



**A University of Sussex PhD thesis**

Available online via Sussex Research Online:

<http://sro.sussex.ac.uk/>

This thesis is protected by copyright which belongs to the author.

This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the Author

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the Author

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given

Please visit Sussex Research Online for more information and further details



**Toward an experimental and  
computational approach to causal  
analysis in behaving zebrafish larvae**

**Daniel Saska**

Submitted for the degree of PhD in Neuroscience

University of Sussex

# Declaration

I hereby declare that this thesis has not been and will not be submitted in whole or in part to another University for the award of any other degree.

**Chapter 3** has been published in *Journal of Neuroscience Methods* with following contributions: **Daniel Saska:** Software, Investigation, Writing – Original Draft, Visualization. **Paul Pichler:** Investigation, Writing – Review and Editing. **Chen Qian:** Investigation. **Christopher Buckley:** Conceptualization, Software, Writing – Review and Editing, Project Administration, Funding Acquisition. **Leon Lagnado:** Conceptualization, Methodology, Software, Visualization, Writing – Review and Editing, Project Administration, Funding Acquisition.

In the remaining chapters the contributions are as follows: **Daniel Saska** was responsible for all aspects of data collection, data analysis, writing of the manuscript and in part for preparation of animal samples for data acquisition; **Paul Pichler** was responsible animal husbandry, preparation of samples for acquisition and in chapter 4 for integration of the presented VR setup into existing two-photon microscope setup. **Christopher Buckley and Leon Lagnado** were responsible for providing feedback and aid in the conceptual design of the presented studies. **Lionel Barnett** provