the sample or the optics, respectively. Due to this inertia the piezo positioner fails to follow high frequency commands (Figure 2.3). This failure can have catastrophic consequences for 3D imaging at high rates. Rapid acceleration during a camera exposure results in blurred images spanning multiple depths in the sample. Unwanted oscillations in the scan system can also result in a sequence of optical sections with depths that are not strictly increasing/decreasing. Particular high-frequency commands may




**Figure 2.3:** Examples of triangle wave scan commands (dashed lines) and measured responses (solid lines) from the piezoelectric positioner of an unoptimized OCPI microscope. One full scan cycle is shown at each of the three scan rates. Fast volume imaging is challenging because the measured response is distorted at higher scan rates due to inertia and imperfect closed-loop (PID) control.

even send the piezo system’s closed-loop controller into an unstable regime that destroys the device.


Were one to resolve the scanning bottleneck, the scan speed would then be limited by the camera framerate. The maximum volume imaging rate, achievable by a hypothetical system limited only by a modern camera (PCO.Edge 4.2), is illustrated in Figure 2.4. The achievable volume rate is inversely proportional to the size of the volume and the density at which the volume is sampled. Maximum readout rates of the scientific CMOS cameras produced by top manufacturers are all similar (Figure 2.5). Moreover these rates have not improved within 6 years, suggesting it may not be fruitful to wait for faster cameras.





**Figure 2.4:** If volume scanning were not rate-limiting, OCPI microscopy volume imaging rate would scale with the “height” of each camera image as well as the thickness of the imaged volume (shown for  $10\times$  magnification with  $5\ \mu\text{m}$  spacing of optical sections, PCO.Edge 4.2 camera). Our first goal was to optimize the scan system so that scan rate was not the limiting factor in most of the parameter space shown.

Below we describe an OCPI system that mitigates the camera rate bottleneck and also removes the scanning bottleneck over much of the parameter space described in Figure 2.4.




**Figure 2.5:** Maximum pixel rates (voxel rates when imaging a volume) of popular 16-bit scientific CMOS cameras are shown. All current widefield microscopy methods, including OCPI, are limited to these maximum rates if there are no other limiting factors such as volume scanning rate. Notably no high-sensitivity camera released within the past 6 years has improved upon the fastest rate shown, motivating our second goal to develop a method that circumvents this limit.

## 2.2 Scanning faster with OCPI

In order to address the scanning bottleneck we enacted five design strategies:

1. minimizing the mass of all translated components
2. optimizing the command signals that drive the piezoelectric actuator
3. acquiring image stacks during both the forward and reverse sweeps of the scan
4. calibrating the timing of camera exposures
5. pulsing the illumination laser during the global exposure period of the camera.

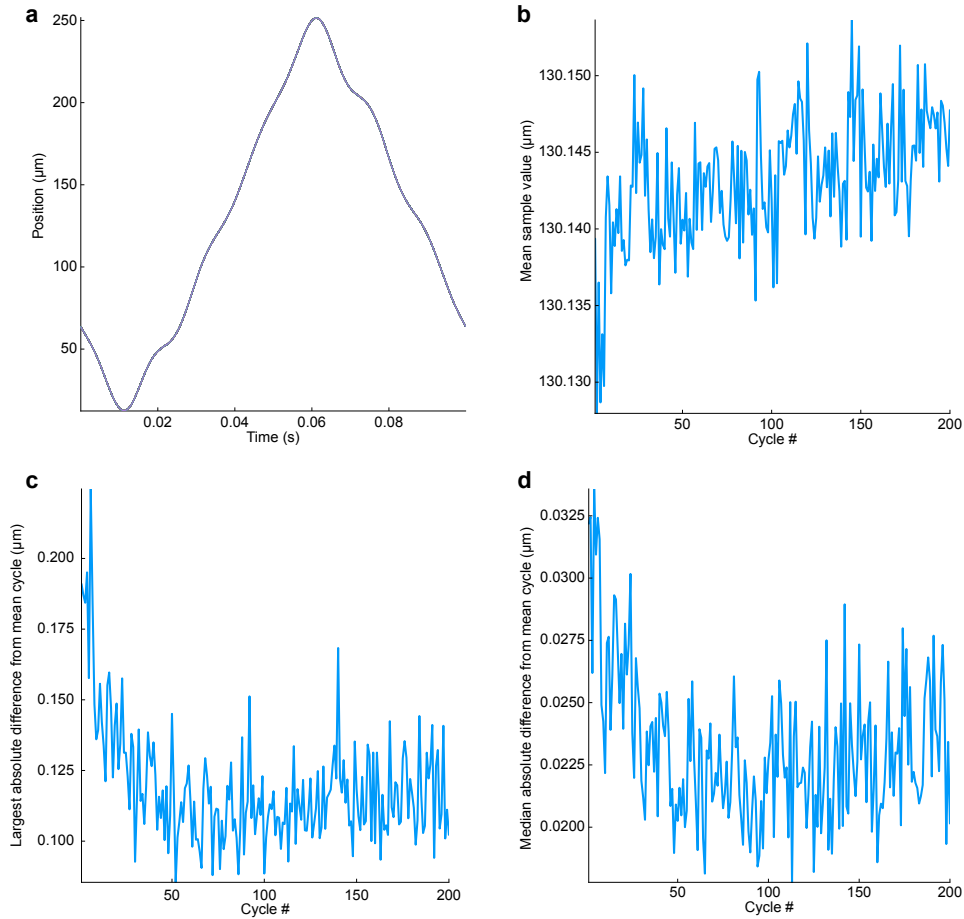


**Figure 2.6:** Photo of the scanning assembly (boxed region of Figure 2.1). Custom optics and machining minimize inertia. Drawings of components are provided in Appendix B.

A photo of our mass-optimized scan assembly is shown in 2.6. Rather than create the lightsheet with a second objective we used custom optics of minimal size[36], and mounting and alignment hardware was machined to minimize mass while maintaining rigidity. All components were attached to an aluminum backplane mounted to the piezoelectric positioner.

Our chosen positioner was able to generate push and pull forces of up to 100 N. This is more than enough force to achieve the accelerations necessary to scan at frequencies up to 20 Hz. Yet despite this potential, to our knowledge OCPI and similar light sheet scanning methods have never demonstrated scan rates greater than 4 Hz, and this was with only a  $32\ \mu\text{m}$  scan scan range[65]. One reason for this is that faster scanning requires more careful tuning of the magnitude and timing of the piezo force output. In many positioners, ours included, force output is calculated by a closed-loop controller utilizing the user’s command signal as well as feedback from a sensor. The controller counteracts the piezoelectric phenomena of creep and hysteresis[2].


Performance of a closed-loop piezo system is particularly sensitive to the mass of translated components, center of gravity, and translation angle relative to gravity. The controller in our system is of the proportional integral derivative (PID) variety, and thus it has three tunable feedback parameters. We requested that the vendor optimize these parameters for the mass of our assembly and angle of translation. Additional tuning was performed manually so that the system's response matched high-frequency scanning commands as closely as possible (see methods). We also verified that the response of the piezo system to a cyclic command is highly consistent across cycles after a brief warmup period (Figure 2.7). This repeatability is crucial to enable stable multi-stack recordings.



**Figure 2.7:** Consistency of the piezoelectric scanner response. **(a)** Overlay of 200 consecutive piezo scan cycles as measured by the capacitive sensor of the device (cycles during the first 20 s of operation are excluded). The command signal was a 10 Hz triangle wave. Samples were acquired at 100 kHz and downsampled to 10 kHz for plotting. **(b)** Closed-loop control of the piezo prevents drift (“creep”) in the mean piezo response measured during each cycle from (a). **(c)** Piezo response cycles are also consistent on a per-sample basis after an initial settling period. First a mean response cycle was created by averaging corresponding samples across the 200 cycles. Plotted is the maximum absolute difference of any sampled value of each cycle from the corresponding value in the mean cycle. Before computing differences each cycle was lowpass filtered with a gaussian kernel of width 100  $\mu\text{s}$  to reduce sampling noise. **(d)** Similar to (c), but the median difference between corresponding samples in each cycle is shown.

### 2.2.1 Optimizing the scanning control system

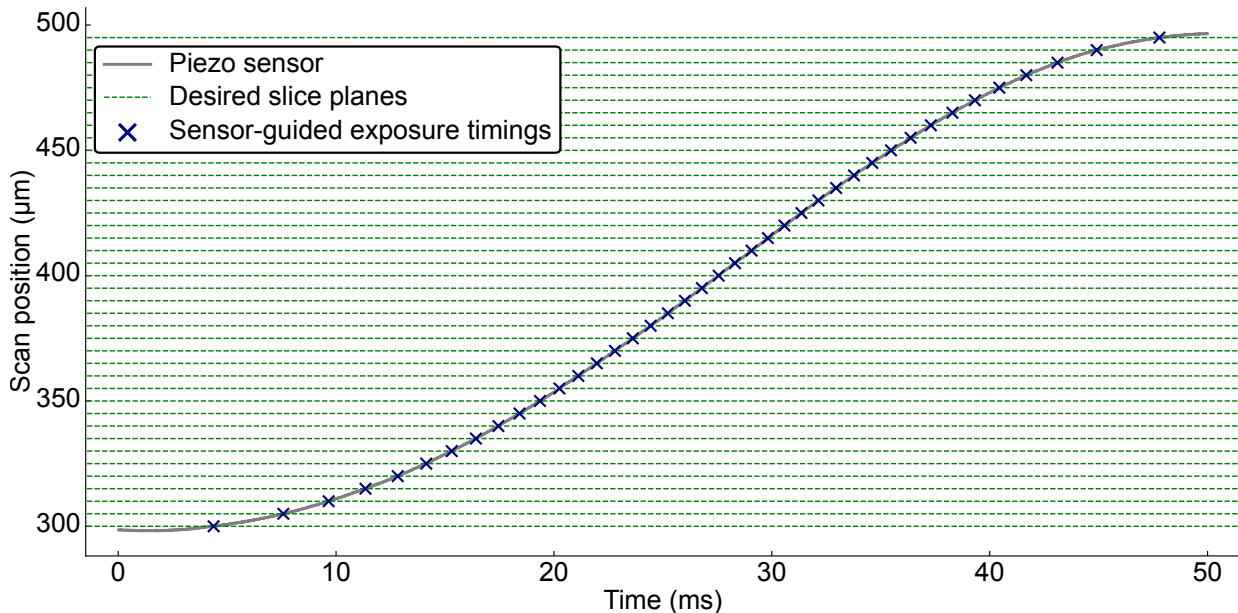
PID control is quite sensitive to large accelerations in the command signal such as those at the extrema of a triangle wave, causing the system to exhibit unfavorable higher frequency oscillations (Figure 2.3). We addressed this issue by utilizing a lowpass filtered triangle wave command, where the filter had a cutoff of  $3.25\times$  the scan frequency (see methods). We also performed a brief iterative optimization of the amplitude and offset of the command signal to achieve the desired scan range as measured by the sensor (see methods). Thus by tuning both the PID and the command waveform we were able to drive the piezo smoothly through a range of up to  $700\ \mu\text{m}$  at frequencies up to 20 Hz (Figure 2.8).



**Figure 2.8:** Distortion of the piezo response is reduced relative to Figure 2.3 by driving a tuned scan system using a lowpass-filtered command. Dashed lines illustrate that the measured amplitude does not match the command. The command was optimized iteratively (see methods) to generate the desired  $700\ \mu\text{m}$  range at a 20 Hz scanning rate.

### 2.2.2 Optimizing camera exposure timing during fast scanning

For later analysis steps such as image registration, it is desirable for the slices of an image stack to be equally spaced in the axial direction. Because of the non-uniform speed of piezo movement, this implies that collecting camera frames with a fixed framerate would result in images that are not evenly sampled in space. Therefore we utilized the measured piezo waveform to time the acquisition of individual slices so that they were equally spaced along the scan axis. Figure 2.9 illustrates this piezo sensor-guided approach: an image is taken whenever the piezo sensor trace intersects with a desired slice plane ( $5\ \mu\text{m}$  slice spacing shown). Notably the time intervals between slices are not uniform; slices near the extrema of the range are separated by longer intervals due to reduced scan velocity. Resonant galvanometer-based imaging systems exhibit similar nonuniformity in the angular velocity of the mirror, and these systems set pixel timing in the same way that we set frame timing[15]. Since piezo cycles are consistent after a warmup period (Figure 2.7), a single measured cycle was sufficient to determine the timing of camera exposures throughout a multi-stack recording.




**Figure 2.9:** A 10 Hz scan of a  $200\ \mu\text{m}$  volume overlaid with dashed lines marking depths at which to acquire images with  $5\ \mu\text{m}$  spacing between optical sections. When output from the piezo sensor is used to guide image acquisition, an image is acquired at each intersection of the sensor trace with a dashed line (marked with x’s). Note the uneven spacing of the x’s along the time axis.


However, we found that the sensor-guided approach to exposure timing was insufficient to yield images in the correct slice plane. Figure 2.10 (left panel) compares two images—one taken during fast scanning and the other taken statically—that were collected at nominally the same plane. Despite the fact that the measured piezo position was the same in both cases, there is a poor correspondence between the images.

Initially we expected this to be explained by lag in the mixed analog and digital circuit that conveyed the piezo sensor signal. However to our surprise we could not explain the inaccuracy with a simple lag, gain, or offset of the sensor signal. We therefore developed a procedure to determine the correct timing for each image slice empirically by acquiring images at various temporal offsets from the “naive” sensor-based timing (Figure 2.11, methods).



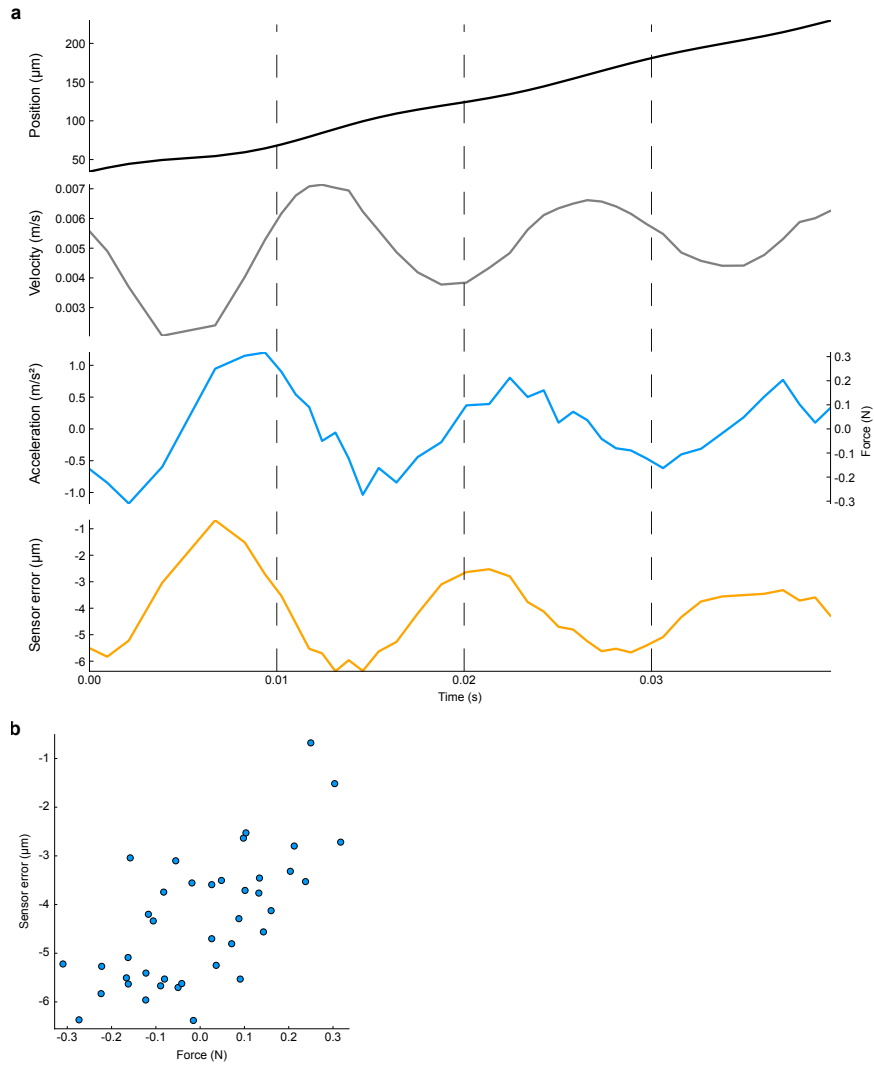


**Figure 2.10:** Sensor-guided image acquisition timing is insufficient to specify the correct image plane during fast scanning. A dynamically-acquired image of fluorescent beads is overlaid with an image taken when the scanner was stationary and, according to the piezo sensor, in the same image plane. An additional image-guided timing calibration corrected this inaccuracy and ensured that each slice of the stack was located in the correct plane (see methods, Figure 2.11). Scalebar: 5  $\mu\text{m}$ .



**Figure 2.11:** (a) A single piezo response cycle as measured by the piezo's built-in capacitive sensor. The dashed line marks a position at which the user desires to acquire an image. The timing of the intersection (marked with arrow) sets the initial guess for the timing of image acquisition. The procedure illustrated in panels b-e is applied independently to refine the timing for each plane to be imaged. (b) The initial guess from (a) is inaccurate for fast dynamic recordings (same as left panel in Figure 2.11). (c) The initial guess is refined by acquiring dynamic images at various temporal offsets and choosing the offset that yields an image matching the static template. The offset that yielded an image with minimal dissimilarity (see methods) is marked with an asterisk. (d) When the images corresponding to the optimal timing from (c) are overlaid, the beads are well aligned. (e) The same procedure is applied to the reverse stack. (f) The timing refinement process is applied to multiple slices. (g) The optimal timing offset for each slice is plotted against slice depth.

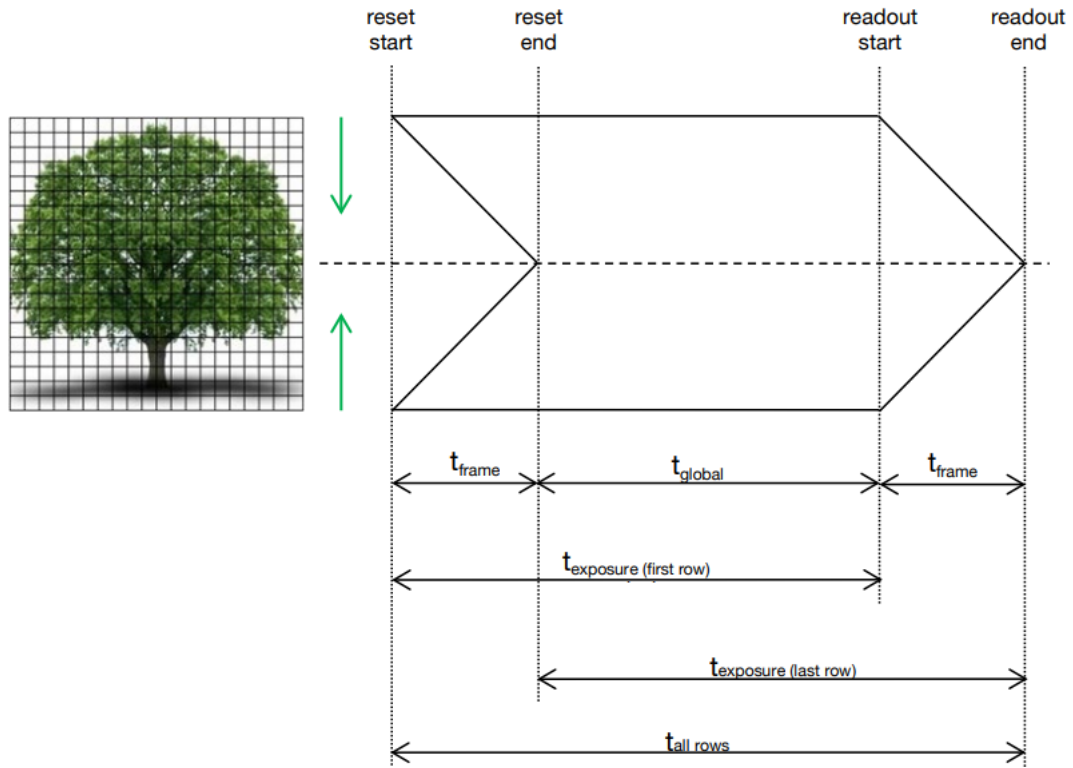
The timing calibration procedure allowed us to associate each timing correction with an axial displacement of the focal plane. The optimal timing offset was not a constant as would be expected if the error were due to a simple temporal lag in the sensor response. Likewise the optimal timing was not predicted by the depth in the sample, as would be expected from a simple error in sensor gain (Figure 2.11). We found that the difference between the piezo sensor value and the true focal plane position was best predicted by the acceleration of the scan system, suggesting that the error may be related to mechanical forces acting on the system (Figure 2.12).



**Figure 2.12:** (a) Error in scan position is better predicted by acceleration/force than by lower-order kinetics. Plotted are kinetic parameters of the piezo during the “forward” sweep of the cycle shown in Figure 2.11. Only timepoints corresponding with the timing of each image are plotted. The raw sensor trace was first lowpass filtered with a gaussian kernel of width  $300\ \mu\text{s}$  to reduce sampling noise. First 3 traces, from top to bottom: position, velocity, acceleration/force (force based on 264 g load of the piezo). The fourth trace shows for each image slice the difference between the measured position and the actual position of the focal plane as determined in the procedure described in Figure 2.11. This error trace aligns better with the acceleration/force trace than with the other traces. (b) Scatter plot of points from the blue force trace in (a) showing that force is predictive of axial displacement of the focal plane.

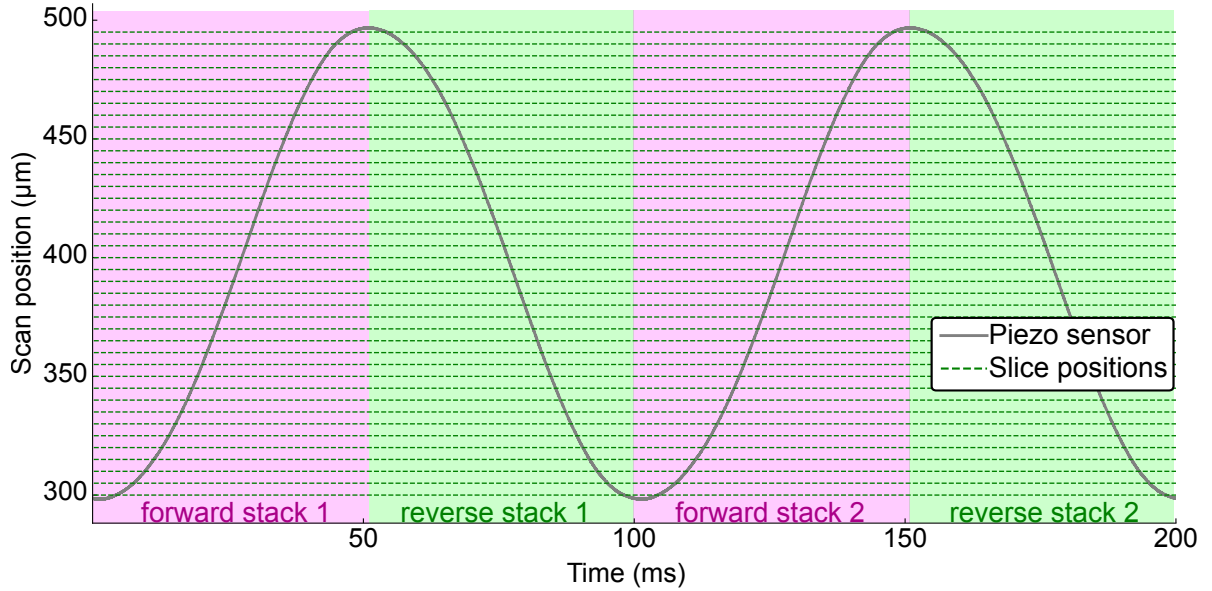
### 2.2.3 Pulsed illumination

In addition to calibrating exposure timing we also employed an optimized illumination protocol to achieve precise optical sectioning during fast scanning. The simplest (suboptimal) approach to illumination is to keep the excitation laser active throughout each stack and rely only on camera exposure pulses to delineate slices. This approach is problematic for CMOS cameras operating with a rolling shutter. With a rolling shutter the start and stop time of the exposure differs for each row of pixels in a frame (see Figure 2.13). Since the scan system is constantly in motion this implies that each row of pixels samples a slightly different axial plane. In order to prevent this contamination of an image with information from multiple axial planes we utilized pulsed illumination. By using a well-timed pulse with a duration much shorter than the camera exposure time we were able to ensure that photon integration occurred only during the “global” exposure period during which all rows of pixels on the camera chip are exposed simultaneously (see methods).



**Figure 2.13:** Shown is a timing diagram of for the rolling shutter of the PCO.Edge family of CMOS cameras (reprinted from camera manual). Each frame is acquired in two overlapping phases, *exposure* and *readout*.  $t_{all\ rows}$  is the time that it takes to expose and read an entire image. Note that this is a longer duration than  $t_{exposure}$  because exposure and readout for each row of pixels starts and stops at a different time. During the interval labeled  $t_{global}$  all rows of pixels are exposing simultaneously, making this an ideal time to flash illumination light. After a row is read exposure can begin again immediately (not shown), meaning that light collected outside of the  $t_{global}$  period will be divided between two camera frames. Note also that  $t_{global}$  is very brief when imaging at the full framerate of the camera (one line time,  $10\ \mu\text{s}$  with the PCO.Edge 4.2).

## 2.2.4 Bidirectional imaging



**Figure 2.14:** Volume imaging rate was increased by an additional factor of two by imaging each plane twice per scan cycle. Thus a 20 Hz imaging rate is achieved for a scanning rate of 10 Hz. We compensated for non-simultaneous and non-uniform temporal spacing by interpolating each adjacent pair of stacks, resulting in virtual stacks aligned with the transitions between colored regions.

Since each depth in the volume is visited twice per scan cycle we further increased speed by utilizing a bidirectional imaging pattern: each scan cycle includes a “forward” stack and a “reverse” stack. Thus the volume imaging rate is equal to twice the scan rate. Image-guided timing was optimized separately for slices of the forward and reverse stacks (Figure 2.11). Figure 2.14 diagrams four image stacks acquired in this manner during two scan cycles. Under the bidirectional paradigm the time interval between consecutive images of a plane is not constant. More precisely, the interval is constant only for the slice in the center of the scan and becomes less uniform closer to the extrema of the scan cycle. Another more common (and commonly ignored) complication with volumetric timeseries data is that the slices within each


stack are not acquired simultaneously. We addressed both of these timing complications by interpolating each consecutive pair of stacks in time, yielding stacks with a virtual sampling time corresponding to the lines of color transition in Figure 2.14 (see methods). While inferior to truly simultaneous sampling, we expect that this correction will improve the fidelity of timing-sensitive analyses. Taken together with the scan capability shown in Figure 2.8, our system would be able to achieve 40 Hz bidirectional imaging of a 700  $\mu\text{m}$  volume if not for the limit imposed by the frame rate of the camera. 40 Hz is a milestone because it approximates the Nyquist sampling rate of the calcium indicator on-transient (The on-transient of GCaMP6f lasts about 50 ms[17]).

### 2.3 Acquiring images faster with MCIS

Thus with careful design and modern hardware, OCPI microscopy is not limited by scan rate when imaging medium-to-large volumes. Therefore further speed improvements for large-scale OCPI microscopy and other fast LSFM variants must come from mitigating the camera framerate bottleneck. We devise MCIS to exploit a feature of modern CMOS camera design: maximum framerate depends on the size of only one dimension of the image (“Image height” in Figure 2.4). This in turn means that volume imaging rate scales with only two of the three image dimensions. MCIS employs hardware to divide the image volume into two halves and relay the halves to two different cameras (Figure 2.15). The volume is divided along the rate-limiting axis of the camera sensor so that each half can be imaged at twice the maximal rate that a single camera can capture the full volume. We cut the images by positioning the apex of a “knife-edged” mirror (KEM) in the focal plane of the tube lens, introducing a 90° fold along the center of the focal plane. The two halves of the focal plane are then


relayed to the two cameras. The exposures of the two cameras were synchronized, and their images were later “stitched” back into a single image with custom software (see methods).





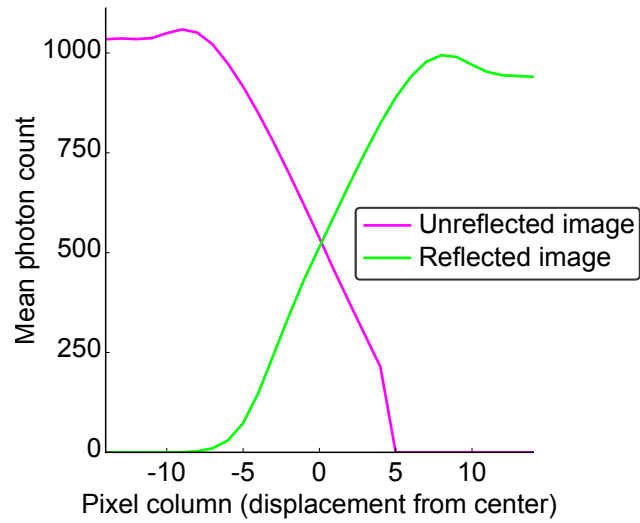
**Figure 2.15:** OCPI system employing Multi-Camera Image Sharing over two cameras. A knife-edged prism mirror (KEM) takes the place of the camera sensor in Figure 2.1. The mirror is aligned so that half of the image is reflected and relayed to a camera above while half passes unimpeded and is relayed to a second camera. Cameras are aligned so that they image a centered horizontal band in the field of view. The two cameras expose synchronously, and their images are later stitched together into a full image. Since the frame rate of a CMOS camera depends only on image “height” (Figure 2.4) this doubles the imaging speed of the system.

If the apex of the KEM is not precisely in the focal plane of the tube lens then a strip of the image will be captured redundantly (but with reduced intensity) on both cameras. Figure 2.16 demonstrates in pseudocolor the alignment and overlap of stitched images of a fluorescent bead sample. Beads along the edge of the KEM are imaged on both cameras.




**Figure 2.16:** (left) Example stitched image of fluorescent beads ( $0.2\ \mu\text{m}$  diameter) with one camera's image in magenta and the other in green. Scale bar:  $20\ \mu\text{m}$ . (right) Zoomed view of the rectangular region marked in the left panel showing a pair of beads in the narrow region imaged by both cameras, corresponding to the apex of the KEM.

We aligned our system so that the redundantly imaged region was only 10 pixels wide, meaning that 99.5% of pixels sample an independent region of space when each camera exposes half of its available pixel region (Figure 2.17). We utilized an off-the-shelf KEM that exhibited roughness at the very edge of the mirror surface due to manufacturing limitations. This roughness scatters incident light and leaves a subtle stripe in the stitched image of a densely fluorescent sample, barely visible in the grayscale image of a larval zebrafish brain slice (Figure 2.18).

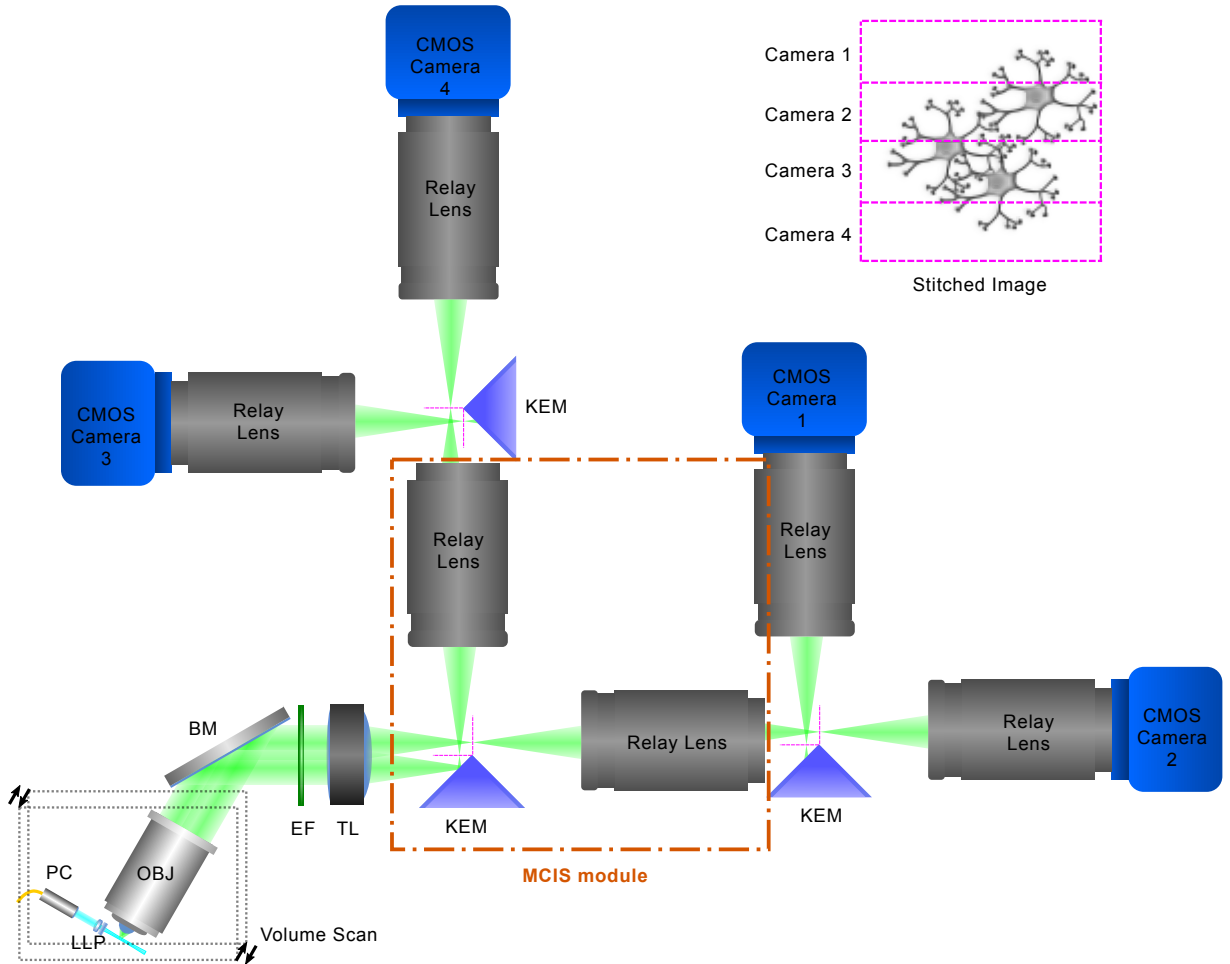


**Figure 2.17:** Quantification of a stitched image of fluorescense dye solution with the same width and location as shown in Figure 2.16. The width of the redundant image region is approximately 10 pixels (less than 0.3% of the camera chip width).



**Figure 2.18:** Pseudocolored and grayscale views of the same stitched slice of a larval (5 dpf, HuC:GCaMP6s) zebrafish brain. Scale bar: 20  $\mu$ m.

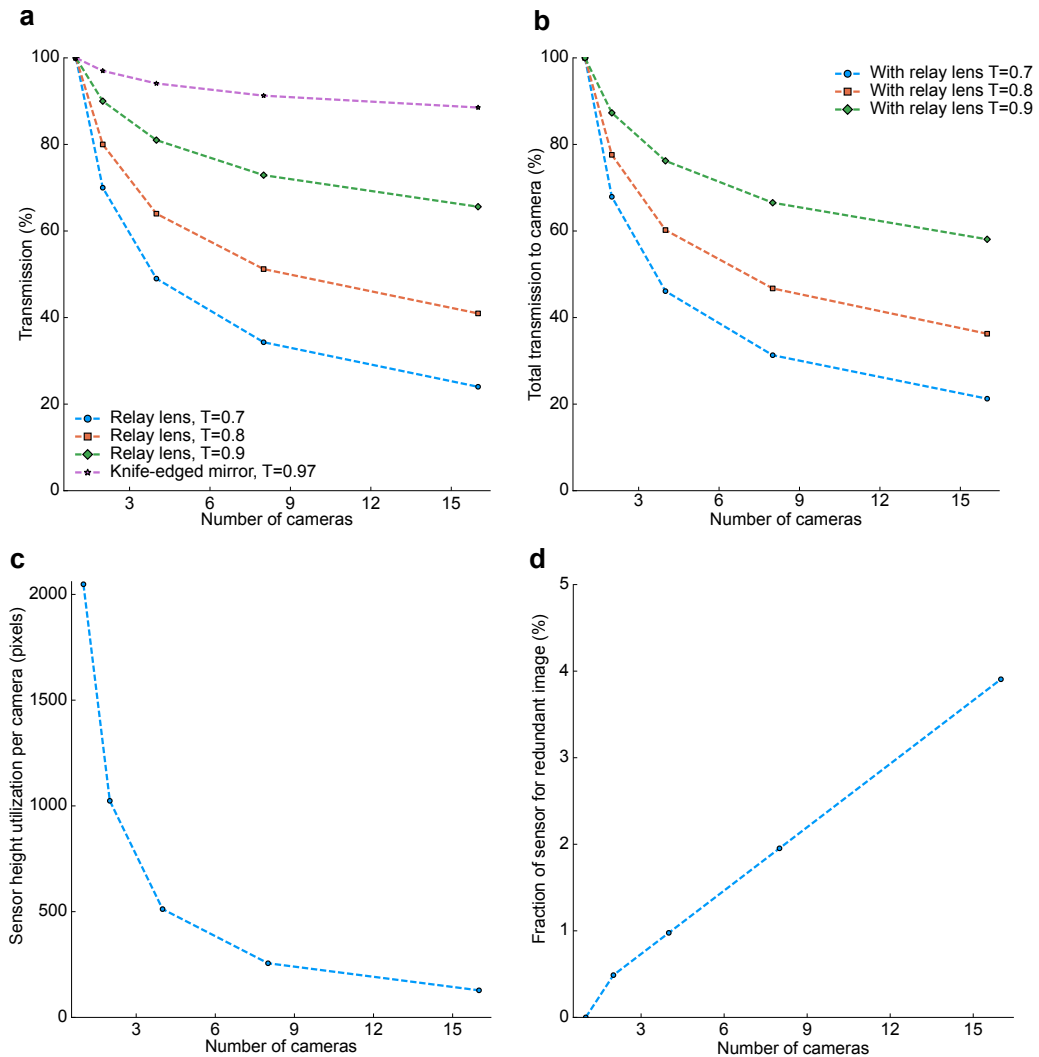
### 2.3.1 Scalability of MCIS



**Figure 2.19:** Multiple MCIS modules can be chained together to split the image to more than two cameras. Shown is a diagram of an OCPI system with MCIS scaled to four cameras. The fundamental repeating unit of the design is outlined in orange. Abbreviations are the same as in Figures 2.1 and 2.15.

Multiple MCIS modules can also be chained to further divide the image and relay the partial images to additional cameras. A 4-camera MCIS system exhibiting nearly  $4\times$  imaging speedup is diagrammed in Figure 2.19. Chaining yields a linear increase in frame rate with each additional camera.


The increase in frame rate comes at the cost of a modest linear increase in image redundancy and a logarithmic decrease in photon efficiency with each additional camera, with the efficiency most sensitive to the transmission efficiency of the relay lens system (Figure 2.20). Based on the analysis shown in Figure 2.20 it would be possible to build a 16-camera system (with nearly 16x speedup) with a total transmission efficiency of 60% if 90% efficient relay lenses are used.



**Figure 2.20:** (a) Each MCIS module compounds losses to imperfections in the transmission and reflectivity of the relay lenses and the KEM, respectively. These losses are plotted separately as a function of the number of cameras in the system for relay lens transmission efficiencies ranging from 70% to 90%. (b) Same as (a) but KEM losses have been combined with transmission losses for each lens plotted. (c) As more cameras are added to the system less of the vertical extent of each sensor is utilized to acquire a full stitched image. As shown in Figure 2.4, the maximum framerate is inversely proportional the utilized vertical extent of the sensor. (d) The fraction of each image that contains redundant image information increases linearly with the number of cameras in the system. Thus the effective size of the field of view is reduced by a small amount when compared with a single-camera system.

## 2.4 Identification and removal of heartbeat artifact in calcium imaging timeseries

In order to demonstrate the impact of overcoming both the scanning and framerate bottlenecks, we used the new system to record neural network dynamics in a zebrafish brain expressing GCaMP6f[17]. We chose to image at  $10\times$  magnification in order to maximize camera framerate. Note that  $10\times$  magnification is not a fixed design parameter; magnification (and field of view) can easily be changed by swapping in a different objective lens. We imaged a volume encompassing 40 slice planes in the forebrain ( $223\ \mu\text{m} \times 127\ \mu\text{m} \times 200\ \mu\text{m}$ ) of a zebrafish larva with a 20 Hz volume rate and  $0.65\ \mu\text{m} \times 0.65\ \mu\text{m} \times 5\ \mu\text{m}$  voxel size over a 20-minute period. Separately we were able to image the whole brain of the fish at a rate of 10 Hz, but the following analysis focuses on the 20 Hz recording of the highly-active forebrain.




**Figure 2.21:** Maximum intensity projection of voxelwise  $F$  (grayscale) and  $\Delta F/F$  (magenta) along the dorsal-ventral axis of the larval zebrafish forebrain (5 dpf) with pan-neuronal GCaMP6f expression (HuC:GCaMP6f) acquired at 20 stacks/s at  $10\times$  magnification with OCPI and MCIS. Only voxels with greater than 15%  $\Delta F/F$  are colored.

Figure 2.21 shows a maximum intensity projection of the voxelwise change in fluorescence relative to baseline ( $\Delta F/F$ ) signal at a single timepoint (see methods). As expected with LSFM, motion of cilia and blood cells in the fish’s body induces time-varying stripe artifacts along the light sheet propagation axis in some image regions. We avoided these regions when we segmented regions of interest (ROIs) containing a subset of 629 neurons throughout the volume and extracted  $\Delta F/F$  timeseries (see methods).

A raster plot of cellular activity during the first 4 minutes of the recording is shown in Figure 2.22, along with a commonplace analysis of neuronal timeseries: a matrix of pairwise correlations computed over the duration of the recording. Before computing correlations, neuron traces were highpass filtered with a cutoff of 1.0 Hz in order to emphasize correspondences over short timescales.

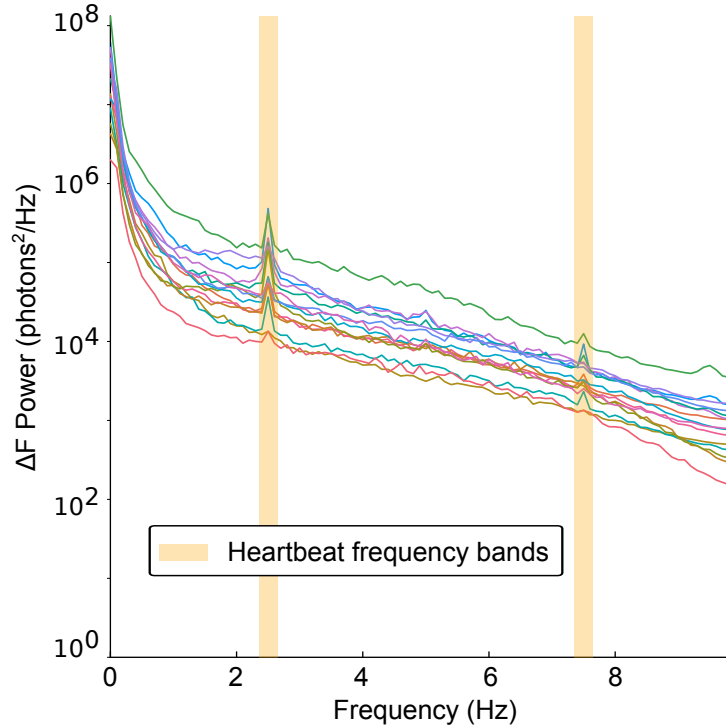




**Figure 2.22:** (a) Raster plot of  $\Delta F/F$  within 629 manually segmented neuron ROIs over a 4 minute period. Shown is an excerpt from a 20 minute recording. ROIs were drawn smaller than the size of each cell in an attempt to minimize the effects of motion artifacts and cross-talk between nearby neurons. (b) Neurons exhibit a range of pairwise correlations in the  $\Delta F$  signal. Correlations were computed with highpass-filtered neuron traces (1.0 Hz cutoff) in order to focus on relationships revealed by high sampling rate. Neurons are ordered by axial depth in the forebrain (dorsal to ventral).

Genetically encoded calcium indicators vary in their decay times. The decay time of GCaMP6f is about 400 ms in dissociated rat hippocampal neurons[17] but this varies by neuron, and the decay time seems to be considerably more brief at temperatures closer to the larval zebrafish body temperature (71 ms at 25 °C). In combination with the fact that indicator rise times are shorter (about 50 ms for GCaMP6f[17]) this suggests that sampling at 20 Hz and beyond will provide useful information. We evaluated this empirically by computing the power spectral density (PSD) of the  $\Delta F$  signal in each neuron. PSDs for 15 neurons are plotted in Figure 2.23.

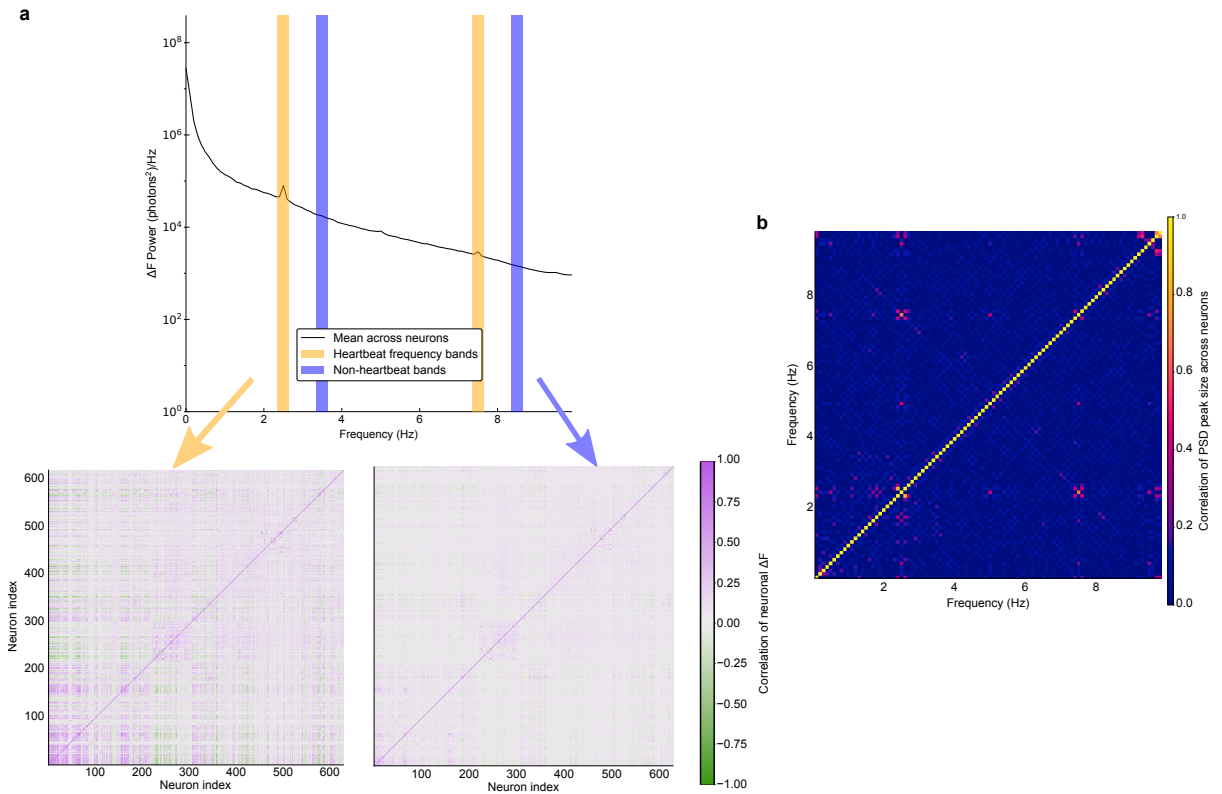
Indeed we found that power diminishes only gradually up to the maximum frequency (10 Hz) permitted by the Nyquist sampling theorem. Moreover we noted peaks at 2.5 Hz and 7.5 Hz in the spectra of many individual neurons. These peaks underlie correlations much larger in magnitude than neighboring frequency bands of the signal, and the size of the 2.5 Hz peak is especially correlated with the size of the 7.5 Hz peak relative to other frequency pairs (Figure 2.24). These attributes suggest that the peaks are two frequency components of the same signal, and that this signal is mixed into the signals of many neurons. The frequencies of the peaks correspond with the primary frequency components found in a fluorescence-based motion-tracking study of the fish's beating heart[51].




**Figure 2.23:** Power spectra of the  $\Delta F$  signals for a subset of neurons show that power diminishes gradually with increasing frequency. These and many other neurons exhibit peaks in their spectra at 2.5 Hz and 7.5 Hz that correspond with the larval zebrafish heart rhythm[51].

Such heartbeat artifacts could bias the results of analyses of neuronal dynamics, just as motion artifacts have been found to bias the results of human brain imaging studies[70]. In order to determine the effect of this artifact on the pairwise correlation measure, we first removed the artifact by applying bandstop filters to remove the heartbeat frequency bands from each neuronal signal. We then computed new pairwise correlations, and for each pair of neurons we recorded the change in correlation magnitude pre vs post artifact removal (Figure 2.25). This analysis suggests that correlation values were overestimated or underestimated by as much as 0.1, highlighting the danger that the beating heart could induce spurious correlations into neuronal network analyses. We conclude that many future zebrafish neuroscience studies would benefit from adopting sampling rates of at least 15 Hz (the Nyquist frequency of the

high frequency heartbeat component) so that the artifact can be filtered out and spurious correlations avoided.



**Figure 2.24:** (a) Comparison of time-domain correlations in the isolated heartbeat frequency bands (lower left) and another nearby set of frequency bands (lower right). The former correlation matrix contains a much higher proportion of strong positive and negative correlations, as would be expected with a broadly distributed heartbeat-induced motion signal. (b) The sizes of spectral deviations at the two frequencies of the putative heartbeat signal are highly correlated across neurons, suggesting that the two frequencies are components of a single signal. For each of the 629 segmented neurons we measured the size of the peak in the PSD at each frequency. Peak amplitude was measured by dividing the PSD amplitude at each frequency by the mean amplitude of the two surrounding frequency bins (thus peak size is not defined for the maximum and minimum frequency bins). Plotted is the correlation between peak sizes across frequencies. Correlations near the diagonal are expected to be strong because they are between similar frequencies. However the strong correlation between the 7.5 Hz and 2.5 Hz peak size stands out.



**Figure 2.25:** Bandstop filters were applied to remove the heartbeat frequency bands (see methods) before recomputing correlations. Shown is the matrix of differences in correlation values obtained before and after heartbeat artifact removal ( $corr_{after} - corr_{before}$ ). Thus this matrix highlights spurious correlations due to heartbeat that could contaminate a naive analysis of neuronal activity.

## 2.5 Discussion

Our results demonstrate that these improvements in hardware and software design allow OCPI microscopy to achieve state-of-the-art fluorescence imaging rates over large volumes. These modifications bring OCPI in line with several other techniques that are limited primarily by camera speed[26, 72, 25, 63, 28, 82, 100, 45, 11, 20, 26, 72, 63, 28, 77]. Considering all tradeoffs, we anticipate that fast OCPI will be a technique-of-choice for many experiments requiring maximal imaging speed of large volumes. The computational complexity of other techniques based on image deconvolution such as light field microscopy can be prohibitive, especially for lengthy imaging sessions. Relative to light field microscopy, direct imaging methods such as OCPI also exhibit a more favorable tradeoff between resolution and imaging rate, and they permit real-time analysis of imaging data[85, 80].

On the other hand when compared to other fast direct-imaging LSFM variants, OCPI can often achieve better image quality because it avoids aberrations caused by imaging outside of the native focal plane of the objective[26, 28]. While “oblique” imaging methods (those that create the excitation sheet with the imaging objective[11, 25, 82, 100]) avoid these aberrations, they are not photon-efficient: only a small fraction of the light cone collected by the objective reaches the image sensor. The highest efficiency demonstrated so far in an oblique configuration is 21%[82], while OCPI microscopy utilizes all of the light collected by the objective. This photon efficiency minimizes photobleaching and phototoxicity that can damage tissue or alter biological processes of interest[46]. Moreover oblique imaging systems require a relatively high-NA objective in order to collect any light at all, nullifying the benefits of low-NA imaging objectives for light sheet microscopy: 1) larger depth of field, permitting a thicker light-sheet with a larger Rayleigh range and thus larger field of view, and 2)

reduced aberrations due to refractive index mismatches between tissue and media. SPIM, perhaps the closest relative to OCPI microscopy, shares these optical advantages and differs only in that the sample is translated rather than the optics. However this makes a crucial difference at high scanning rates: we avoid jostling the sample, and we can optimize the response of the scan system without regard for the mass of the sample being imaged.

OCPI and SPIM also have the unique advantage that the size of the imaged volume is limited only by the range of the linear actuator used, whereas all other techniques mentioned are limited to a volume set by the field of view of the objective. This advantage is currently underutilized because the opacity of samples limits imaging depth. Integrating a two-photon light sheet[90] into an OCPI microscope would alleviate this problem to some extent, but in practice 2P LSFM does not offer the same depth advantage as point-scanning 2P. This is because with 2P LSFM, as with 1P LSFM, scattered emission cannot be attributed to a precise location in the sample [88]. A greater scan range can also be utilized by incorporating structured illumination[42, 49] or multi-view imaging[44] into future OCPI microscopes, but these solutions decrease the rate of image acquisition. Alternatively, instead of modifying the microscope one could modify the tissue to match the refractive index of the media (i.e. reduce scattering), but as of now this is only possible in fixed tissue[89, 74, 99]. Another promising direction for future work would be to alter the OCPI scan direction to be parallel to the sample surface, avoiding scanning deep in the tissue.

Our scan system was tuned manually in an attempt to simultaneously satisfy multiple scan rates, ranges, and amplitudes, and we conservatively utilized only a fraction of its 100N maximum force output (Figure 2.12). In the future an automated procedure (either empirical or simulated) will allow one to choose optimal parameters for



each recording session and to more easily swap in an objective lens with a different mass. An automated procedure should optimize the control system so that the piezo response is closer to an ideal triangle wave. A triangle wave is optimal because image slices can be distributed uniformly in both space and time. When images are not spaced uniformly in time, as diagrammed in Figure 2.9, the camera spends a fraction of each stack idle and thus its maximum framerate is not fully utilized. The achievable framerate of the camera depends on the shortest interval between frames, 0.8 ms in Figure 2.9. When compared to the 1.25 ms interval that would be possible with a true triangle wave it is apparent that the camera spent 37% of the time idle. Future optimizations will prioritize reducing this idle time.

MCIS addresses the bottleneck resulting from limits in camera framerate, providing an increase in framerate proportional to the number of cameras used. It will be straightforward to integrate MCIS into other LSFM variants because only the components downstream of the tube lens need to be modified. Furthermore other techniques that do not image a plane to a plane, such as light field microscopy (discussed in Section 1.3.2), can still benefit from MCIS by combining multiple camera sensors into one large virtual sensor.

These improvements establish a solid foundation for studying fast dynamic processes—such as signal transmission between neurons—at scale. We also demonstrated that a high sampling rate allows one to remove physiological artifacts such as heartbeat that could bias fluorescence timeseries analyses. We expect that heartbeat artifact removal will become a standard preprocessing step when analyzing zebrafish imaging timeseries. In combination with advances in fluorescent indicators and in large-scale image analysis, these improvements to microscope hardware bring us closer to a more comprehensive understanding of brain-wide activity.

## 2.6 Methods

### 2.6.1 Calculation of image quality in non-native focal planes

See Appendix A.

### 2.6.2 Off-the-shelf components

A laser system (Spectral LMM 5) output a collimated gaussian beam (1 mm diameter) via a pigtailed fiber-optic collimator. This beam was passed first through an achromatic doublet lens (Edmund Optics 45-262) and then a cylindrical lens (see next section), and finally a coverslip before reaching the sample. The lenses were sealed inside a housing so that water could not enter into the lens space when submersed in the sample dish. The alignment of the light sheet was adjusted with a set of small stages (Elliot Scientific MDE266 and MDE269). Either the Olympus UMPLFLN10X/W or Olympus LUMPLFLN40X/W objective collected emission from the sample. The light sheet and objective were mounted  $60^\circ$  from the horizontal axis and scanned together (Piezosystem Jena NanoSX800 piezoelectric positioner, 30DV300 amplifier). A stationary broadband mirror reflected the output from the objective horizontally to a 200 mm tube lens (Thorlabs ITL200) placed at the  $1f$  distance from the objective's back focal plane. The KEM (Thorlabs MRAK25-G01) was placed in the image plane behind the tube lens, and sometimes swapped with a 50/50 beamsplitter (Thorlabs BSW10R) for alignment purposes (see below). Two telecentric relay lenses (Edmund Optics 62-902) relayed the divided image to the cameras (PCO Edge 4.2). Analog and digital I/O to the positioner, cameras, and laser was managed by a PCI data acquisition device (National Instruments PCI-6259) with a single sample clock shared

across all signals. A PC with two RAID arrays (10 hard drives each) streamed the output of the cameras to disk (up to 1GB/s per camera at maximum framerate). The sample was positioned on a physiology breadboard (Thorlabs PHYS24BB) mounted to a lab jack (Newport 281) and XY stage (Scientifica).

### **2.6.3 Custom components**

The only custom optic was the small cylindrical lens used to form the light sheet. This lens was ordered from Tower Optical with a specified focal length of  $-6.25$  mm and diameter of 3 mm. Full lens specifications are included in the appendix. Custom mechanical components were designed collaboratively and refined iteratively in collaboration with the Washington University Medical School Machine Shop. Hardware for the scan system was designed to minimize weight. A parts list, schematics, and labeled photos are also included in the appendix. We found that the small dovetail stages used to align the lightsheet exhibited a few microns of motion in their joints when scanning at high rates. This motion defocused the image and required correction by modifying the dovetail slides to add a locking screw. A magnetic swappable Thorlabs filter cube insert (Thorlabs DFMT1) was modified to hold the knife-edged prism mirror. The 1" apertures of the filter cube itself (Thorlabs DFM) were widened to 1.1" with a standard boring tool to prevent vignetting of the relayed image.

### **2.6.4 Calibration of piezo closed-loop controller**

Initial calibration of the piezo control system (NanoSX800 with 30DV300 amplifier) was performed by the manufacturer (Piezosystem Jena). We requested that they optimize the control system for the highest achievable frequency and amplitude of

operation with a triangle wave command, 400 g load, and a translation angle of 30 degrees from vertical. They tuned PID parameters as follows:  $k_p = -0.3$ ,  $k_i = 50$ ,  $k_d = 0.1$ . The load used in the experiments detailed in this article was smaller (264 g), so we further refined the calibration using the iterative procedure described in the product manual. One parameter at a time was manually adjusted by serial command, and the response of the system was measured. If a parameter update drove the system into oscillations, then the system was immediately switched to open-loop mode and the parameters reset. A detailed PID tuning procedure is available in the product manual. The final PID parameters after this secondary tuning were:  $k_p = -0.37$ ,  $k_i = 50$ ,  $k_d = 0.11$ .

### **2.6.5 Generation of smoothed triangle wave commands**

A triangle wave with the desired frequency, amplitude, and offset was lowpass filtered with a cutoff of 3.25x the triangle wave frequency (32.5 Hz for a 10 Hz triangle wave). This resulted in erosion of the triangle peaks and a reduction in the range of the command. In order to compensate for this reduction the original triangle wave was expanded and filtered again iteratively until the filtered wave matched the desired range.

### **2.6.6 Scan range tuning**

The piezo command waveform was adjusted iteratively until the maximum and minimum values of the piezo cycle (as measured by its built-in sensor) matched those requested by the user. The initial guess for the command range was set to 10% smaller than the target range to guard against potential damage from overshoot.

The piezo was then operated with this repeated waveform for a 20 s warmup period before measuring the sensor response for a cycle. The lower and upper limits were then updated independently based on sensor feedback with the same procedure: error was calculated as  $(target - measured)$ , a value equal to 90% of this error was added to the limit used to generate the command signal, and a new command signal was generated. This was continued until both the upper and lower limits matched the user’s request within a 0.1  $\mu\text{m}$  margin of error. These stopping criteria were met within 5 iterations or less.

### 2.6.7 Pulsed illumination

Since the cameras operate with a rolling shutter, only the latter part of the exposure interval corresponded to simultaneous (“global”) exposure of all CMOS sensor lines. Laser pulses were timed to occur only within this global interval in order to prevent image information from bleeding into adjacent slices of the stack. When the camera is operated at maximum framerate the duration of the global shutter period is only one line time (9.76  $\mu\text{s}$  for PCO.Edge 4.2). At sub-maximal framerates the global period is equal to difference between the chosen exposure time and the shortest possible exposure time that the camera can sustain. Therefore one can prevent bleeding of image information into adjacent slices of a stack during dynamic recordings by using brief illumination pulses aligned with the end of each frame and operating the camera slightly below its maximum framerate. We pulsed the illumination laser only during the last 5% of the exposure interval, which required that the camera operate 5% slower than its maximal framerate. During the imaging session shown in Figure 4 the exposure time was 580  $\mu\text{s}$  and the excitation pulse duration was 30  $\mu\text{s}$ . Peak laser power was set to 18 mW for an average sustained laser power of 0.864 mW. Detailed

diagrams of the camera’s timing system can be found in the PCO.Edge 4.2 camera manual online.

### 2.6.8 Image timing calibration

After sensor-based timing of images proved inadequate we adopted an approach utilizing the camera to perform further calibration of the timing of each image of the stack separately so that the images corresponded to the intended planes of the sample. First a ground-truth image stack of  $0.2\ \mu\text{m}$  fluorescent beads was acquired at a very slow ( $0.1\ \text{Hz}$ ) scan rate so that factors such as dynamic forces and lag in the sensor circuit did not affect the appearance of images. Then, a fast dynamic recording was performed in which each image in both the “forward” and “reverse” stacks was acquired with various timing offsets relative to the sensor-based timing. The search space of timing offsets ranged from  $0\ \mu\text{s}$  to  $1.2\ \text{ms}$  at  $50\ \mu\text{s}$  intervals. Each dynamic slice image was then compared with the corresponding ground-truth “static” image and scored by similarity. The timing offset that produced the highest similarity score for a slice was chosen as the corrected timing for that slice. We noticed that fast dynamic operation produces not only an axial shift in each slice but also a lateral shift of less than  $2\ \mu\text{m}$  that varied by slice (likely due to compression and flexion of components). Therefore in order to calculate the similarity score we first performed 2D rigid image registration to shift the trial image laterally into alignment with the ground-truth image. After alignment the similarity score was calculated as the sum of squared differences between each pixel in the trial image and the ground-truth image normalized by the sum of squared intensity of pixels in the overlap region between images. Both the optimal temporal offset and the optimal lateral shift were recorded for each slice in the forward and reverse stack and used to acquire and align images of

the zebrafish specimen. When using MCIS, only the camera receiving the unreflected image was used for this alignment procedure.

### **2.6.9 Interpolation of bidirectional images**

Bidirectional image acquisition produces image slices that are not uniformly spaced in time. A uniform sampling rate was simulated by interpolating new slices from each forward and back slice pair in the timeseries. Each interpolated timepoint was midway between the sample times, so each interpolated pixel intensity was simply the mean of the corresponding intensities of the two slices. This method also emulates simultaneous sampling of each image slice in the stack.

### **2.6.10 KEM alignment with focal plane for MCIS**

The camera receiving the unreflected image was aligned first using a procedure common to any OCPI microscope. The knife-edged mirror was installed in a modified magnetic filter cube insert so that it could be swapped easily with a 50/50 dichroic mirror (see image stitching method). The filter cube was incorporated into a cage system to maintain its alignment with the tube lens. This knife-edged mirror insert was placed in the cube and the distance between the tube lens and the cube was set by translating the cube along the cage axis and observing a dense bead sample on the camera receiving the unreflected image. The cube was approximately aligned when the image of the bead sample on this camera was in focus and centered on the aperture of the cube (this also required lateral repositioning of the camera). A finer alignment of the cube was achieved by observing closely on both cameras the strip of the image corresponding to the knife edge of the mirror: the farther the edge of

the mirror is from the relay image plane the larger the region of the image that is redundantly imaged on both cameras. Alignment was complete when the width of the redundant image region was minimized (Figure 2.17).

### **2.6.11 Camera alignment for MCIS**

Since a row of pixels on the camera sensor is only  $6.5\ \mu\text{m}$  wide, slight rotations of the relative image planes of the two cameras can result in misalignment of the knife-edge line with the rows of the camera. Therefore in addition to the focal alignment above we also adjusted the relative rotations of the two cameras until the imaged knife edge line was horizontal. Since the KEM was mounted in a highly-repeatable kinematic insert, it was only necessary to perform this angular adjustment once.

The following procedure was used whenever the vertical size of the camera's active pixel region was changed: With the KEM installed, the active pixel region of both cameras was set as desired (settings on both cameras must match). Both cameras were activated to stream live updates of an image of a dense fluorescent bead sample. The camera receiving the unreflected image was translated up so that the KEM edge in its live image feed corresponded with the bottom edge of the pixel region. The other camera was likewise translated laterally so that the KEM edge aligned with the bottom edge of the pixel region. Since the image is reflected this edge corresponds to the top side of the unreflected image. For this reason all reflected images were flipped in software before performing the alignment steps described in the next section.



### 2.6.12 Stitching MCIS images

After aligning the cameras and setting the desired pixel region, both cameras recorded an image of the beads simultaneously. Then the KEM was swapped with a 50:50 plate beamsplitter (Thorlabs BSW10R) mounted in a kinematic filter cube insert, the pixel region of the camera receiving the reflected image was set to full size (2060 x 2040 pixels), and another image was recorded. It was critical that the bead sample remained stationary during the interim between these image snapshots. The full-size image spans the region where the two smaller images meet, and thus contains the information needed to align and stitch the smaller images. The bead sample was then replaced with the sample of interest, and the full imaging session was completed.

After all data were recorded, an image transformation was found to align the smaller bead images, and this same transformation was applied offline to stitch all images recorded in the zebrafish specimen. This transformation was found with the following 3-step procedure. First, a rigid 2D transform was found to align the smaller reflected image with the full image. This transformation was minimal in magnitude because both images were acquired with the same camera and centered on the same region of space (only the reflective surface was different). Second, a 2D affine transformation was found to align the full image with the smaller reflected image. A full affine transformation was allowed because subtle differences in alignment or manufacturing of the cameras and relay lenses cannot be captured by a rigid transform. The non-rigid component of the transformation was small (less than 1% scaling factor) but important to maintain alignment of beads throughout the image. In the third and final step the two transforms were composed into one transform (rigid first, affine second). This composite transform was applied to each reflected image of the zebrafish specimen to align it with the unreflected image. Before combining the two images the

camera’s constant bias intensity was subtracted from each pixel (by design the black level is not zero but a constant value).

### **2.6.13 Zebrafish imaging**

HuC:GCaMP6f and HuC:GCaMP6s zebrafish larvae[24] were crossed with “Casper” larvae[97] for two generations to obtain transparent fluorescent larvae for imaging. Embryos were raised at 28.5C, screened for green fluorescence at 3dpf, and imaged at 5dpf. The larva was transferred by pipette into a drop of 1.5% low-melting-point agarose gel while the gel was still warm and in a liquid state. A syringe was used to suck the larva tail-first into a segment of Fluorinated Ethylene Propylene (FEP) tubing, and the gel was allowed to solidify. FEP was chosen because its refractive index closely matches that of water, and therefore aberrations are minimized when imaging through the material. The tubing segment was then mounted in a custom water-filled chamber at an angle of 60° from horizontal so that the rostrocaudal axis aligned with the light sheet. Excess tubing in front of and above the fish’s head was cut away with a razor blade. All protocols were approved by the Institutional Animal Care and Use Committee at the Washington University School of Medicine.

### **2.6.14 Stripe removal**

As is common with LSFM, we observed stripes in the zebrafish images due to scattering, absorption, and interference as the light sheet propagated through the sample. We applied a destriping filter to attenuate these stripes as a preprocessing step for the images in Figure 2.18 and Figure 2.21. Each image slice was first log-transformed to account for the multiplicative (rather than additive) nature of the stripes. Next

the slice was Fourier transformed, and the magnitude of bins corresponding with the angle of the stripes was attenuated until the stripes were no longer visible when the image was reconstructed with an inverse Fourier transform. We did not apply the filter when performing timeseries analyses due to concern that the filter could introduce subtle artifacts. Most of the stripes observed were static and therefore could not be expected to influence the  $\Delta F/F$  calculation. Brain regions exhibiting dynamic stripe patterns (for example regions behind the motile cilia of the olfactory rosettes) were not analyzed.

### **2.6.15 Image registration**

The zebrafish, embedded in agarose and FEP tubing, moved little during the recording. However image registration was required to compensate for slow drift of the specimen. A single image stack was chosen as a “fixed” reference stack and each of the other “moving” stacks in the timeseries was registered to this stack. A custom algorithm found the simple shift (translation) that maximized the overlap between the fixed and moving stack.

### **2.6.16 Extraction of neuronal calcium traces**

Neuron ROIs were selected by hand using a graphical software tool. A single 2D rectangular region was marked for each neuron in the plane that best appeared to capture the neuron’s activity. The outer edges of the cells were excluded from ROIs in order to minimize crosstalk in the calcium traces of nearby neurons and to reduce the effect of motion artifacts. For each timepoint of each ROI the raw fluorescence value was calculated as the sum of voxel intensities within the ROI.  $\Delta F/F$  was calculated

by subtracting baseline fluorescence from this raw value and dividing the result by baseline. Baseline for each timepoint was calculated as the moving average of raw ROI fluorescence during the 60 s interval centered on the timepoint (the first and last 30 s of the recording were not analyzed).

### **2.6.17 Power spectral density and filtering of neuronal calcium traces**

All PSDs were computed with the Welch method with a rectangular window function on baseline-subtracted fluorescence traces. When filtering the signals to remove heartbeat frequency bands a Butterworth filter of order 5 was applied for each stop-band. The filters were applied in both the forward and reverse directions in order to preserve phase.

### **2.6.18 Software**

We wrote several software modules to accomplish microscope control, PID adjustment, image timing calibration, image stitching, temporal interpolation, and manual cell segmentation. Web links to these modules on Github are collected here: <https://github.com/HolyLab/FastScanningAndMCIS>.

# Chapter 3

## Analyzing data from a fast OCPI microscope

As discussed in Section 1.4, leveraging the full capabilities of a fast OCPI microscope requires careful management of very large imaging datasets. These datasets are now so large that examining raw images by eye is of limited use, and algorithms are increasingly important for extracting insights from the images (However manual inspection remains critically important for quality control of raw images and the algorithms applied to them). This chapter describes important considerations when generating and processing large imaging datasets with the fast OCPI microscope described in Chapter 2. It begins by describing factors that must be considered *before* any data is acquired, along with ways that we have enabled flexible experimental design and data acquisition (Section 3.1). Section 3.2 outlines a set of general strategies for scientific computing with large-scale imaging data. Sections 3.3 and 3.4 then explore special considerations arising when processing light sheet images and video imaging of live animals, respectively. Where relevant we also propose microscope hardware modifications that will simplify preprocessing in future work. In Section 3.5 we describe a full preprocessing pipeline developed for optical neurophysiology. Finally in Section

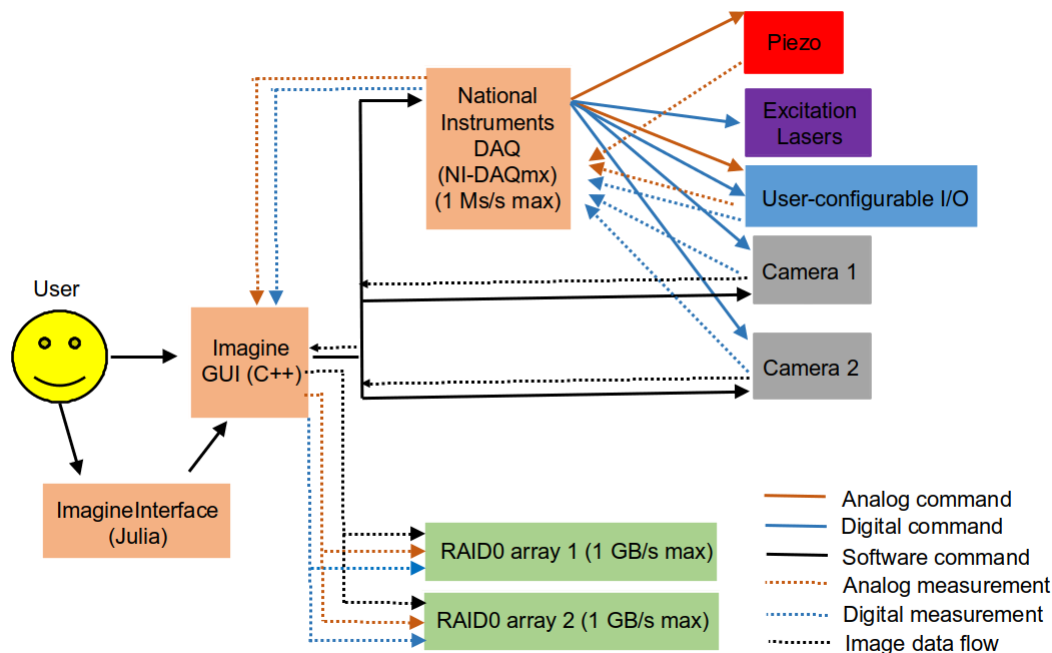
3.6 we discuss promising directions for future work at the intersection of microscope hardware and algorithms.

### 3.1 Before the experiment

The most critical scientific decisions are usually made *before* an experiment is performed, and this is perhaps especially true when using a fast OCPI microscope. This is because extracting a scientific result from a large imaging dataset usually requires several time-consuming processing steps. Currently it's typical for these steps to take much more time than it takes to collect the data. Thus the cost of a poorly designed experiment can be quite high if the design flaw is not discovered until after lengthy processing.

Experimental design is of course highly specialized to the scientific question at hand, and we refrain from discussing specific designs here. Rather, our goal with the fast OCPI microscope is to design a system flexible enough that the experimenter is free to make good design choices. Flexibility is a key design consideration both for microscope hardware (i.e. easy choice of objectives, fluorescence filters) and software. The remainder of this section describes flexible software developed for designing and running experiments.

Figure 3.1 illustrates the flow of imaging data as well as analog/digital signals in the dual-camera fast OCPI microscope. The user interacts with a graphical interface (shown in Figure 3.2) that coordinates several hardware devices: cameras, piezo scanner, excitation lasers, hard drives recording data, and any other devices (analog/digital, input/output) that the user connects with the microscope. This ability



**Figure 3.1:** The microscope user interacts with the Imagine graphical interface (shown in Figure 3.2). Optionally, if more precise control is desired the user can design an experiment with the Julia package `ImagineInterface` and then load that experiment in Imagine. Imagine coordinates several hardware devices: cameras, piezo scanner, excitation lasers, hard drives recording data, and any other devices (analog/digital, input/output) that the user connects with the microscope. During a recording Imagine streams images from the camera(s) to arrays of hard drives (RAID0 configuration for speed), one array per camera.

to interface with user-specific devices affords the microscopist a great deal of flexibility in experimental design. For example, neuroscience experiments often involve delivery of stimuli during imaging as well as concurrent recording of other physiological signals beyond images. The graphical interface, named Imagine, was redesigned in key ways to support this flexibility with the fast OCPI microscope:

1. Upgrading to the modern Qt toolkit (frontend).
2. Switching to hardware-timing and synchronized sampling for all analog and digital input and output channels (backend).

3. Optimizing multithreading to enable streaming of data from multiple cameras operating at maximum frame rate (backend). See also Figure 3.3 for thread diagram.
4. Support for saving/loading all instructions for the experiment to/from a file, optionally generated by a script in Julia (frontend and backend).

Item #4 references a new software package written in Julia to aid in the rapid design of experiments (<https://github.com/HolyLab/ImagineInterface>). The selection of Julia is further discussed in Section 3.2.2. The package gives full access to all of the analog and digital channels of the microscope shown in Figure 3.1 and also facilitates loading of input channels recorded during the experiment. Moreover both `ImagineInterface.jl` and `Imagine` support multiple existing OCPI microscopes differing in hardware configuration, and the software is designed to facilitate extension to new microscopes. Upgrades to existing microscopes, such as the point-scanning feature discussed in 3.6 that requires control of two galvanometers, can be integrated easily.

Currently instructions for an experiment must be saved to a file from Julia and then loaded and run from the `Imagine` GUI. In the near future we plan to make this process entirely scriptable: `Imagine` will run as a “daemon” process for streaming images to disk/screen while a new package, `Imagine.jl`, will control all microscope hardware. The iterative calibrations of the piezo scan command and camera exposure timing described in Chapter 2 made heavy use of `ImagineInterface.jl` and an early version of `Imagine.jl`.



## 3.2 Computing with terabyte-scale imaging data

After imaging data is saved locally to the computer controlling the microscope it is typically moved to another location for processing. In the future it may be advantageous to perform some data processing “online” while recording with the microscope, but this is only feasible for relatively simple processing steps. A large fraction of the computation involved in an experiment is performed afterward with tools chosen by the experimenter. This section describes choices faced by users regarding hardware and software for performing this analysis, focusing on our particular choices in these areas when processing datasets produced by the fast OCPI microscope.

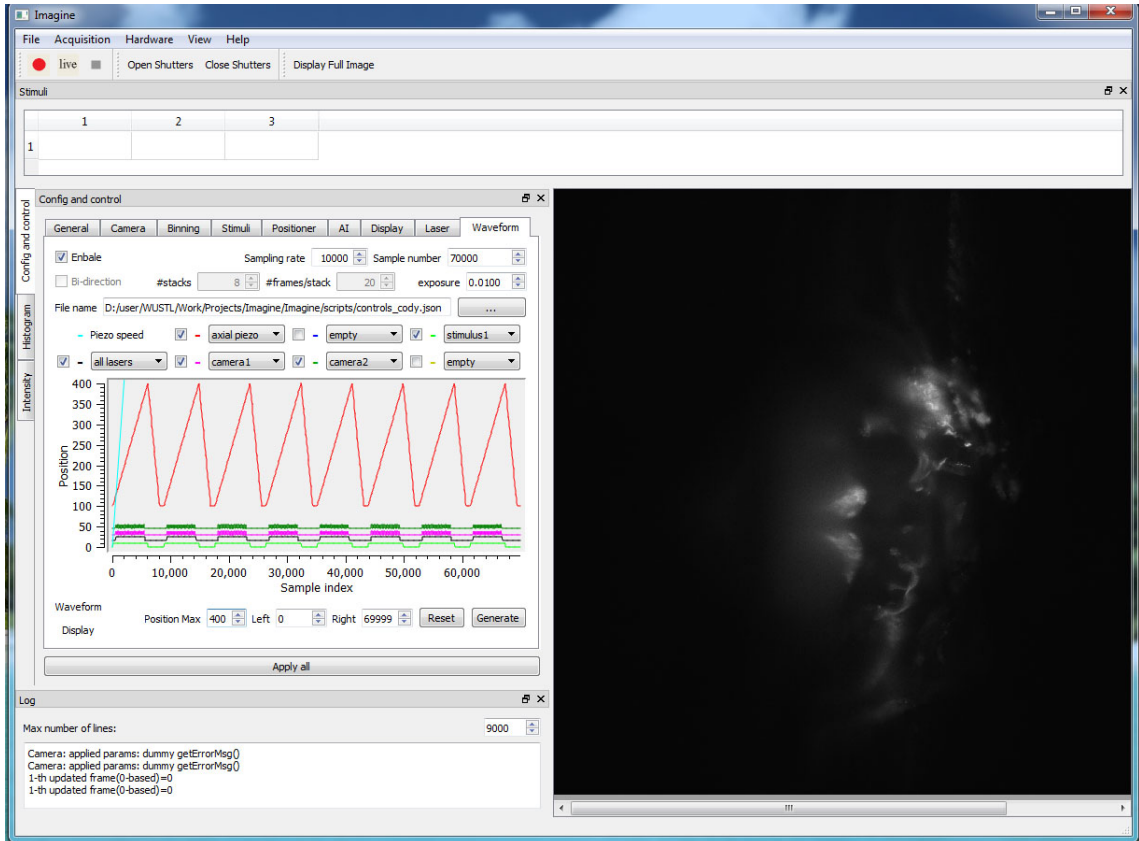
### 3.2.1 Computing hardware

Perhaps the most consequential choice in hardware is between a single server and a distributed system (i.e. a cluster or supercomputer). This is consequential because an algorithm that works well on a server may need to be adapted to work well on a distributed system, or vice versa. Distributed systems have access to greater computing power and memory, but because these resources are spread over many computing nodes it is relatively slow to move data from one node to another. Another major difference is that computing clusters require a higher level of expertise to maintain and operate, often necessitating involvement of a third party computing service. At least one research group has already committed to developing distributed computing solutions for large LSFM datasets[29]. Our group, on the other hand, has focused on developing a software pipeline that can be run on a single server. Most of the algorithmic considerations discussed in the following sections apply to both single-server and distributed environments.

Another important decision is what computing hardware to include in the single-server. A fast central processing unit (CPU) is highly useful. However CPU speeds have essentially stopped increasing due to manufacturing limitations, and as a result computing systems are increasing in *parallelism* rather than speed. Therefore it will be fruitful to choose CPUs with high core counts and to design software accordingly. In addition recent years have seen the rise of powerful *coprocessors* that can perform specialized computing tasks highly efficiently. The graphics processing unit (GPU), for example, was developed for rendering images but has proven extremely useful for certain general-purpose processing tasks. Likewise the more recently developed tensor processing unit (TPU) greatly accelerates many linear algebra operations[78]. Both of these coprocessors excel at parallelism, integrating 1000s of simple processing units rather than the 10s of complex cores available in modern CPUs.

Regardless of the chosen processing unit, computation on large imaging datasets is often *I/O bound*, meaning that the rate of loading/writing the data from/to storage is the limiting factor in the runtime of an algorithm. Therefore it is beneficial to use a fast data storage device. Unfortunately the speed of a storage device tends to be inversely proportional to the capacity of the device; choices may be limited for very large imaging datasets. However this is an area of rapid development; recently multi-terabyte solid state storage devices with speeds of over 1 GB/s have entered the market. Such devices, while too expensive for permanent storage, may work well as temporary storage for mitigating I/O bottlenecks during processing. Note that even the fastest storage devices are much slower than RAM, both in terms of access time (latency) and throughput. Modern RAM has an access time of around 50 ns, compared to about 20  $\mu$ s for the current fastest storage devices. Likewise RAM can stream data at 60,000 GB/s, compared with 3 GB/s for the fastest sustained reads

from storage. RAM will remain much faster than storage for the foreseeable future, and algorithms must be written accordingly.

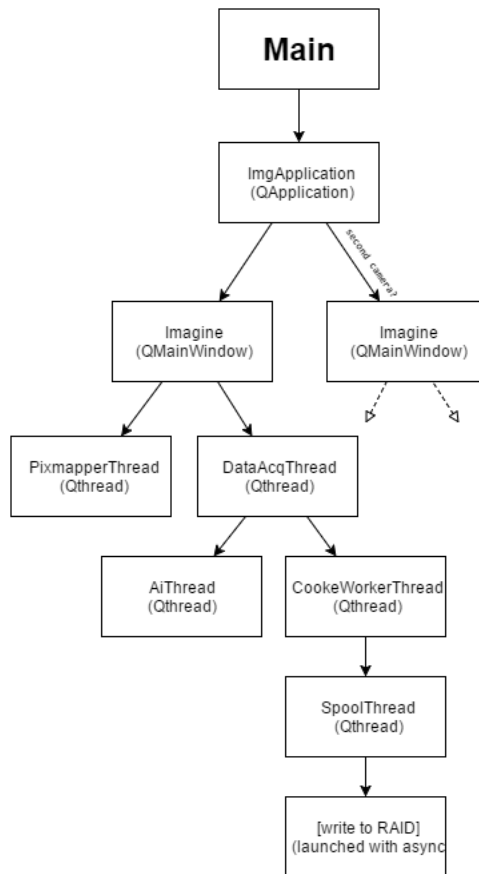


**Figure 3.2:** Screenshot of the graphical user interface (GUI), “Imagine”, for controlling the fast OCPI microscope. The GUI was originally written in C++ in 2008 and utilizes the Qt windowing toolkit for graphical elements. Several upgrades were made in recent years. The tabs shown left of the image window control aspects of the recording. The selected “Waveform” tab allows loading and visualizing a full set of instructions for running an experiment from a file. A support package in the Julia language, “ImagineInterface.jl”, provides a scripting interface to design and export these instructions to Imagine.

ImagineInterface.jl credit: Cody Greer

Original Imagine credit: Zhongsheng Guo, Timothy E. Holy

Upgraded Imagine credit: Ben Ackland, Cody Greer, Dae Woo Kim




**Figure 3.3:** Diagrammed is the thread hierarchy employed by the Imagine software. One Imagine graphical window (third level in hierarchy) is created for each camera connected to the microscope. During data acquisition the PixmapperThread (updates images on screen) and the DataAcqThread run simultaneously. The DataAcqThread delegates the work of streaming data from the camera to the CookeWorkerThread, which then communicates with other threads to write data to the hard drives (RAID array). Thread communication and prioritization are orchestrated so that streaming image data takes priority over updates to the graphical interface during fast recordings.

### 3.2.2 Software design strategies

The stagnation of single-CPU performance and the rise of parallelism has two main implications when selecting computing infrastructure that can scale with increasing rates of image acquisition. First, parallelism should be the preferred way to solve difficult problems. Second, when it is difficult or impossible to develop a parallel solution, single-threaded CPU operations should be performed as efficiently as possible. Achieving peak efficiency often requires additional time and expertise from the programmer, a difficult tradeoff to balance in a highly experimental environment like academic science. The Julia language, described below, seeks to make a favorable tradeoff in this regard.

**The Julia language** Any algorithm can be implemented in any Turing complete programming language, but a careful choice of programming language can make programming tasks *easier*. The Julia language, launched in 2012 but only recently reaching version 1.0, was designed to make programming easier for scientists, mathematicians, and engineers[6]. We have chosen Julia for our most computationally intensive tasks due to its highly favorable tradeoff between program runtime and the time it takes to write code. One proxy for programmer investment time is simply the length of the text that specifies a program. When this quantity is plotted against runtime for several benchmark algorithms Julia fares very well compared with other popular languages, demonstrating performance approaching that of C with many fewer lines of code (Figure 3.4).

Julia achieves high performance through multiple strategies, but perhaps the most important one is that Julia blurs the line between *compiled* and *interpreted* languages. Many languages that are popular for exploratory data analysis (Python, Matlab, R)



**Figure 3.4:** Shown is a comparison of performance (runtime, smaller is better) and conciseness for a diverse set of benchmark algorithms written in Julia and other popular languages. Note that Julia (blue dots) exhibits the favorable combination of short runtimes and concise code. This combination is highly desirable in scientific applications where code is rapidly prototyped, executed, and rewritten. (Reprinted from london-julia.org, original source was benchmarksgame.alioth.debian.org but is now offline.)

are interpreted languages, meaning that the programmer’s instructions are compiled to machine code line-by-line as the program is executed. This allows a great deal of flexibility because the types of variables need not be known until runtime, but the tradeoff is that these languages spend a lot of time compiling code. This contrasts with efficient compiled languages like C in which the entire program is translated to machine code before execution begins. This means that the types of all variables must be known at compile time, making it difficult to write code that is flexible and reusable. It can be said that Julia supports both compiled and interpreted modes of operation. If the types involved in the execution of a function can be inferred by the Julia compiler (LLVM) at compile time then Julia compiles a specialized version

of the function automatically (and only once), allowing the function to run with the efficiency of a compiled language. If, on the other hand, types cannot be inferred, then Julia behaves as an interpreted language with slower performance. A great number of crucial language design decisions (not discussed here) have centered around making it easy to write Julia code that compiles well while preserving a simple syntax familiar to Matlab and Python users.

Another important factor in choosing a language is simply its popularity. A language's popularity is strongly correlated with the availability and maturity of software libraries. Moreover collaborations are much easier when scientists share a language. In this regard Julia has not yet reached the maturity of Python, but this is changing quickly. Academics and major corporations alike are increasingly adopting Julia.

Regardless of language choice there are common software design considerations when processing large imaging datasets. The following paragraphs discuss some language-agnostic design strategies that we have found useful.

**Optimize read and write operations** As mentioned in 3.2.1, the most severe bottleneck when processing large imaging datasets is the time spent communicating with the storage device. Small datasets can be loaded entirely into system memory (RAM) to exploit its faster read/write times, but this is currently impractical for multi-terabyte datasets with a single server. When a single conventional hard disk is used, the time required to sequentially read each pixel from a dataset into memory is around  $10\times$  longer than it took to record the data with the microscope (for the microscope described in Chapter 2). The speed difference between memory and storage is especially extreme when a random access pattern is needed rather than a sequential pattern.



This suggests three design criteria for algorithms:

- Read the dataset from storage as few times as possible.
- Use a sequential access pattern whenever reading from storage.
- Avoid writing intermediate (partially-processed) copies of the dataset whenever possible.

We defer discussion of specific algorithms until Section 3.5, but one general strategy that we have found useful is that of *lazy evaluation*. The “lazy” programming paradigm delays operating on data until the point in the program in which the result of the operation is required. The main advantage of this approach is that a sequence of image processing steps can be composed without re-reading data from storage at each step. An example of the lazy approach when applying image transformations is shown in Figure 3.5. The syntax is nothing special, but each of the functions shown is implemented such that the operation is delayed until the last line when data is accessed. Moreover *only* the computations needed to generate the requested array element (in this case a pixel) are performed. With this approach raw data is read from disk only once and the results of intermediate steps need not be saved to storage. If one wanted to inspect results at an intermediate step (as often happens) then one can do so with only a minor modification of the code.

```
img = stitch(img1, img2, tform)
img = warp_timeseries(img, tforms)
img = dfof(img)
pixval = img[1] # Above computations are delayed until now
```

**Figure 3.5:** Shown is Julia code demonstrating “lazy” evaluation of functions for stitching images, correcting for movement, and calculating fluorescence change from baseline (operations needed for the results shown in Chapter 2). The syntax gives no indication of lazy evaluation; rather, each function shown is written so that the action it performs is delayed until pixel data is accessed (in the last line). *Only* the data needed for the requested pixel is accessed and processed. In Julia lazy evaluation has no overhead relative to conventional evaluation.

**Patchwise algorithms** As already mentioned, advances in computing capability are now predominantly driven by increases in parallelism. Therefore in order to continue to scale image processing to ever-larger datasets it will be necessary to design algorithms that harness this parallelism. Fortunately image processing tasks are relatively well-suited to parallelization. This is because images often contain a great deal of *local* information that can be treated independently of information in distant parts of the image and thus processed by separate hardware devices in parallel. Coprocessors like the GPU were designed with this parallelism in mind.

In many cases the processing of a local “patch” of an image is only partially independent from other patches. Any non-independent computations require communication between parallel processing streams, potentially slowing the processing significantly. A crucial design goal is therefore to design algorithms that operate on image patches and require minimal communication between threads. Our approaches to image registration and neuron segmentation (discussed further in Section 3.5) adopted this design philosophy.

## 3.3 Special considerations for light sheet microscopy data

The previous section discussed some computational strategies that apply generally to large imaging datasets. This section focuses on computational considerations that are specific to fluorescence imaging datasets, especially those produced by LSFM.

### 3.3.1 Point spread function

The point spread function (PSF) is a key property of any imaging system: it specifies the distribution of light intensity for an image of a single infinitely-small point in the object space. Note that when the PSF is determined experimentally, a bead is taken as an approximation of “infinitely-small”. This is a fair approximation if the bead is smaller than the diffraction limit set by the NA of the imaging system. Fluorescence-based imaging systems have two PSFs of importance: the excitation PSF and the emission PSF. These PSFs are necessarily different because excitation and emission light differ in wavelength and thus have different diffraction limits (defined in Section 1.2.3). The effective PSF of a fluorescence microscope can be described in terms of the combined effect of the excitation and emission PSFs on the image of a single point. In a confocal or 2-photon microscope (point-scanning devices) the effective PSF depends mostly on the excitation PSF because it’s not necessary to focus emission light when scanning one point at a time. In a well-designed and aligned system the PSF does not differ significantly for different points within the field-of-view of the microscope.

In a widefield microscope, on the other hand, both excitation and emission must be considered. Moreover LSFM has an additional caveat: excitation and emission PSFs

differ not only due to wavelength but also due to different NAs and different orientations in space. In order to create a light sheet a low NA lens is used, and typically the delivery light path is rotated  $90^\circ$  relative to the collection path. This means that the size of the lateral PSF in LSFM is determined by the emission wavelength and the NA of the objective.


**Lateral PSF** While the size of the lateral PSF is uniform for points within the field-of-view, the peak *intensity* of the PSF may not be constant. The most common method to form a light sheet (and the one used for the microscope in Chapter 2) employs an excitation beam with a Gaussian intensity distribution. This results in non-uniform image brightness in the LSFM field-of-view. If this non-uniformity is not accounted for computationally, several negative consequences may ensue:

- Quantification of absolute fluorescence intensity throughout the field-of-view will be biased.
- Moving specimens will exhibit changing brightness distributions that can undermine motion correction via image registration (discussed in subsequent sections).
- The rate of photobleaching/phototoxicity (discussed next section) will not be uniform throughout the field-of-view.

The first two consequences are the most problematic, and fortunately it is relatively straightforward to minimize their impact by measuring the light sheet intensity distribution before the experiment and normalizing pixels by the known excitation intensity. A hardware solution is also available: generating the lightsheet via fast scanning

of a Gaussian beam[41]. Note that these factors exerted minimal effect in the analysis shown in Chapter 2 because a normalized fluorescence measure was used, and the specimen moved little during recording. Future efforts to image freely-moving specimens must give greater attention to these factors.

**Axial PSF** Unlike the lateral PSF the axial PSF in LSFM is determined primarily by the thickness of the light sheet. This provides LSFM the unusual ability to achieve an axial resolution much better than what might be expected based on the NA of the imaging objective. Unfortunately due to diffraction there is a tradeoff between minimum thickness (“waist” size) of the light sheet and the size of the practical field-of-view. A diagram illustrating the geometry of the light sheet is shown in Figure 3.6.



**Figure 3.6:** Shown is the geometry of a light sheet generated with a Gaussian beam, viewed from the side. The horizontal axis corresponds to one lateral dimension of the image while the horizontal axis corresponds to the axial dimension. The constants  $z_r$  and  $w_0$  are determined by the diameter of the Gaussian input beam  $D$ . These constants specify the best axial resolution and the rate of change in axial resolution throughout the field-of-view, respectively. Due to diffraction, a decrease in  $w_0$  corresponds with a decrease in  $z_r$ , resulting in a tradeoff between axial resolution and field-of-view. (image from [64])

Thus the axial width of the PSF is not constant throughout the field-of-view. This may seem inconsequential as long as the worst-case axial resolution is sufficient for the purpose at hand. However there can still be negative consequences due to the fact that the appearance of an object viewed by the microscope is not *translation invariant*. As mentioned in the previous section, this variance with translation can complicate attempts to analyze specimens that move during imaging. Correcting for this variance is more complex than correcting for the brightness difference described previously. One conservative method would be to blur 3D images along the axial dimension with kernels of various widths so that the axial resolution matched the worst-case axial resolution throughout the field-of-view. One could also apply deconvolution to improve the worst-case axial resolution, but this would likely result in artifacts.

Again a hardware solution to this problem has been developed. It is possible to quickly change the focal length of the light sheet optics during a camera exposure so that the waist visits every location in the field-of-view[21]. This produces a nearly uniform axial resolution intermediate between the best-case and the worst-case resolution. It is likely that uniformity could be further improved by using an excitation system that meets the Herschel condition (See Section 1.2.3).

### 3.3.2 Depth effects

When imaging real biological tissue, image quality usually deviates significantly from what might be expected based on PSF measurements. This is because light rays change direction while passing through the tissue due to *scattering*. Scattered light cannot be effectively focused by the imaging system. It may or may not arrive at the image sensor. If it does arrive, it usually arrives at the incorrect location, resulting in degradation of image resolution in a widefield imaging device.

Scattering occurs whenever light passes between media of different refractive indices. Living biological specimens have many components with refractive indices different from that of the typical imaging medium, water. Figure 3.7 lists approximate refractive index values for various cellular components. Longer light paths through the specimen cross more scattering interfaces, resulting in decreasing image quality as a function of depth in the sample.

water	$n \sim 1.33$
extra-cellular fluid	$n \sim 1.34\text{--}1.35$
intra-cellular fluid	$n \sim 1.35\text{--}1.36$
proteins	$n \sim 1.40?$
lipids	$n \sim 1.42?$
DNA	$n \sim 1.44?$
bi-lipid membrane	$n \sim 1.46$ , but BLM is very thin $\therefore$ scattering very weak from cell walls
melanin	$n \sim 1.65?$

**Figure 3.7:** A list of basic biological components and their estimated refractive indices is shown. Image quality is degraded when excitation light or emission light passes through a component with a refractive index differing from the imaging medium (water for live imaging). Question marks indicate that a wide range of values has been published, and the precise value depends on molecular form. (reprinted from [7])

This depth effect is common to *any* existing fluorescence microscope. One implication is similar to that of the PSF inhomogeneities already discussed: the appearance of the image is not invariant to motion of the specimen, particularly to changes in orientation. This effect can be quite severe, especially for widefield imaging methods applied to relatively opaque tissues.

LSFM is atypical in that there are *two* consequential depth axes: the imaging optical axis and the excitation optical axis. Thus 3D LSFM images of scattering tissue exhibit a 2D gradient in image quality. Multi-view LSFM was developed in part to mitigate these depth effects by delivering excitation light along two axes and collecting emission light along two axes [44]. This reduces depth effects in small and thin specimens like

larval zebrafish, but is still problematic for larger specimens. Another promising methodological direction is to apply reagents that homogenize the refractive indices of the tissue and the medium. This has been successful in non-living tissue[89, 99] but no approach has yet been found that is safe for living specimens.


### 3.3.3 Stripe artifacts

LSFM images often exhibit stripe artifacts. These artifacts arise due to interactions between excitation light and tissue. In principle stripes can be caused by both absorption and interference phenomena. Absorption can occur when the light sheet passes through absorptive substances like pigments, leaving a line of shadow that obscures anything behind the pigment. In tissues without pigment there is little absorption at wavelengths commonly used for fluorescence microscopy, and stripes are predominantly generated by interference. When excitation light gets deflected at a scattering interface, excitation intensity is decreased in a line behind the interface and *increased* along other (not necessarily parallel) lines.

The impact of stripes on the analysis of LSFM images ranges from inconsequential to catastrophic depending on the context. Stripe artifacts undermine attempts to precisely quantify fluorescence intensity within striped regions. Since the location and severity of stripes depend on the specific path of excitation light through the sample, sometimes it is possible to reposition the sample so that stripes do not obstruct crucial regions of the image. Fortunately in many cases, such as when imaging calcium in neurons, relative changes in fluorescence are more important than absolute measurements. In these cases stripe artifacts are benign, provided that the sample remains stationary. Any movement of the sample, however, can change the distribution of stripes and disrupt even relative brightness measurements.



Stripes can be removed post-acquisition using computational methods. One such method, described in Section 2.6.14, was applied to destripe the image shown in Figure 2.18. Figure 3.8 shows this image before and after destriping.



**Figure 3.8:** Light sheet microscope images often exhibit stripe artifacts due to absorption and interference phenomena along the axis of light sheet propagation. Shown is the image from Figure 2.18 before and after applying the stripe removal algorithm described in Section 2.6.14.

Our destriping algorithm and several other similar algorithms[91] demonstrate good qualitative results. However timeseries analysis of calcium data must reliably track fluorescence changes of just a few percent, and thus analysis is highly sensitive to artifacts and biases introduced by the destriping procedure. In our experience the artifacts introduced by destriping are too severe to be used in tandem with calcium imaging analyses. Fortunately destriping was not necessary for the neuronal analyses in Chapter 2 for reasons already mentioned. Motile cells within the fish’s body created problematic regions with dynamic stripe patterns, but these regions were avoided when manually selecting ROIs.

We expect that improved destriping algorithms will facilitate analysis of these problematic regions. Once again, however, this problem can be minimized by modifying the microscope hardware. The first microscope designed to address this issue was multidirectional SPIM (mSPIM), a method that rapidly changes the propagation angles within the plane of the light sheet during image acquisition[38]. The final image acquired with mSPIM can be understood as an average of many striped images, each with a different dominant stripe angle. Thus the stripes are averaged out of the final image. Multiview SPIM[44] exhibits less striping for a similar reason.

Recently a method has been developed to prevent stripes by inserting a passive element (a line diffuser) into the excitation path[87]. This approach is comparatively simple to implement but it also introduces a modest increase in light sheet thickness due to manufacturing imperfections.

Neither active nor passive stripe removal can deliver excitation at the full set of angles that would be required to *completely* eliminate artifacts. Therefore in the future highly sensitive analyses may still benefit from algorithms that compensate for the stripe artifact.

### **3.4 Special considerations for live imaging**

In this section we discuss factors affecting the analysis of imaging datasets acquired from living neural tissue. We describe ways in which the imaging process can affect the specimen and how to prevent these effects from biasing quantitative image analyses.

### 3.4.1 Photodamage, Photobleaching, and Phototoxicity

All fluorescence imaging methods introduce energy into the sample by way of excitation light. This energy can damage the tissue, altering the specimen during the imaging process. This damage is especially troublesome in live specimens because the living system may respond actively to the damage, disrupting the process under study. Photodamage can occur without fluorescence: heat damage can result from absorption of illumination light by the specimen[8].

The fluorescence process can result in another type of damage, phototoxicity. Phototoxicity is tissue damage induced chemically as a result of the photobleaching process. Photobleaching is the process by which a fluorophore changes chemical structure due to excitation, losing its fluorescence properties. The mechanism of this molecular change is not fully understood, but it involves a reaction with molecular oxygen that produces free radicals[23]. These free radicals can damage or kill living cells in the specimen. Regardless of damage, the photobleaching also results in a steady decrease in the fluorescence quantified in images. Since fluorescence intensity is often used as a proxy for a biological variable of interest, the photobleaching process can bias results[46].

LSFM exhibits exceptionally low photodamage, photobleaching, and phototoxicity because it confines excitation to the image plane and samples pixels in parallel. However these damaging processes must still be considered with LSFM, especially in the context of high speed recordings: higher imaging rates require higher photon flux in order to maintain signal-to-noise ratio. Photodamage varies depending on properties of the specimen and fluorophore, and thus it should be accounted for on a case-by-case basis. Photobleaching, on the other hand, can often be addressed with postprocessing. One simple way that timeseries analyses compensate for photobleaching is by

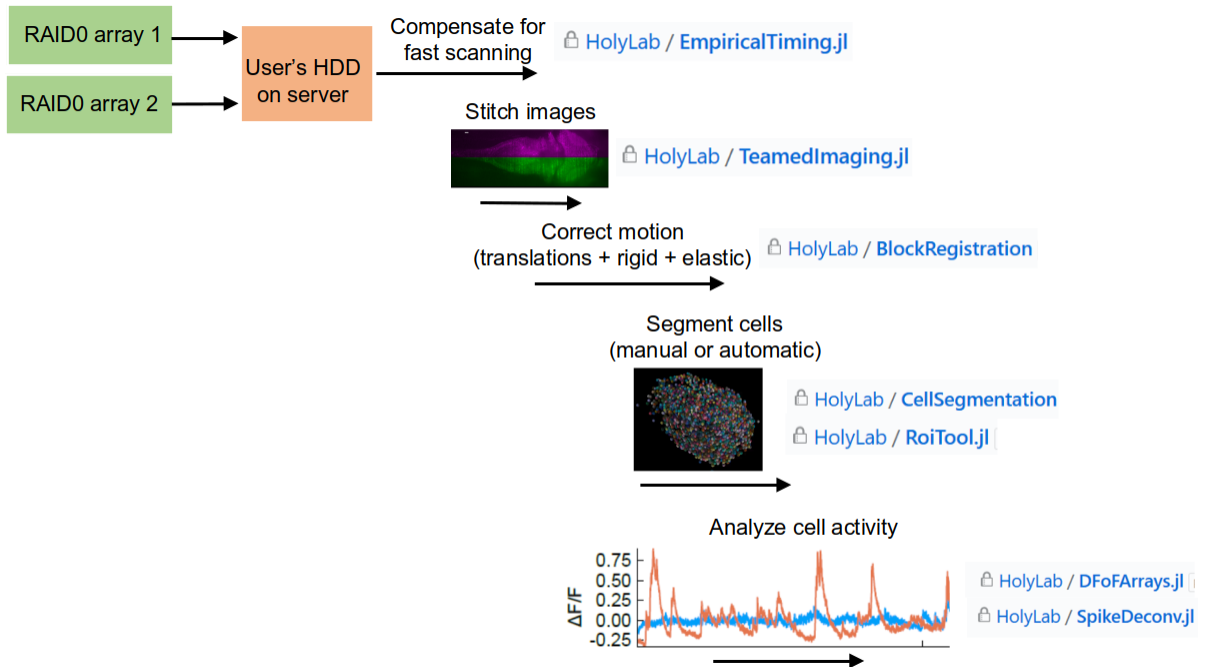
quantifying fluorescence relative to a baseline computed as a moving average (as done in Chapter 2).

### 3.4.2 Unintentional stimulation of neural systems

Neural systems present a special consideration for light microscopy: most organisms exhibit some degree of light sensitivity, and this information is conveyed by the nervous system. Therefore the imaging process itself can be a potent stimulus for the neural system, even when it does not cause photodamage. For this reason neuroscientific studies of retinal function typically avoid imaging with excitation wavelengths that can be sensed by the retinal cells. 2P excitation, which has been integrated with LSFM[90], is a method-of-choice in these cases due to its use of infrared excitation. Due to the higher energy (and thus higher damage) involved with 2P LSFM it is not a universal replacement for 1P LSFM. 1P LSFM remains attractive for fast imaging of the many neural subsystems that exhibit no light sensitivity.

## 3.5 Preprocessing pipeline for optical neurophysiology

In this section we outline a full preprocessing pipeline for optical neurophysiology with the fast OCPI microscope, illustrated in Figure 3.9. The first two steps in the pipeline, fast scanning compensation and image stitching, were already described in Chapter 2. We focus instead on our computational approaches to image registration, neuron segmentation, and deconvolution of spiking activity from calcium imaging data.



**Figure 3.9:** Shown is a preprocessing pipeline for optical neurophysiology data gathered by the fast OCPI microscope. The pipeline begins by moving data from storage on the microscope computer (diagrammed in 3.1) to the user’s storage. The user can then process data on their own server in several steps. The first two steps, already explained in Chapter 2, are specific to the fast OCPI microscope. The remaining steps, described in this section, are appropriate whenever imaging calcium in neurons. The names of relevant code repositories on Github are shown to the right.

### 3.5.1 Image registration

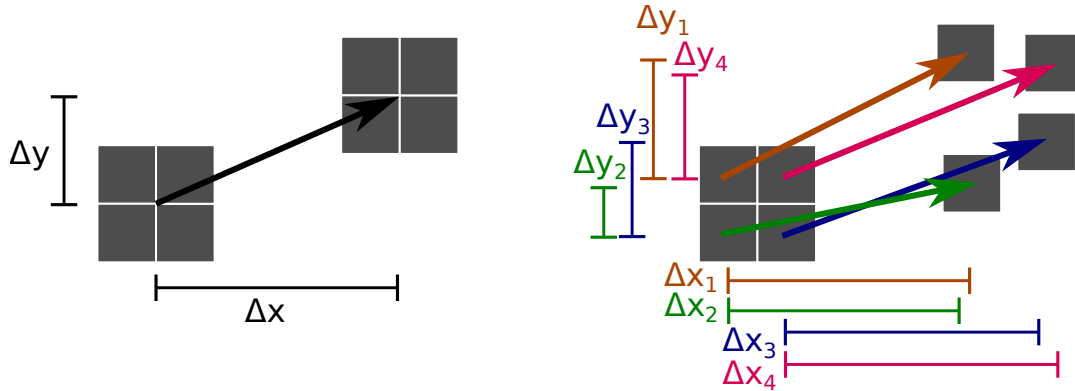
Complications arising due to motion of the specimen were mentioned in Section 3.3 in the contexts of the non-uniformity of the LSFM PSF, depth effects, and dynamic striping artifacts. Even when the PSF is uniform, scattering is minimal, and striping is avoided, specimen motion presents a significant challenge for timeseries analysis of neuronal activity. In the absence of motion, one can track the activity of a neuron over time by simply quantifying moment-to-moment changes at a set of voxel coordinates. In a moving sample the voxel coordinates of a neuron change over time, complicating

tracking of dynamics. Live animals may move significantly during imaging. Even detached tissue samples tend to slowly drift and change shape during recording.

Typically a motion correction preprocessing step is performed to align each image in the timeseries before analyzing dynamic signals. This step is performed by *image registration* algorithms. General image registration is a difficult problem, and remains a very active area of algorithms research. The problem is computationally intensive, but perhaps the more insidious difficulty is that few algorithms are guaranteed to find the *globally optimal* transformation that aligns two images. More insidious still, an image transformation that is numerically optimal may be *impossible* when considering the constraints imposed by the physics of specimen motion. Therefore the art of image registration entails constraining the search space to transformations that are physically possible *and* finding the best transformation within that subspace. Figure 3.10 illustrates these challenges with a simple 2D image consisting of 4 pixels.

As Figure 3.10 suggests, it is most practical to optimize an image transformation model with intermediate expressivity. We, and others independently[68], have chosen an intermediate model that groups pixels into many image blocks. For each block we then optimize a simple translation (diagram in Figure 3.11). These translations together express a complex transformation that allows stretch and compression of images. A regularization is also implemented in order to prevent selection of transformations that are physically impossible. This registration method is highly efficient for two reasons:

1. The elementary blockwise transformation is limited to a translation and therefore can be found efficiently.

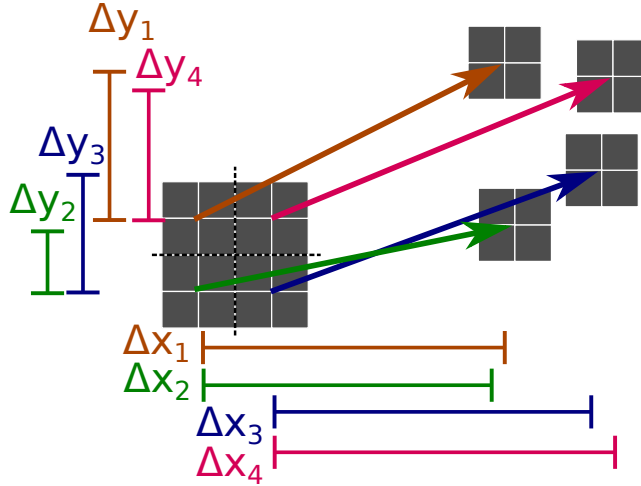


**Figure 3.10:** Shown are two parameterizations of motion for a 4-pixel “specimen”. The left parameterization is highly constrained, describing the motion of *every* pixel with a single pair of displacements along the x and y axes. An efficient algorithm exists for finding the globally optimal  $\Delta x$  and  $\Delta y$ , but this transformation model can only describe simple translations. Biological tissue movement tends to be more complex. A maximally general model is shown on the right, where the motion of each pixel is parameterized independently. This model is *too* general without constraints on the parameter space. For instance, a transformation that swaps the top-right and bottom-left pixels is impossible without physical tearing of the specimen. Such an impossible transformation may result when a registration algorithm is *overfitting* the model to the data. The propensity for overfitting is high given that the number of parameters required for the model on the right is  $2 \times$  the number of datapoints (pixels in the image).

2. The transformation for each block can be computed independently and therefore *in parallel*, as discussed in Section 3.2.2.

An example result of our registration method is shown in Figure 3.12. While this approach is computationally efficient, it does not always find good transformations and requires that the analyst carefully tune hyperparameters. This difficulty seems common to all “deformable” image registration algorithms, and a breakthrough in this field would be a boon for neuroimaging.

Image registration is not only useful for aligning images in a timeseries; it can also be applied to align multiple specimens for comparison. Specimens with stereotyped




**Figure 3.11:** Shown is a parameterization of specimen motion intermediate to the two extremes shown in Figure 3.10. Pixels are grouped into local blocks, and a separate translation describes motion in each block. In this way elastic transformations can be represented with a manageable number of parameters. (Note that the number of parameters no longer outnumbers the number of datapoints as in the right panel of Figure 3.10.) This method is amenable to parallelization because each block can be processed independently. Our full registration algorithm also includes a subsequent step that reconciles blockwise transformations so that the composite transformation respects physical constraints. Also note that the final deformation vector field is interpolated smoothly between block centers.

structure (such as zebrafish brains) can then be compared in a common coordinate system. Comparisons are facilitated by brain “atlases” that serve as templates for alignment of multiple datasets from multiple laboratories. Atlases have enabled substantial advances in the human neuroimaging field[47].

Recently, microscale atlases for aligning larval zebrafish brains have become available[73, 76, 53, 54]. These atlases show great promise, giving access to expert anatomical annotation of brain regions and allowing cross-referencing of molecular information gleaned from confocal imaging of a multitude of fluorescent reporter lines. This includes molecular information that may provide meaningful distinctions between types of neurons based on, for example, the dominant neurotransmitter used by the






**Figure 3.12:** Shown is the result of image registration applied to a single optical section of the larval zebrafish brain, acquired with LSFM. The left image shows an overlay of the section at two timepoints (one green, one red) of the recording. The images align poorly due to motion of the sample during the interval between snapshots. The right image shows the overlay again after applying an image transformation found by a registration algorithm. Alignment is exceptionally good because the sample only moved in the lateral direction; typically there is also motion in the axial direction.

neuron. We see a particularly exciting opportunity to align *functional* LSFM datasets with these atlases. This alignment is termed *cross-modal* registration. Cross-modal registration is more challenging because of the additional variability across microscopy methods combined with individual differences across specimens.

We developed a cross-modal registration pipeline to align our functional LSFM datasets to zebrafish brain atlases. Shown in Figure 3.13 is a set of neurons for which we have both functional data and expert anatomical classifications from the atlas. Similar overlays can be generated for molecular variables such as, for example, the probability that each neuron expresses GABA based on its anatomical location. These atlases continue to improve with new layers of information, and we expect that atlas alignment will become a standard step in many zebrafish functional imaging pipelines.



**Figure 3.13:** Shown is a set of neurons in a larval zebrafish brain that were segmented based on calcium activity recorded with LSM (see Section 3.5.2) and then aligned with an atlas containing expert anatomical annotation of brain regions. Registration to an atlas allows comparisons across animals and laboratories. It also enables cross-referencing of function with molecular information gleaned from the many images of reporter zebrafish lines integrated into the atlas.

### 3.5.2 Extracting neuronal timeseries

Image registration enables the next step in the processing pipeline, extraction of neuronal timeseries from the (motion-corrected) image timeseries. The class of algorithms suited to this task are *image segmentation* algorithms. The majority of image segmentation algorithms have been designed to locate and extract interesting features from static images. Extracting time-varying signals of individual neurons from calcium imaging data requires a very different algorithm.

This is because typically any single image from a calcium imaging timeseries contains insufficient information to locate neurons. The baseline fluorescence of modern calcium indicators[17] is quite low, meaning that neurons may only become visible when

active. Moreover neurons are often densely packed with broadly distributed and overlapping fluorescence profiles, meaning that a pixel cannot “belong” unambiguously to a single neuron. This complication can be circumvented by expressing calcium indicators in a sparse subset of all neurons in order to minimize overlap. In these cases, neurons can be segmented manually by the microscopist with a graphical tool for drawing regions of interest (ROIs). We developed a simple tool for this in Julia and applied it to extract the calcium signals analyzed in Chapter 2. This manual segmentation approach has a number of drawbacks:

- It is very time consuming. 30 person-hours were spent drawing and verifying the 629 neuronal ROIs analyzed in Chapter 2, representing a minority of the total number of active neurons visible in this dataset.
- Much of the neuronal signal-to-noise ratio (SNR) is lost. This is because neuron ROIs must be drawn conservatively, excluding fluorescence measured near cell boundaries in order to prevent contamination by signals from adjacent cells.
- It relies on the expertise/preferences of a single analyst, and therefore it is not easily repeatable.

Cell segmentation is simplified by confining calcium indicator expression to the cell nucleus[43], but this only partially addresses the above issues. Moreover nuclear expression results in calcium signals with slower onsets and offsets, obfuscating fast dynamics. These issues have motivated numerous recent efforts to perform neuron segmentation with algorithms[58, 69, 55, 101, 31, 13].

A highly successful approach has been to formulate the spatiotemporal segmentation problem as a matrix factorization problem:

$$I = S \times T \tag{3.1}$$

Matrix  $I$  represents the imaging dataset with spatial dimensions *unrolled* so that  $I$  has dimensions  $n_{pixels} \times n_{timepoints}$ . Matrices  $S$  and  $T$  are unknown, with sizes  $n_{pixels} \times n_{neurons}$  and  $n_{neurons} \times n_{timepoints}$ , respectively. Also unknown is  $n_{neurons}$ . Several recent algorithms[69, 55, 101, 31, 13] have focused on applying appropriate constraints to this factorization problem and inferring  $n_{neurons}$  so that the factorization procedure produces:

- a matrix  $S$  with the spatial profile of each neuron occupying a different column of  $S$ , and possibly with a single column devoted to a “background” fluorescence distribution.
- a matrix  $T$  with the temporal dynamics of each neuron occupying a different row of  $T$ , and possibly with a single row devoted to dynamic “background” fluorescence.

Thus this formulation models the fluorescence measured at each pixel in the image as a linear combination of fluorescence from multiple partially-overlapping sources.

The background terms mentioned above can be understood as the baseline neuronal calcium levels when neurons are inactive. Allowing the background to vary over time enables modeling the gradual decrease in baseline fluorescence from photobleaching (discussed in Section 3.4). In order to get good results from this approach it is helpful to constrain the factorization in several ways:

1. Force elements of  $S$  and  $T$  to be non-negative, respecting the fact that negative fluorescence has no physical basis. This converts the problem to a non-negative matrix factorization (NMF).
2. Require sparsity in  $S$ , reflecting the expectation that each neuron only occupies a small fraction of the image space.
3. Require locality in  $S$ , reflecting the expectation that each neuron occupies a contiguous region of space.
4. Require signals in  $T$  to respect the expected dynamics of the spike-induced calcium transient (discussed further in next subsection).
5. Employ some mechanism to reject signals with insignificant SNR (reduce  $n_{neurons}$ ).

We refrain from discussing specific mechanisms for implementing the constraints listed above. Progress in this area has been rapid, and we see opportunities for further improvements in algorithmic efficiency. For example, the locality constraint described above (item #3) suggests that segmentation process might be efficiently parallelized (see also discussion of patchwise algorithms in Section 3.2.2). We also note that constraint #4 above may be problematic. Recent data suggests that the time constants associated with the calcium response function vary significantly from neuron-to-neuron[92]. Our own NMF-based segmentation algorithm achieves good results without enforcing this constraint. Uncertainty in the calcium response function also complicates efforts at spike deconvolution, described next.

### 3.5.3 Spike deconvolution

Before the development of calcium imaging, analyses of neural systems predominantly focused on the timing and rate of action potentials measured through electrophysiology. This focus on spiking activity is perhaps justified given that the action potential seems to be the dominant mode of information transmission among many types of neurons (graded potentials and slower neuromodulatory processes are examples of other modes). Thus calcium signals are often viewed as an indirect measure of neural spiking activity, and there is great interest in extracting information about spike timing from calcium signals.

Individual spiking neurons have been shown to exhibit calcium responses with characteristic rise times and fall times following a single action potential (See also Section 1.2.2). This repeatability suggests that, given a known calcium response function, it should be possible to *deconvolve* the calcium signal to find the timing of individual spiking events. When performed successfully, calcium deconvolution affords a significant increase in temporal resolution, matching or even exceeding the calcium sampling rate and bringing calcium imaging closer to the resolution of electrophysiology[95]. Several calcium deconvolution methods have been developed with varying degrees of computational efficiency.

We developed a new algorithm that lies at the efficient extreme of this spectrum, running in  $O(n)$  time, where  $n$  is the number of timepoints in a neuronal timeseries. The algorithm, derived in Appendix C, performs efficient Wiener deconvolution. Wiener deconvolution is a long-established deconvolution method applied in many engineering fields[98]. The method can be applied to produce a timeseries of “spikes” guaranteed to be the least-squares optimal explanation for the observed calcium signal. Naive Wiener deconvolution requires operating on the timeseries in the Fourier domain.

Thus a fast Fourier transform is required, which runs in  $O(n \log n)$  time. We show that when the decay of the calcium response is modeled as an exponential (as is standard practice), it is possible to perform Wiener deconvolution entirely in the time domain with just two passes through the dataset. The algorithm in Appendix C extends another  $O(n)$  algorithm derived using the continuous time Fourier transform[36]. Our new method, derived instead using the discrete time Fourier transform, shows improved performance for timeseries data acquired at sampling rates much slower than the action potential.

Despite a long history of algorithm development[83, 36, 95, 92, 30], calcium deconvolution has not yet become a standard preprocessing step. This may be in part due to the limited sampling rates of microscopes, a barrier that we sought to lower with the microscope described in Chapter 2. However in our experience there is another significant barrier, the inconsistency of calcium response parameters across neurons. A recent study estimated that the decay constants of neurons in the larval zebrafish brain commonly differ by as much as 1 s[92]. Typically these time constants are required inputs to deconvolution algorithms. While methods exist for estimating these constants from calcium signals, our experience with a recent method[30] is that estimation is unreliable. Another approach to estimation was recently proposed[92]; we expect that this will remain an active area of research. It's also possible that advances in voltage-sensitive fluorophores will obviate the need for calcium deconvolution, but at the moment these fluorophores are not stable enough to replace calcium indicators for recordings lasting several minutes[32, 84, 67].

## 3.6 Future work interpreting neuronal dynamics

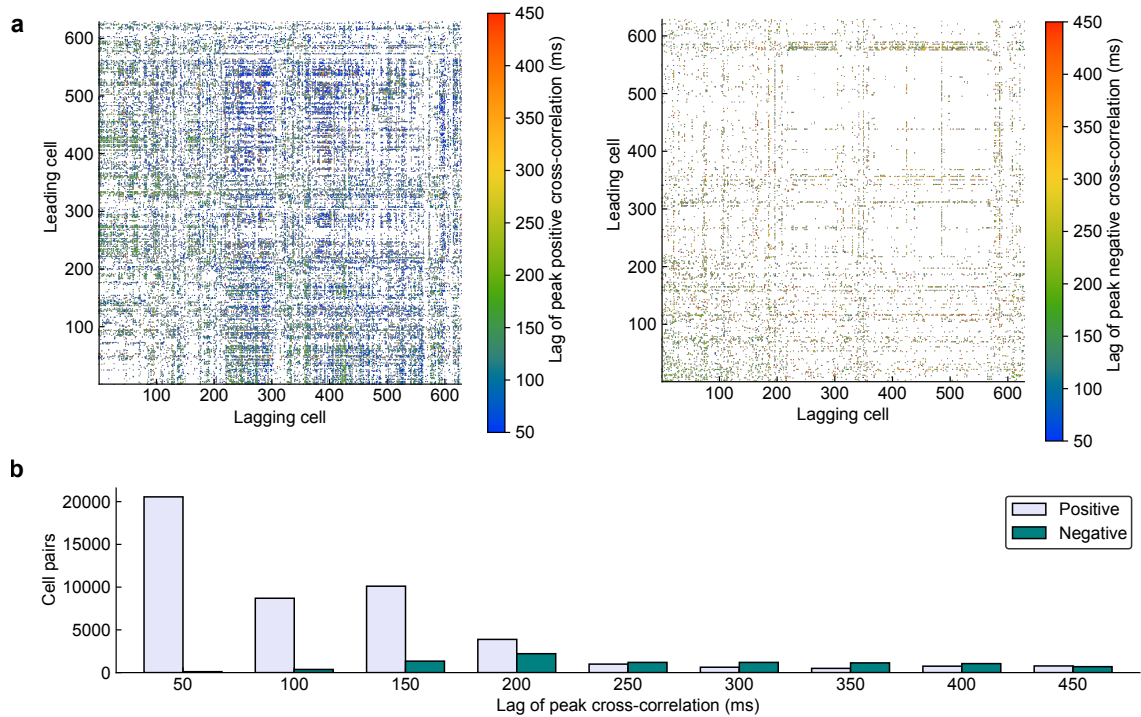
The flexible design of the fast OCPI microscope hardware and software paves the way for adding new capabilities to the system in the future. Likewise our development of a preprocessing pipeline for calcium imaging lays the groundwork for ambitious analyses of neural networks including thousands of neurons. In this chapter we discuss analyses designed to extract circuit parameters from preprocessed calcium imaging data as well as ways that these analyses will be enhanced by the addition of a new point-scanning module allowing targeted stimulation of the zebrafish brain during OCPI imaging.

In Section 2.4 we presented a correlation analysis of a set of neurons recorded at high speed in the larval zebrafish forebrain. Correlation is one of the simplest and most common metrics used to describe relationships between neurons or brain regions. We chose correlation in order to estimate the impact of the heartbeat artifact on analyses; if correlations are affected then more complex analyses are also likely affected.

With greater confidence in our ability to manage artifacts contaminating neuronal timeseries we now consider more complex analyses. One attractive direction is to move from metrics, like correlation, that *describe* the data and seek instead to *explain* observations with a model. If we can fit a model that successfully predicts the behavior of the neural system in a variety of contexts then this model will likely shed light on the mechanisms generating the data. Before discussing an explicit model first let us examine the explanatory power gained through an incremental change in our analysis, computing *cross-correlations* instead of correlations.



### 3.6.1 Cross-correlation analysis



**Figure 3.14:** (a) Cross-correlation analysis reveals predictive relationships between neurons. Cross-correlations were calculated for each pair of neurons for lag values in the interval 0 s to 1 s (21 lag values per pair). First we identified the subset of statistically significant ( $p \leq 0.01$ ) cross-correlation values (statistics below). If a neuron pair had significant cross-correlations at multiple lag values then a color was assigned by the lag at the most extreme (“peak”) cross-correlation value. Only the subset of lags that are nonzero and less than 500 ms are colored, with the remainder in white. Positive and negative peaks in the cross-correlation are shown separately (left and right, respectively) (b) Histogram summary of the data shown in (a).

**Statistical significance** Significance of cross-correlations was determined as follows: A null hypothesis was computed separately for each neuron pair by first reshuffling the neurons’ traces and computing the correlation 10000 times. The threshold for significant positive correlation with a single measurement was then the 99th percentile value of the resampled correlations. However since correlation was measured for each of the 21 lags we applied a Bonferroni correction, and the corrected threshold was then the 99.952th percentile value. The same procedure determined the significance of negative cross-correlations, with the 0.047th percentile value as an upper bound.

Figure 3.14 shows a cross-correlation analysis of the dataset analyzed in Section 2.4. A cross-correlation is simply a correlation computed after shifting two timeseries forward or backward in time relative to one another. The extent of this forward or backward shift is labeled as “lag” in Figure 3.14. Further details are discussed in the figure caption. This analysis illustrates that a subset of neuron pairs exhibit an asymmetric predictive relationship: positive changes in one neuron’s fluorescence predict either positive or negative changes for another neuron *at a later time*. Interestingly, significant positive peaks in the cross-correlation tend to occur at shorter lags than negative peaks. Regardless of how this is interpreted it is clear that unlike (zero-lag) correlation, cross-correlation can identify neurons that tend to precede or follow one another in their activity changes.

### 3.6.2 Predictive modeling

While cross-correlation can reveal predictive relationships it does not model network dynamics. What gains might be made by moving from cross-correlation to an explicit model? Consider a model that predicts a neuron’s dynamics based on linear regression with previous activity two other neurons as input:

$$c_t = w_a a_{t-\Delta t} + w_b b_{t-\Delta t} \tag{3.2}$$


where  $a$ ,  $b$ , and  $c$  are neuronal timeseries vectors indexed by  $t$  and the  $w$  values are scalar weights that must be fit to the data. This can be generalized to a matrix equation for predicting the activity of *each* neuron via the previous states of all neurons

$$Y = Y_{shift}W \tag{3.3}$$

where  $Y$  is a  $n_{timepoints} \times n_{neurons}$  matrix of neuronal timeseries,  $Y_{shift}$  is the same data but time-shifted by a lag (as is done for cross-correlation), and  $W$  is a  $n_{neurons} \times n_{neurons}$  matrix of weights to be fitted. This formulation has two advantages over cross-correlation:

- The fitted weights in  $W$  have a natural interpretation as the *influence* that neurons exert on network dynamics, and under certain conditions may approximate synaptic weights (discussed further below).
- The regression procedure tends to produce a  $W$  that is more sparse (and thus easier to interpret) than the cross-correlation matrix.

Neither of the above statements is entirely obvious. The second statement we justify analytically in Appendix D. Empirical results, shown in Figure 3.15, also support the first statement. The first statement, that  $W$  may contain cause-and-effect relationships, requires more supporting evidence. While we offer no mathematical proof, the following paragraphs discuss key constraints and ways to improve the likelihood of extracting concrete circuit parameters (such as synaptic weights) with predictive modeling.



**Figure 3.15:** Shown is the weight matrix resulting from solving equation 3.3 for the same dataset analyzed in Figures 2.22 and 3.14 with a lag of a single sample (50 ms). Note that the matrix is quite sparse relative to the correlation (and cross correlation) matrix. The reader is encouraged to zoom in with a pdf viewer. Self-prediction weights (values along the diagonal) tend to be large and positive, and they are omitted so that the colorscale emphasizes interaction weights between neurons. The largest weights tend to be close to the diagonal, which is perhaps not surprising given that neurons are ordered by increasing slice depth. This agrees with the intuition that nearby neurons are more likely to influence one another.

**Fast imaging** The potential to model causal relationships hinges on the fact that presynaptic activation must precede postsynaptic activation (by definition). If the sampling rate of the microscope is too slow then activations may appear to occur simultaneously, and the temporal relationship cannot be modeled. How fast is fast enough? In the case of a single synaptic interaction it may be necessary to sample as often as *every millisecond*. While the microscope described in Chapter 2 made strides in this regard, this is still impossible for all but the smallest image volumes. There are, however, reasons for optimism. Quite often a single spike in a presynaptic neuron is insufficient to trigger a spike in the postsynaptic neuron. Instead, temporal summation (multiple presynaptic spikes) and/or spatial summation (presynaptic

spikes from multiple input neurons) is required. This may alleviate sampling rate constraints, at least for some of the synaptic interactions in the circuit.

**Prediction at multiple timescales** Another optimistic point is that if the sampling rate is too slow to detect single-synapse relationships it may still be highly useful for detecting relationships between neurons separated by multiple synapses. The same logic applies for neuromodulatory interactions known to transpire over seconds or even minutes. These points highlight a weakness of modeling activity based on a *single* lag; ideally a model should be able to explain interactions at multiple timescales. Multiple lags can be modeled simply:

$$Y = Y_{shift1}W_1 + Y_{shift2}W_2 + Y_{shift3}W_3\dots \quad (3.4)$$

Note that all shifted  $Y$  matrices can be combined into one matrix, and the same is true for the  $W$  matrices, meaning that the least-squares solution can be obtained in exactly the same way as in the single-lag case. However the size of the matrices involved grows quickly, scaled by the number of lags in the model. This scaling presents a challenging computational problem. For instance, in order to model activity of 10,000 neurons based on 100 lags one must fit a  $W$  of size  $1,000,000 \times 10,000$ . It may be possible to develop tractable fitting methods for data of this scale, but perhaps a more serious problem is the potential for *overfitting* with a large model. As was discussed in the context of image registration (Section 3.5), fitting a more complex model (larger  $W$ ) without increasing the number of datapoints is a recipe for overfitting.

**The need for ground truth data** It is impossible to recognize and prevent overfitting without some notion of “ground truth” against which to evaluate the performance

of the model. Classically this is done by splitting the data into “training” and “test” sets, omitting the test set from training, and quantifying performance only on the test set. This will be an important step when evaluating models of neural dynamics, but this would be complemented by additional ground truth in other forms. If, for instance, in addition to recording function we could also map all synaptic connections in the same specimen with electron microscopy (EM), then we could effectively evaluate the model’s ability to capture synaptic information.

EM is perhaps the most time-intensive and labor-intensive way to perform this evaluation. A more efficient way may be to image a circuit for which connections have already been mapped in previous work. In this case it would be important to choose a circuit with stereotyped connectivity. *C. elegans* may then be an obvious choice, but any system for which we have statistical information (if not precise counts) about neuron types and connections would be conducive to this analysis.

**Broad sampling of circuit contexts** Fitting an accurate functional model of a neural circuit will require more than a high sampling rate; it will also be necessary to sample network activity in a sufficient set of *contexts*. If a neuron never activates during a recording, then it cannot possibly enter into the model. Moreover there may be multiple network states that could activate a neuron, with each state employing a different set of synapses. If the network does not visit all of these states during the recording then information will be missing from the model. These issues suggest that experiments should be specifically designed to push the network into a diverse set of states.

In an intact animal one way to do this would be to deliver a variety of stimuli to the animal’s sensory systems. Perhaps a well-designed stimulus set could traverse a

broad swath of the state space of sensory neural networks. The space traversed in other networks (i.e. motor) may be comparatively smaller. Thus generating a more robust and *comprehensive* model may require activating neurons by other means in addition to delivering sensory stimuli.

**Targeted neural activation with optogenetics** Optogenetics has recently emerged as a tool to activate neurons by delivering light, activating light-sensitive ion channels in the neural membrane[22]. Precise activation of neurons can be achieved either by targeting expression of the actuator genetically or by focusing light on specific neurons in a circuit. We have taken the latter approach to enable targeted optogenetic stimulation while imaging with the fast OCPI microscope, integrating a new point-scanning module (Figure 3.16).

Using this point-scanning module in tandem with optogenetics will allow us to selectively activate any neuron (or set of neurons in a confined region) in the specimen. Even with this ability we may only be able to traverse a small region of a state space involving 1000s of neurons. A conservative approach would then be to focus on small and relatively isolated neural subsystems including 10s of neurons rather than 1000s. A less-conservative view is that we may be able to infer important circuit characteristics without comprehensive sampling of network states. Moreover the natural state space of the system may be *low dimensional*, a possibility suggested by repeated observations that network activity occupies a state space much smaller than that suggested by the number of neurons recorded[16, 18, 50, 5, 59]. These observations may be misleading estimates of dimensionality, arising due to the limited set of contexts explored during experiments. Ultimately this issue will be resolved empirically, and tools like those presented in this dissertation will enable this empirical process.



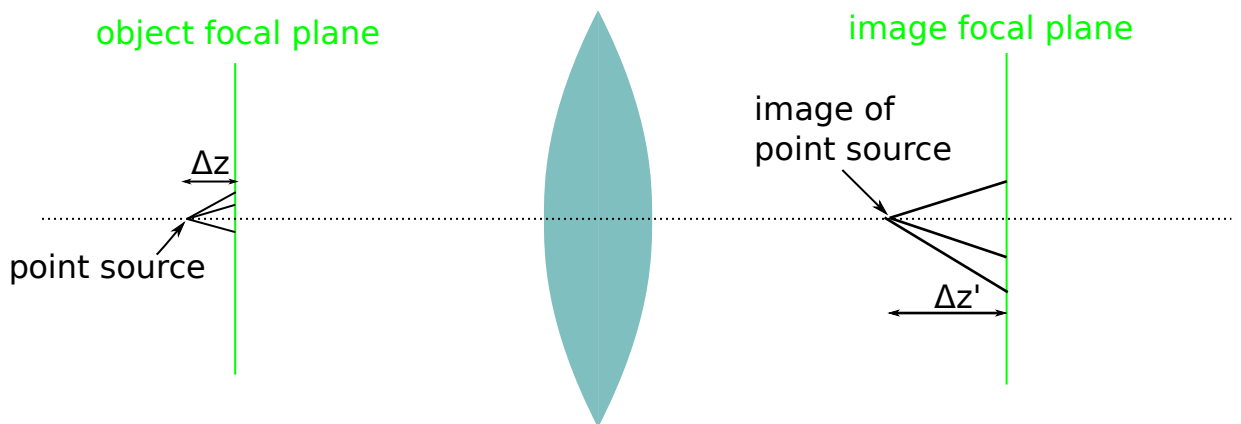


# Appendix A

## Spherical aberrations from remote focusing

Cody Greer and Timothy E. Holy

Consider an optical system that satisfies the Abbe sine condition for two planes, called the object focal plane and image focal plane. Consider a point source along the optic



axis separated by  $\Delta z$  from the object focal plane. Rays are emitted from this object over a range of angles and can be traced to the object focal plane.

The sine condition states that each ray satisfies

$$n \sin \theta = M n' \sin \theta', \quad (\text{A.1})$$

where  $n$  is the refractive index of the object immersion medium,  $n'$  that of the imaging immersion medium (here taken to be 1 for air),  $M$  the local linear magnification near the axis, and  $\theta$ ,  $\theta'$  are the ray angles in the object and image space, respectively. For a ray of angle  $\theta$ , let  $h$  denote the height of the strike position of each ray in the object focal plane, where  $h = \Delta z \tan \theta$ . We can calculate the position  $\Delta z'_\theta$  where this intersects the optic axis:

$$\Delta z'_\theta = \frac{h'}{\tan \theta'} \quad (\text{A.2})$$

$$= \frac{M h}{\frac{\sin \theta'}{\sqrt{1 - \sin^2 \theta'}}} \quad (\text{A.3})$$

$$= \frac{M \Delta z \tan \theta}{\frac{\frac{1}{M} n \sin \theta}{\sqrt{1 - \frac{1}{M^2} n^2 \sin^2 \theta}}} \quad (\text{A.4})$$

$$= \frac{M^2}{n} \Delta z \frac{\sqrt{1 - \frac{1}{M^2} n^2 \sin^2 \theta}}{\sqrt{1 - \sin^2 \theta}}. \quad (\text{A.5})$$

Note that the strike position is a function of  $\theta$ , and therefore in general we do not have perfect focus. The only time the angle-dependence disappears is when the ratio of square-roots cancels, which requires  $|M| = n$ . This is the condition derived more generally by the Maxwell perfect-imaging theorem. (We can also see that the longitudinal magnification is  $\frac{M^2}{n}$ .)

When this condition does not hold, the fractional range of strike positions (up to the maximum angle consistent with the numerical aperture NA) is

$$\frac{\sqrt{1 - \frac{\text{NA}^2}{M^2}}}{\sqrt{1 - \frac{\text{NA}^2}{n^2}}} - 1 \quad (\text{A.6})$$

For  $M = 20$  and water immersion, at NA 0.5 we get approximately 8%, and at NA 1 we get a 50% spread.

The axial RMS spot radius offers a more informative summary of the imaging performance. This quantity is the square root of the mean squared error (MSE) in strike position for every ray within the collection cone of the objective lens. For a ray with angle  $\theta$  emitted from a source at  $\Delta z$  the squared error is

$$(\Delta z'_\theta - z'_0)^2 = \left( \Delta z'_\theta - \frac{M^2}{n} \Delta z \right)^2 \quad (\text{A.7})$$

For each angle  $\theta$  an ideal point source at  $\Delta z$  emits a circle of rays. It can easily be shown with trigonometry that the total light collected at angle  $\theta$  is proportional to  $\tan(\theta)$ . Therefore we scale the error at each angle  $\theta$  accordingly to get the sum squared error and then divide to calculate MSE:

$$\text{MSE} = \frac{\int_0^{\theta'} \tan(\theta) \left( \Delta z'_\theta - \frac{M^2}{n} \Delta z \right)^2 d\theta}{\int_0^{\theta'} \tan(\theta) d\theta} \quad (\text{A.8})$$

$$= \frac{\int_0^{\theta'} \tan(\theta) \left( \Delta z'_\theta - \frac{M^2}{n} \Delta z \right)^2 d\theta}{-\ln(\cos(\theta'))} \quad (\text{A.9})$$

The upper limit,  $\theta'$  is the maximum angle that the objective can collect as specified by the NA. Note that the above calculates the RMS spot radius in image space; in order to convert this to object space units divide by the axial magnification.

# Appendix B


## Microscope components and drawings

**Parts List** (ordered approximately from excitation to collection of emission)


Number in diagram	Part description	Vendor	Custom?	Part number	Price (USD)	Count	
not shown	Laser system w/5 wavelengths, AOTF	Spectral (now Andor)	no	Laser Merge Module (LMM)	90000	1	*discontinued, estimated price
1	Pigttailed fiber collimator	OZ Optics	no	LPC-01-405/640-3/125-P-1-6AC-40-3S-3-2	146.5	1	
2	Pigttailed fiber collimator holder	(machine shop)	yes		NA	1	
3	Lightsheet lens module ("cigarette")	(machine shop)	yes		NA	1	
not shown	Lightsheet achromat lens	Edmund Optics	no	45-262	87.5	1	
not shown	Lightsheet cylinder lens	Tower Optical	yes*		99	1	
4	Rotation collar	(machine shop)			NA	1	
5	Clamp for lightsheet lens holder	(machine shop)			NA	1	
not shown	10x, 0.3 N.A. objective (UMPLFLN10X/W)	Olympus	no	1-U2M583	774.25	1	
6	20x, 0.5 N.A. objective (UMPLFLN20X/W)	Olympus	no	1-U2M585	1464.64	1	
not shown	40x, 0.8 N.A. objective (LUMPLFLN 40X/W)	Olympus	no		2500	1	
7	Objective holder (RMS)	(machine shop)	yes				
8	Mini-dovetail stage (2-axis)	Lightspeed Technologies Inc.	no	MDE266	304	1	
9	Mini-dovetail stage (3-axis)	Lightspeed Technologies Inc.	no	MDE269	573	1	
10	dovetail stabilizer	(machine shop)	yes		NA	1	
11	Front piezo plate	(machine shop)	yes		NA	1	
12	Piezo positioner, 800um range	Piezosystem Jena	no	NanoSX800	11205	1	
not shown	Piezo amplifier (digital control)	Piezosystem Jena	no	30DV300	5616	1	
13	Rear piezo plate	(machine shop)	yes		NA		
not shown	Precision broadband mirror	Edmund Optics	no	48-017	395	1	
not shown	200mm tube lens	Thorlabs	no	ITL200	450	1	
not shown	Knife edged mirror	Thorlabs	no	MRAK25-G01	125.46	1	
not shown	50/50 beamsplitter (25 x 36 mm)	Thorlabs	no	BSW10R	110	1	
not shown	Filter cube w/insert	Thorlabs	no	DFM	304	1	
not shown	Extra filter cube insert	Thorlabs	no	DFMT1	201	1	
not shown	0.9x telecentric relay lens	Edmund Optics	no	62-902	2265	2	
not shown	CMOS camera	PCO	no	Edge 4.2	16400	2	
not shown	Stages for aligning cameras	Thorlabs	no	DTS25	179.5	4	
not shown	DAQ board (PCI-6259)	National instruments	no	779072-01	1592.1	1	
not shown	Physiology stage surface	Thorlabs	no	PHYS24BB	2500	1	
not shown	Lab jack	Newport	no	281	999.94	1	
not shown	Breadboard for connecting lab jack	Thorlabs	no	MB1224	259	1	
not shown	XY microscope translation stage	Scientifica	no	N/A (Quote ref: QLS-30894)	4042.5	1	
not shown	RAID hard drives	Seagate	no	4221403	78.19	20	
not shown	Breadboard for vertical mounting of system	Thorlabs	no	MB1824	400	1	
				Price total:	163760.69		

**\*Cylindrical lens specifications**


Diameter	3mm +0/-0.2mm
Center Thickness	1mm ±0.1mm
Edge Thickness	1.37mm
Effective Focal Length	-6.25mm
Back Focal Length	-6.91mm
Focal Length Tol	±3%
Radius	-3.24mm
Surface quality	60/40 both sides
Coating	MgF2







view from front



view from left




view from rear left




Created by: <b>Cody Greer</b>		Title: <b>Pigtailed collimator holder</b>		
Supplementary information:  Clamps onto pigtailed collimator Attaches to MDE266 dovetail slide Addtl authors: T.E. Holy J. Kreitler K. Poenicke				
Size: <b>A3</b>	Sheet: <b>X / Y</b>	Scale: <b>mm</b>		
Part number: <b>NA</b>				
Drawing number: <b>DN</b>				
Date: <b>17/10/2017</b>		Revision: <b>REV A</b>		







DESIGNED BY: C. Greer, T.E. Holy		Lightsheet Lens Module		G	-
DATE: 21 Dec 2017		"cigarette"		F	-
SIZE: <b>A4</b>				E	-
				D	-
SCALE	WEIGHT (kg)	DRAWING NUMBER	SHEET	C	-
		Addtl designers: J. Kreitler, K. Poerlicke		B	-
This drawing is our property; it can't be reproduced or communicated without our written consent.				A	-










1	3.0 x 0.5
2	counterbore, 3mm depth
3	alignment pin hole, 6mm depth
4	4.1mm dia through
5	M32 threaded hole


DESIGNED BY: J Kreittler, TE Holy, C Greer		I H G F E D C B A
DATE: 2014		
CHECKED BY: SUPERVISOR NAME DATE: CHECK DATE		1 1 1 1 1 1 1 1 1
SIZE <b>A3</b>		
SCALE	WEIGHT (kg)	DRAWING NUMBER NUMBER
SCALE	WEIGHT	
SHEET		SHEET
SHEET		
This drawing is our property; it can't be reproduced or communicated without our written agreement.		



1	3.0 x 0.5
2	counterbore, 3mm depth
3	alignment pin hole, 6mm depth
4	4.1mm dia through
5	M34 threaded hole

DESIGNED BY: J Kreittler, TE Holy, C Greer		H H G F E D C B A
DATE: 2014		
CHECKED BY: SUPERVISOR NAME OCPI 2		1 1 1 1 1 1 1 1 1
DATE: CHECK DATE		
SIZE: <b>A3</b>		
SCALE	WEIGHT (Kg)	
SCALE	WEIGHT	DRAWING NUMBER
		SHEET
This drawing is our property; it can't be reproduced or communicated without our written agreement.		






DESIGNED BY:  
C. Greer, K. Poenicke  
DATE:  
1 November 2017  
CHECKED BY:  
NA  
DATE:  
NA

### Dovetail stabilizer

Applies tension to dovetail joint

SIZE  
**A3**



SCALE  
mm

WEIGHT (kg)  
NA

DRAWING NUMBER  
NA


SHEET  
NA

I	-
H	-
G	-
F	-
E	-
D	-
C	-
B	-
A	-

This drawing is our property; it can't be reproduced or communicated without our written agreement.








DESIGNED BY:  
Cody Greer  
DATE:  
19 February 2017  
CHECKED BY:  
SUPERVISOR NAME  
DATE:  
CHECK DATE

**Rear Piezo Plate**  
OCPI 2 Microscope

SIZE  
**A3**



SCALE  
SCALE

WEIGHT (kg)  
WEIGHT

DRAWING NUMBER  
NUMBER

SHEET  
SHEET

This drawing is our property; it can't be reproduced or communicated without our written agreement.

I	-
H	-
G	-
F	-
E	-
D	-
C	-
B	-
A	-

4  
3  
2  
1

4  
3  
2  
1

H G F E D C B A

# Appendix C

## Derivation of Wiener filter for calcium deconvolution

Cody Greer and Timothy E. Holy

An “exponential decay” discrete-time impulse response may be written

$$r_a(n) = u(n)a^n, \tag{C.1}$$

where  $u$  is the Heaviside function

$$u(n) = \begin{cases} 1 & (n \geq 0) \\ 0 & (n < 0) \end{cases} \tag{C.2}$$

and  $|a| < 1$ . Given a “source signal”  $s(n)$ , the measured signal  $m(n)$  can be written

$$m(n) = \sum_{n'=-\infty}^{\infty} s(n-n')r_a(n') + \eta(n), \quad (\text{C.3})$$

where  $\eta(n)$  is the noise in the measurement (here taken to have zero mean). When the noise is zero, the particular case of Eq. (C.1) can be written more simply as

$$m(n) = s(n) + am(n-1). \quad (\text{C.4})$$

Now suppose we have  $m(n)$  and want to infer  $s(n)$ . In the absence of noise

$$s(n) = m(n) - am(n-1). \quad (\text{C.5})$$

However, this is not optimal in the presence of noise.

We will tackle this by Wiener deconvolution. Eq. (C.1) has DFT

$$\hat{r}_a(\omega) = \sum_{n=-\infty}^{\infty} r_a(n)e^{-i\omega n} = \sum_{n=0}^{\infty} (az)^n = \frac{1}{1-az}, \quad (\text{C.6})$$

where we have set  $z = e^{-i\omega}$ . The optimal (Wiener) inverse filter  $w(n)$  has DFT

$$\hat{w}(\omega) = \frac{\hat{r}^*}{|\hat{r}|^2 + \eta^2} = \frac{1/\hat{r}}{1 + \eta^2/|\hat{r}|^2}, \quad (\text{C.7})$$

where  $*$  means complex conjugation and here we’re using  $\eta^2$  as a constant encoding mean-square noise. (Here we assume the noise is additive and thus independent of

s.) Substituting Eq. (C.6) into Eq. (C.7), we obtain

$$\hat{w}(\omega) = \frac{1 - az}{1 + \eta^2(1 - az)(1 - a/z)} \quad (z^* = 1/z) \quad (\text{C.8})$$

$$= -\frac{z(1 - az)}{a\eta^2 z^2 - [1 + \eta^2(1 + a^2)]z + a\eta^2}. \quad (\text{C.9})$$

The denominator can be expressed as  $cz^2 + bz + c$  if we set

$$b = -[1 + \eta^2(1 + a^2)] \quad (\text{C.10a})$$

$$c = a\eta^2, \quad (\text{C.10b})$$

and may therefore be written

$$c(z - \gamma_+)(z - \gamma_-) \quad (\text{C.11})$$

where

$$\gamma_{\pm} = \frac{-b \pm \sqrt{b^2 - 4c^2}}{2c}. \quad (\text{C.12})$$

Since  $|a| < 1$ , we have

$$b^2 - 4c^2 > \eta^4 [(1 + a^2)^2 - 4a^2] > 0, \quad (\text{C.13})$$

so both roots will always be real and distinct from one another. Computing  $\gamma_{\pm}$  seems best done numerically. Note that if  $c^2 \ll b^2$  there can be a delicate cancelation in  $\gamma_-$ .

The more numerically stable form is

$$\gamma_- = \frac{-b - \sqrt{b^2 - 4c^2}}{2c} \quad (\text{C.14})$$

$$= \frac{-b - \sqrt{b^2 - 4c^2}}{2c} \frac{-b + \sqrt{b^2 - 4c^2}}{-b + \sqrt{b^2 - 4c^2}} \quad (\text{C.15})$$

$$= \frac{b^2 - (b^2 - 4c^2)}{2c(-b + \sqrt{b^2 - 4c^2})} \quad (\text{C.16})$$

$$= \frac{2c}{-b + \sqrt{b^2 - 4c^2}} \quad (\text{C.17})$$

$$= 1/\gamma_+. \quad (\text{C.18})$$

Therefore we will always have one  $\gamma_{\pm}$  satisfying  $|\gamma_{\pm}| < 1$  and the other  $|\gamma_{\mp}| > 1$ . Since  $b < 0$  and  $c > 0$ , we have  $\gamma_+ > 1$  and  $0 < \gamma_- < 1$ . Therefore let us simply define

$$\gamma = \frac{2c}{-b + \sqrt{b^2 - 4c^2}} \quad (\text{C.19})$$

as the one for which  $|\gamma| < 1$ .

Substituting Eq. (C.11) into Eq. (C.9), we obtain

$$\hat{w}(\omega) = -\frac{z(1 - az)}{c(z - 1/\gamma)(z - \gamma)} \quad (\text{C.20})$$

$$= \frac{Az + B}{z - 1/\gamma} + \frac{Cz + D}{z - \gamma} \quad (\text{C.21})$$

$$= \frac{(A + C)z^2 + (B - \gamma A + D - C/\gamma)z - B\gamma - D/\gamma}{(z - 1/\gamma)(z - \gamma)}, \quad (\text{C.22})$$

which implies

$$A + C = a/c \quad (\text{C.23})$$

$$B - \gamma A + D - C/\gamma = -1/c \quad (\text{C.24})$$

$$B\gamma + D/\gamma = 0. \quad (\text{C.25})$$

Since there are four constants and only three constraints, we may freely choose  $B = 0$ , which then sets  $D = 0$  and we obtain

$$A = -\frac{\gamma}{c} \frac{1 - a/\gamma}{1 - \gamma^2}; \quad (\text{C.26a})$$

$$C = \frac{\gamma}{c} \frac{1 - a\gamma}{1 - \gamma^2}. \quad (\text{C.26b})$$

Consequently we can write this as

$$\hat{w}(\omega) = \frac{Az}{z - 1/\gamma} + \frac{Cz}{z - \gamma} \quad (\text{C.27})$$

$$= -\frac{A\gamma z}{1 - \gamma z} + \frac{C}{1 - \gamma z^*}, \quad (1/z = z^*) \quad (\text{C.28})$$

where we have chosen the latter form because Eq. (C.6) holds only for  $|a| < 1$ , so we have chosen the coefficient of a term involving  $z$  to always satisfy this constraint.

First we treat the conjugation. Note that if we define  $g(n) = f(-n)$  (the time-reversal of  $f$ ), we have

$$\hat{g}(\omega) = \sum_{n=-\infty}^{\infty} f(-n)e^{-i\omega n} \quad (\text{C.29})$$

$$= \sum_{n'=-\infty}^{\infty} f(n')e^{i\omega n'} \quad (n' = -n) \quad (\text{C.30})$$

$$= \hat{f}^*(\omega). \quad (\text{C.31})$$

Therefore the second term of Eq. (C.28) represents a backwards-filtered decaying exponential.

For the first term, we note that it is a product of  $z$  and Eq. (C.6); this indicates that this is the DFT of the convolution between a delta-function (a translation) and an exponential decay. Rather than use the convolution theorem, it's easier to compute

directly: if  $h(n) = r(n - n_0)$ , then

$$\hat{h}(\omega) = \sum_{n=-\infty}^{\infty} a^{n-n_0} u(n - n_0) e^{-i\omega n} \quad (\text{C.32})$$

$$= z^{n_0} \sum_{n=-\infty}^{\infty} a^{n-n_0} u(n - n_0) z^{n-n_0} \quad (e^{-i\omega n} \rightarrow z^n) \quad (\text{C.33})$$

$$= \frac{z^{n_0}}{1 - az}. \quad (\text{C.34})$$

Consequently, Eq. (C.28) can be recognized as the DFT of

$$w(n) = -A\gamma r_\gamma(n - 1) + Cr_\gamma(-n). \quad (\text{C.35})$$

Let's check this in the limit of low noise  $\eta^2 \ll 1$ , for which we have  $b \approx -1$ ,  $c \ll 1$ , and consequently  $\gamma \approx c = a\eta^2$ . So in this limit  $|\gamma| \ll 1$ , and the filters in Eq. (C.35) decay very rapidly, nearly becoming delta-functions. Since  $A \rightarrow 1/\eta^2$  and  $C \rightarrow 1$ , Eq. (C.35) becomes approximately

$$w(n) \approx -A\gamma\delta(n - 1) + C\delta(n) \quad (\text{C.36})$$

$$= -a\delta(n - 1) + \delta(n). \quad (\text{C.37})$$

This is the filter version of Eq. (C.5), so it checks out.

Thus, the final recipe is:

- Given  $a$  and  $\eta^2$ , compute  $\gamma$  from Eqs. (C.10) and Eq. (C.19)



- Compute the forward- and reverse-filtered signals  $s_+$  and  $s_-$ :

$$s_+(n) = m(n) + \gamma s_+(n-1); \quad (\text{starting from the beginning}) \quad (\text{C.38})$$

$$s_-(n) = m(n) + \gamma s_-(n+1). \quad (\text{starting from the end}) \quad (\text{C.39})$$

Near the edges these have “issues,” so don’t take the edges too seriously. (These issues decay with time constant  $\gamma$ .)

- Compute  $A$  and  $C$  from Eqs. (C.26)
- Compute

$$s(n) = -A\gamma s_+(n-1) + C s_-(n). \quad (\text{C.40})$$

# Appendix D

## Sparsity of the weight matrix with linear regression

Consider a dataset consisting of three vectors  $a$ ,  $b$ , and  $c$ . Suppose that these vectors are temporal traces of the activity of three neurons, and that one is interested in fitting a linear model to predict  $c$  based on samples from  $a$  and  $d$  at a previous timepoint. Fitting the model consists of solving the following equation for predictive weights  $w_a$  and  $w_b$  for a particular time lag  $\Delta t$

$$c_t = w_a a_{t-\Delta t} + w_b b_{t-\Delta t} \tag{D.1}$$

Typically no  $w_a$  and  $w_b$  exist to solve this equation exactly, so instead we can find the  $w_a$  and  $w_b$  that minimize the sum of squared prediction errors

$$\min_{w_a, w_b} \sum_t (c_t - (w_a a_{t-\Delta t} + w_b b_{t-\Delta t}))^2 \quad (\text{D.2})$$

Equivalently we can minimize the mean squared error, which is simply the above quantity divided by the number of timepoints. We will present an argument that this optimization naturally leads to sparsity in weights, often driving either  $w_a$  or  $w_b$  to zero even in situations where the time-shifted  $a$  and  $b$  vectors have similar correlations with  $c$ .

First we restate the mean squared error optimization problem in statistical language

$$\min_{w_a, w_b} \text{E}[(C - (w_a A_{shift} + w_b B_{shift}))^2] \quad (\text{D.3})$$

casting  $A$ ,  $B$ ,  $C$  and their time-shifted versions as random variables.

Suppose  $B_{shift}$  is partially redundant with  $A_{shift}$ , differing only by the addition of another signal  $S$

$$B_{shift} = A_{shift} + S \quad (\text{D.4})$$

Substituting into Equation D.3

$$\begin{aligned} & \min_{w_a, w_b} \text{E}[(C - (w_a A_{shift} + w_b (A_{shift} + S)))^2] \\ &= \min_{w_a, w_b} \text{E}[(C - ((w_a + w_b) A_{shift} + w_b S))^2] \end{aligned} \quad (\text{D.5})$$

For simplicity we will define weight  $w_{ab} = w_a + w_b$  and substitute, yielding

$$\begin{aligned}
& \min_{w_{ab}, w_b} \mathbb{E}[(C - (w_{ab}A_{shift} + w_bS))^2] \\
&= \min_{w_{ab}, w_b} \mathbb{E}[C^2 + w_{ab}^2 A_{shift}^2 + w_b^2 S^2 - 2w_{ab}CA_{shift} - 2w_bCS + 2w_{ab}w_bA_{shift}S] \\
&= \min_{w_{ab}, w_b} \mathbb{E}[C^2] + \mathbb{E}[w_{ab}^2 A_{shift}^2] + \mathbb{E}[w_b^2 S^2] - \\
& \quad \mathbb{E}[2w_{ab}CA_{shift}] - \mathbb{E}[2w_bCS] + \mathbb{E}[2w_{ab}w_bA_{shift}S]
\end{aligned} \tag{D.6}$$

We seek the optimal  $w_b$  in this situation. We can solve for this by setting the derivative of the objective function w.r.t.  $w_b$  to zero and solving for  $w_b$ .

$$\begin{aligned}
0 &= \frac{d}{dw_b} (\mathbb{E}[w_b^2 S^2] - \mathbb{E}[2w_bCS] + \mathbb{E}[2w_{ab}w_bA_{shift}S]) \\
&= \frac{d}{dw_b} (w_b^2 \mathbb{E}[S^2] - 2w_b \mathbb{E}[CS] + 2w_{ab}w_b \mathbb{E}[A_{shift}S]) \\
&= 2w_b \mathbb{E}[S^2] - 2\mathbb{E}[CS] + 2w_{ab} \mathbb{E}[A_{shift}S] \\
w_b &= \frac{\mathbb{E}[CS] - w_{ab} \mathbb{E}[A_{shift}S]}{\mathbb{E}[S^2]}
\end{aligned} \tag{D.7}$$

Equation D.7 gives the value of  $w_b$  for *any*  $C$ ,  $A_{shift}$ , and  $S$ . We next consider the scenario when  $S$  is *independent* of  $C$ , contributing no information useful for prediction. This does not mean that  $S$  is “noise”, though it may include noise.  $S$  may also be independent if it consists of neural activity unrelated to activity in  $C$ , such as might be expected when activity is spontaneous (self-generated) or initiated through a different neural pathway. In addition to this independence let us assume that  $\mathbb{E}[C] = 0$  and  $\mathbb{E}[A] = 0$ . This is equivalent to preprocessing  $C$  and  $A$  by subtracting their means,

akin to subtracting baseline fluorescence when calcium imaging.

$$\begin{aligned}
 w_b &= \frac{\mathbb{E}[C]E[S] - w_{ab}\mathbb{E}[A_{shift}S]}{\mathbb{E}[S^2]} \quad (\text{by independence}) \\
 &= \frac{-w_{ab}\text{Cov}[A_{shift}, S]}{\mathbb{E}[S^2]} \quad (\text{because } \mathbb{E}[C] = 0 \text{ and } \mathbb{E}[A] = 0)
 \end{aligned} \tag{D.8}$$

Thus under these conditions the optimal  $w_b$  is independent of  $C$  and is nonzero only when  $A_{shift}$  and  $S$  have nonzero covariance. We can also see that  $w_b$  decreases in proportion to the average power in  $S$ . If, in addition,  $S$  and  $A_{shift}$  are also independent then  $\text{Cov}[A_{shift}, S] = 0$  and therefore  $w_b$  must equal zero. The crucial point is that under these same conditions the correlations of  $A_{shift}$  with  $C$  and  $B_{shift}$  with  $C$  are both nonzero, and indeed may be almost equal.

While these independence conditions may not be met precisely in real neuronal recordings, we see that the regression approach declines to weight relationships that provide redundant predictive information.

The above argument extends easily to regression with an arbitrary number of variables (neurons), leading to sparsity in the weight matrix describing their interactions.

# References

- [1] E. Abbe. “Beiträge zur Theorie des Mikroskops und der mikroskopischen Wahrnehmung”. In: *Archiv für Mikroskopische Anatomie* 9.1 (1873), pp. 413–418. DOI: 10.1007/bf02956173.
- [2] S.S. Aphale, B. Bhikkaji, and S.O.R. Moheimani. “Minimizing Scanning Errors in Piezoelectric Stack-Actuated Nanopositioning Platforms”. In: *IEEE Transactions on Nanotechnology* 7.1 (Jan. 1, 2008), pp. 79–90. DOI: 10.1109/tnano.2007.910333.
- [3] Rembrandt Bakker, Thomas Wachtler, and Markus Diesmann. “CoCoMac 2.0 and the future of tract-tracing databases”. In: *Frontiers in Neuroinformatics* 6 (2012). DOI: 10.3389/fninf.2012.00030.
- [4] Cornelia I. Bargmann. “Beyond the connectome: How neuromodulators shape neural circuits”. In: *BioEssays* 34.6 (2012), pp. 458–465. DOI: 10.1002/bies.201100185.
- [5] Peter Bartho et al. “Population coding of tone stimuli in auditory cortex: dynamic rate vector analysis”. In: *European Journal of Neuroscience* 30.9 (2009), pp. 1767–1778. DOI: 10.1111/j.1460-9568.2009.06954.x.
- [6] Jeff Bezanson et al. “Julia: A Fast Dynamic Language for Technical Computing”. In: *CoRR* abs/1209.5145 (2012). arXiv: 1209.5145.
- [7] Irving J. Bigio and Stephen G. Bown. “Spectroscopic Sensing of Cancer and Cancer Therapy: Current Status of Translational Research”. In: *Cancer Biology & Therapy* 3.3 (2004), pp. 259–267. DOI: 10.4161/cbt.3.3.694.
- [8] Joel N. Bixler et al. “Assessment of tissue heating under tunable near-infrared radiation”. In: *Journal of Biomedical Optics* 19.7 (2014), p. 070501. DOI: 10.1117/1.jbo.19.7.070501.
- [9] Martin J. Booth, Delphine Débarre, and Alexander Jesacher. “Adaptive Optics for Biomedical Microscopy”. In: *Optics and Photonics News* 23.1 (2012), p. 22. DOI: 10.1364/opn.23.1.000022.
- [10] E.J. Botcherby et al. “An optical technique for remote focusing in microscopy”. In: *Optics Communications* 281.4 (2008), pp. 880–887. DOI: 10.1016/j.optcom.2007.10.007.

- [11] Matthew B. Bouchard et al. “Swept confocally-aligned planar excitation (SCAPE) microscopy for high-speed volumetric imaging of behaving organisms”. In: *Nature Photonics* 9.2 (Jan. 19, 2015), nphoton.2014.323. ISSN: 1749-4893. DOI: 10.1038/nphoton.2014.323.
- [12] Michael Broxton et al. “Wave optics theory and 3-D deconvolution for the light field microscope”. In: *Optics Express* 21.21 (2013), p. 25418. DOI: 10.1364/oe.21.025418.
- [13] E. Kelly Buchanan et al. “Penalized matrix decomposition for denoising, compression, and improved demixing of functional imaging data”. In: (2018). DOI: 10.1101/334706.
- [14] Ramón y Cajal and S Gac Méd Cat. “Sobre la existencia de células nerviosas especiales en la primera capa de las circonvoluciones cerebrales”. In: vol. 13. 1890, pp. 737–739.
- [15] N. Callamaras and I. Parker. “Construction of a confocal microscope for real-time x-y and x-z imaging”. In: *Cell Calcium* 26.6 (1999), pp. 271–279. DOI: 10.1054/ceca.1999.0085.
- [16] Luis Carrillo-Reid et al. “Encoding Network States by Striatal Cell Assemblies”. In: *Journal of Neurophysiology* 99.3 (2008), pp. 1435–1450. DOI: 10.1152/jn.01131.2007.
- [17] Tsai-Wen Chen et al. “Ultrasensitive fluorescent proteins for imaging neuronal activity”. In: *Nature* 499.7458 (July 18, 2013), pp. 295–300. ISSN: 0028-0836. DOI: 10.1038/nature12354.
- [18] John P Cunningham and Byron M Yu. “Dimensionality reduction for large-scale neural recordings”. In: *Nature Neuroscience* 17.11 (2014), pp. 1500–1509. DOI: 10.1038/nn.3776.
- [19] Lleres David et al. “Detecting Protein-Protein Interactions In Vivo with FRET using Multiphoton Fluorescence Lifetime Imaging Microscopy (FLIM). Current protocols in cytometry”. In: editor:editorial board, J. Paul Robinson, managing editor. 2007.
- [20] Kevin M. Dean and Reto Fiolka. “Lossless Three-Dimensional Parallelization in Digitally Scanned Light-Sheet Fluorescence Microscopy”. In: *Scientific Reports* 7.1 (Aug. 24, 2017). DOI: 10.1038/s41598-017-08113-8.
- [21] Kevin M. Dean and Reto Fiolka. “Uniform and scalable light-sheets generated by extended focusing”. In: *Optics Express* 22.21 (2014), p. 26141. DOI: 10.1364/oe.22.026141.
- [22] Karl Deisseroth. “Optogenetics”. In: *Nature Methods* 8.1 (2010), pp. 26–29. DOI: 10.1038/nmeth.f.324.

- [23] Ram Dixit and Richard Cyr. “Cell damage and reactive oxygen species production induced by fluorescence microscopy: effect on mitosis and guidelines for non-invasive fluorescence microscopy”. In: *The Plant Journal* 36.2 (2003), pp. 280–290. DOI: 10.1046/j.1365-313x.2003.01868.x.
- [24] Timothy W Dunn et al. “Brain-wide mapping of neural activity controlling zebrafish exploratory locomotion”. In: *eLife* 5 (2016-03-22). DOI: 10.7554/elife.12741.
- [25] C. Dunsby. “Optically sectioned imaging by oblique plane microscopy”. In: *Optics Express* 16.25 (Dec. 8, 2008), pp. 20306–20316. ISSN: 1094-4087. DOI: 10.1364/OE.16.020306.
- [26] Martí Duocastella et al. “Fast Inertia-Free Volumetric Light-Sheet Microscope”. In: *ACS Photonics* 4.7 (2017), pp. 1797–1804. DOI: 10.1021/acsp Photonics.7b00382. eprint: <http://dx.doi.org/10.1021/acsp Photonics.7b00382>.
- [27] David C. Van Essen et al. “The WU-Minn Human Connectome Project: An overview”. In: *NeuroImage* 80 (2013), pp. 62–79. DOI: 10.1016/j.neuroimage.2013.05.041.
- [28] F. O. Fahrbach et al. “Rapid 3D light-sheet microscopy with a tunable lens”. In: *Optics Express* 21.18 (Sept. 9, 2013), pp. 21010–21026. ISSN: 1094-4087. DOI: info:doi/10.1364/OE.21.021010.
- [29] Jeremy Freeman et al. “Mapping brain activity at scale with cluster computing”. In: *Nature Methods* 11.9 (2014), pp. 941–950. DOI: 10.1038/nmeth.3041.
- [30] Johannes Friedrich, Pengcheng Zhou, and Liam Paninski. “Fast online deconvolution of calcium imaging data”. In: *PLOS Computational Biology* 13.3 (2017). Ed. by Joshua Vogelstein, e1005423. DOI: 10.1371/journal.pcbi.1005423.
- [31] Andrea Giovannucci et al. “OnACID: Online Analysis of Calcium Imaging Data in Real Time”. In: (2017). DOI: 10.1101/193383.
- [32] Yiyang Gong et al. “High-speed recording of neural spikes in awake mice and flies with a fluorescent voltage sensor”. In: *Science* (Nov. 19, 2015), aab0810. ISSN: 0036-8075, 1095-9203. DOI: 10.1126/science.aab0810.
- [33] Grynkieicz Grzegorz, Poenie Martin, and Tsien Roger J Biol Chem. “A new generation of calcium indicators with greatly improved fluorescence properties”. In: 1985, pp. 260–3340.
- [34] Suzanaerculano-Houzel. “The human brain in numbers: a linearly scaled-up primate brain”. In: *Frontiers in Human Neuroscience* 3 (2009). DOI: 10.3389/neuro.09.031.2009.
- [35] J. F. W. Herschel. “On the Aberrations of Compound Lenses and Object-Glasses”. In: *Philosophical Transactions of the Royal Society of London* 111.0 (1821), pp. 222–267. DOI: 10.1098/rstl.1821.0018.



- [36] Terrence F. Holekamp, Diwakar Turaga, and Timothy E. Holy. “Fast Three-Dimensional Fluorescence Imaging of Activity in Neural Populations by Objective-Coupled Planar Illumination Microscopy”. In: *Neuron* 57.5 (Mar. 13, 2008), pp. 661–672. ISSN: 0896-6273. DOI: 10.1016/j.neuron.2008.01.011.
- [37] J. Huisken and D. Y. R. Stainier. “Selective plane illumination microscopy techniques in developmental biology”. In: *Development* 136.12 (2009), pp. 1963–1975. DOI: 10.1242/dev.022426.
- [38] Jan Huisken and Didier Y. R. Stainier. “Even fluorescence excitation by multi-directional selective plane illumination microscopy (mSPIM)”. In: *Optics Letters* 32.17 (2007), p. 2608. DOI: 10.1364/ol.32.002608.
- [39] Jan Huisken et al. “Optical Sectioning Deep Inside Live Embryos by Selective Plane Illumination Microscopy”. In: *Science* 305.5686 (Aug. 13, 2004), pp. 1007–1009. ISSN: 0036-8075, 1095-9203. DOI: 10.1126/science.1100035.
- [40] James J. Jun et al. “Fully integrated silicon probes for high-density recording of neural activity”. In: *Nature* 551.7679 (2017), pp. 232–236. DOI: 10.1038/nature24636.
- [41] P. J. Keller et al. “Reconstruction of Zebrafish Early Embryonic Development by Scanned Light Sheet Microscopy”. In: *Science* 322.5904 (2008), pp. 1065–1069. DOI: 10.1126/science.1162493.
- [42] Philipp J. Keller et al. “Fast, high-contrast imaging of animal development with scanned light sheet-based structured-illumination microscopy”. In: *Nature Methods* 7.8 (Aug. 2010), pp. 637–642. ISSN: 1548-7091. DOI: 10.1038/nmeth.1476.
- [43] Christina K. Kim et al. “Prolonged, brain-wide expression of nuclear-localized GCaMP3 for functional circuit mapping”. In: *Frontiers in Neural Circuits* 8 (2014). DOI: 10.3389/fncir.2014.00138.
- [44] Uros Krzic et al. “Multiview light-sheet microscope for rapid in toto imaging”. In: *Nature Methods* 9.7 (July 2012), pp. 730–733. ISSN: 1548-7091. DOI: 10.1038/nmeth.2064.
- [45] Manish Kumar et al. “Integrated one- and two-photon scanned oblique plane illumination (SOPi) microscopy for rapid volumetric imaging”. In: *Optics Express* 26.10 (2018), p. 13027. DOI: 10.1364/oe.26.013027.
- [46] P. Philippe Laissue et al. “Assessing phototoxicity in live fluorescence imaging”. In: *Nature Methods* 14.7 (July 2017), pp. 657–661. ISSN: 1548-7091. DOI: 10.1038/nmeth.4344.
- [47] Jack L. Lancaster et al. “Automated Talairach Atlas labels for functional brain mapping”. In: *Human Brain Mapping* 10.3 (2000), pp. 120–131. DOI: 10.1002/1097-0193(200007)10:3<120::aid-hbm30>3.0.co;2-8.

- [48] Marc Levoy et al. “Light Field Microscopy”. In: *ACM SIGGRAPH 2006 Papers*. SIGGRAPH '06. New York, NY, USA: ACM, 2006, pp. 924–934. ISBN: 978-1-59593-364-5. DOI: 10.1145/1179352.1141976.
- [49] Daryl Lim et al. “Optically sectioned *in vivo* imaging with speckle illumination HiLo microscopy”. In: *Journal of Biomedical Optics* 16.1 (Jan. 2011), p. 016014. ISSN: 1083-3668, 1560-2281. DOI: 10.1117/1.3528656.
- [50] L. Lin et al. “Identification of network-level coding units for real-time representation of episodic experiences in the hippocampus”. In: *Proceedings of the National Academy of Sciences* 102.17 (2005), pp. 6125–6130. DOI: 10.1073/pnas.0408233102.
- [51] Elisa De Luca et al. “ZebraBeat: a flexible platform for the analysis of the cardiac rate in zebrafish embryos”. In: *Scientific Reports* 4.1 (2014). DOI: 10.1038/srep04898.
- [52] Werner Manz et al. “Widefield deconvolution epifluorescence microscopy combined with fluorescence in situ hybridization reveals the spatial arrangement of bacteria in sponge tissue”. In: *Journal of Microbiological Methods* 40.2 (2000), pp. 125–134. DOI: 10.1016/S0167-7012(99)00103-7.
- [53] Gregory D. Marquart et al. “A 3D Searchable Database of Transgenic Zebrafish Gal4 and Cre Lines for Functional Neuroanatomy Studies”. In: *Frontiers in Neural Circuits* 9 (2015). DOI: 10.3389/fncir.2015.00078.
- [54] Gregory D. Marquart et al. “High-precision registration between zebrafish brain atlases using symmetric diffeomorphic normalization”. In: *GigaScience* 6.8 (2017). DOI: 10.1093/gigascience/gix056.
- [55] Ryuichi Maruyama et al. “Detecting cells using non-negative matrix factorization on calcium imaging data”. In: *Neural Networks* 55 (2014), pp. 11–19. DOI: 10.1016/j.neunet.2014.03.007.
- [56] Jonathan S Marvin et al. “An optimized fluorescent probe for visualizing glutamate neurotransmission”. In: *Nature Methods* 10.2 (2013), pp. 162–170. DOI: 10.1038/nmeth.2333.
- [57] James Clerk Maxwell. “On the General Laws of Optical Instruments”. In: *Quarterly Journal of Pure & Applied Mathematics* 2 (1858), pp. 233–246.
- [58] Eran A. Mukamel, Axel Nimmerjahn, and Mark J. Schnitzer. “Automated Analysis of Cellular Signals from Large-Scale Calcium Imaging Data”. In: *Neuron* 63.6 (2009), pp. 747–760. DOI: 10.1016/j.neuron.2009.08.009.
- [59] Nandakumar S. Narayanan and Mark Laubach. “Delay Activity in Rodent Frontal Cortex During a Simple Reaction Time Task”. In: *Journal of Neurophysiology* 101.6 (2009), pp. 2859–2871. DOI: 10.1152/jn.90615.2008.
- [60] M.A.A. Neil, R. Juškaitis, and T. Wilson. “Real time 3D fluorescence microscopy by two beam interference illumination”. In: *Optics Communications* 153.1-3 (1998), pp. 1–4. DOI: 10.1016/S0030-4018(98)00210-7.

- [61] Hamid R. Noori et al. “A multiscale cerebral neurochemical connectome of the rat brain”. In: *PLOS Biology* 15.7 (2017). Ed. by Hongkui Zeng, e2002612. DOI: 10.1371/journal.pbio.2002612.
- [62] Seung Wook Oh et al. “A mesoscale connectome of the mouse brain”. In: *Nature* 508.7495 (2014), pp. 207–214. DOI: 10.1038/nature13186.
- [63] Omar E. Olarte et al. “Decoupled illumination detection in light sheet microscopy for fast volumetric imaging”. In: *Optica* 2.8 (Aug. 20, 2015), pp. 702–705. ISSN: 2334-2536. DOI: 10.1364/OPTICA.2.000702.
- [64] Omar E. Olarte et al. “Light-sheet microscopy: a tutorial”. In: *Advances in Optics and Photonics* 10.1 (2018), p. 111. DOI: 10.1364/aop.10.000111.
- [65] Thomas Panier et al. “Fast functional imaging of multiple brain regions in intact zebrafish larvae using Selective Plane Illumination Microscopy”. In: *Frontiers in Neural Circuits* 7 (2013). DOI: 10.3389/fncir.2013.00065.
- [66] J. B. Pawley. “Fundamental and practical limits in confocal light microscopy”. In: *Scanning* 13.2 (1991), pp. 184–198. ISSN: 0161-0457. DOI: 10.1002/sca.4950130205.
- [67] Jelena Platisa and Vincent A Pieribone. “Genetically encoded fluorescent voltage indicators: are we there yet?” In: *Current Opinion in Neurobiology* 50 (2018), pp. 146–153. DOI: 10.1016/j.conb.2018.02.006.
- [68] Eftychios A. Pnevmatikakis and Andrea Giovannucci. “NoRMCorre: An online algorithm for piecewise rigid motion correction of calcium imaging data”. In: *Journal of Neuroscience Methods* 291 (2017), pp. 83–94. DOI: 10.1016/j.jneumeth.2017.07.031.
- [69] Eftychios A. Pnevmatikakis et al. “Simultaneous Denoising, Deconvolution, and Demixing of Calcium Imaging Data”. In: *Neuron* 89.2 (2016), pp. 285–299. DOI: 10.1016/j.neuron.2015.11.037.
- [70] Jonathan D. Power et al. “Methods to detect, characterize, and remove motion artifact in resting state fMRI”. In: *NeuroImage* 84 (2014), pp. 320–341. DOI: 10.1016/j.neuroimage.2013.08.048.
- [71] Robert Prevedel et al. “Simultaneous whole-animal 3D imaging of neuronal activity using light-field microscopy”. In: *Nature Methods* 11.7 (July 2014), pp. 727–730. ISSN: 1548-7091. DOI: 10.1038/nmeth.2964.
- [72] Sean Quirin et al. “Calcium imaging of neural circuits with extended depth-of-field light-sheet microscopy”. In: *Optics Letters* 41.5 (Mar. 1, 2016), pp. 855–858. ISSN: 1539-4794. DOI: 10.1364/OL.41.000855.
- [73] Owen Randlett et al. “Whole-brain activity mapping onto a zebrafish brain atlas”. In: *Nature Methods* 12.11 (2015), pp. 1039–1046. DOI: 10.1038/nmeth.3581.

- [74] Nicolas Renier et al. “iDISCO: A Simple, Rapid Method to Immunolabel Large Tissue Samples for Volume Imaging”. In: *Cell* 159.4 (Nov. 6, 2014), pp. 896–910. ISSN: 0092-8674. DOI: 10.1016/j.cell.2014.10.010.
- [75] Mike Richardson. “Deconvolution of Images and Spectra, 2nd Edition”. In: *Optical Engineering* 36.11 (1997), p. 3224. DOI: 10.1117/1.601524.
- [76] Olaf Ronneberger et al. “ViBE-Z: a framework for 3D virtual colocalization analysis in zebrafish larval brains”. In: *Nature Methods* 9.7 (2012), pp. 735–742. DOI: 10.1038/nmeth.2076.
- [77] Peter Rupprecht et al. “Remote z-scanning with a macroscopic voice coil motor for fast 3D multiphoton laser scanning microscopy”. In: *Biomedical Optics Express* 7.5 (2016), p. 1656. DOI: 10.1364/boe.7.001656.
- [78] Kaz Sato, Cliff Young, and David Patterson. “An in-depth look at Google’s first Tensor Processing Unit (TPU)”. In: *Google Cloud Big Data and Machine Learning Blog* 12 (2017).
- [79] Manish Saxena, Gangadhar Eluru, and Sai Siva Gorthi. “Structured illumination microscopy”. In: *Advances in Optics and Photonics* 7.2 (2015), p. 241. DOI: 10.1364/aop.7.000241.
- [80] Benjamin Schmid and Jan Huisken. “Real-time multi-view deconvolution”. In: *Bioinformatics* 31.20 (Oct. 15, 2015), pp. 3398–3400. ISSN: 1367-4803. DOI: 10.1093/bioinformatics/btv387.
- [81] H. Sebastian Seung. “Towards functional connectomics”. In: *Nature* 471.7337 (2011), pp. 171–172. DOI: 10.1038/471170a.
- [82] Markus B. Sikkell et al. “High speed sCMOS-based oblique plane microscopy applied to the study of calcium dynamics in cardiac myocytes”. In: *Journal of Biophotonics* 9.3 (Oct. 21, 2015), pp. 311–323. DOI: 10.1002/jbio.201500193.
- [83] Diana Smetters, Ania Majewska, and Rafael Yuste. “Detecting Action Potentials in Neuronal Populations with Calcium Imaging”. In: *Methods* 18.2 (1999), pp. 215–221. DOI: 10.1006/meth.1999.0774.
- [84] François St-Pierre et al. “High-fidelity optical reporting of neuronal electrical activity with an ultrafast fluorescent voltage sensor”. In: *Nature Neuroscience* 17.6 (Apr. 22, 2014), nn.3709. ISSN: 1546-1726. DOI: 10.1038/nn.3709.
- [85] Johannes Stegmaier et al. “Real-Time Three-Dimensional Cell Segmentation in Large-Scale Microscopy Data of Developing Embryos”. In: *Developmental Cell* 36.2 (Jan. 25, 2016), pp. 225–240. ISSN: 1534-5807. DOI: 10.1016/j.devcel.2015.12.028.
- [86] Yong Tang et al. “Total regional and global number of synapses in the human brain neocortex”. In: *Synapse* 41.3 (2001), pp. 258–273. DOI: 10.1002/syn.1083.

- [87] Michael A. Taylor et al. “Diffuse light-sheet microscopy for stripe-free calcium imaging of neural populations”. In: *Journal of Biophotonics* (2018), e201800088. DOI: 10.1002/jbio.201800088.
- [88] Patrick Theer and Winfried Denk. “On the fundamental imaging-depth limit in two-photon microscopy”. In: *Journal of the Optical Society of America A* 23.12 (Dec. 1, 2006), p. 3139. DOI: 10.1364/josaa.23.003139.
- [89] Raju Tomer et al. “Advanced CLARITY for rapid and high-resolution imaging of intact tissues”. In: *Nature Protocols* 9.7 (July 2014), pp. 1682–1697. ISSN: 1754-2189. DOI: 10.1038/nprot.2014.123.
- [90] Thai V. Truong et al. “Deep and fast live imaging with two-photon scanned light-sheet microscopy”. In: *Nature Methods* 8.9 (Sept. 2011), pp. 757–760. ISSN: 1548-7091. DOI: 10.1038/nmeth.1652.
- [91] Fuan Tsai and Walter W. Chen. “Striping Noise Detection and Correction of Remote Sensing Images”. In: *IEEE Transactions on Geoscience and Remote Sensing* 46.12 (2008), pp. 4122–4131. DOI: 10.1109/tgrs.2008.2000646.
- [92] Jérôme Tubiana, Sébastien Wolf, and Georges Debregeas. “Blind sparse deconvolution for inferring spike trains from fluorescence recordings”. In: (2017). DOI: 10.1101/156364.
- [93] Diwakar Turaga and Timothy E. Holy. “Miniaturization and defocus correction for objective-coupled planar illumination microscopy”. In: *Optics Letters* 33.20 (Oct. 15, 2008), pp. 2302–2304. ISSN: 1539-4794. DOI: 10.1364/OL.33.002302.
- [94] Lav R. Varshney et al. “Structural Properties of the Caenorhabditis elegans Neuronal Network”. In: *PLoS Computational Biology* 7.2 (2011). Ed. by Olaf Sporns, e1001066. DOI: 10.1371/journal.pcbi.1001066.
- [95] Joshua T. Vogelstein et al. “Spike Inference from Calcium Imaging Using Sequential Monte Carlo Methods”. In: *Biophysical Journal* 97.2 (2009), pp. 636–655. DOI: 10.1016/j.bpj.2008.08.005.
- [96] J. G. White et al. “The Structure of the Nervous System of the Nematode *Caenorhabditis elegans*”. In: *Philosophical Transactions of the Royal Society B: Biological Sciences* 314.1165 (1986), pp. 1–340. DOI: 10.1098/rstb.1986.0056.
- [97] Richard Mark White et al. “Transparent adult zebrafish as a tool for in vivo transplantation analysis”. In: *Cell stem cell* 2.2 (Feb. 7, 2008), pp. 183–189. ISSN: 1934-5909. DOI: 10.1016/j.stem.2007.11.002.
- [98] Norbert Wiener et al. “Extrapolation, interpolation, and smoothing of stationary time series: with engineering applications”. In: (1949).
- [99] Bin Yang et al. “Single-Cell Phenotyping within Transparent Intact Tissue through Whole-Body Clearing”. In: *Cell* 158.4 (Aug. 14, 2014), pp. 945–958. ISSN: 0092-8674. DOI: 10.1016/j.cell.2014.07.017.

- [100] Hang Yu et al. “Two-photon Swept Confocally Aligned Planar Excitation Microscopy (2P-SCAPE)”. In: *Optics in the Life Sciences Congress*. OSA, 2017. DOI: 10.1364/ntm.2017.nw4c.3.
- [101] Pengcheng Zhou et al. “Efficient and accurate extraction of in vivo calcium signals from microendoscopic video data”. In: *eLife* 7 (2018). DOI: 10.7554/elife.28728.

**Fast Objective Coupled Planar Illumination Microscopy, Greer, Ph.D. 2018**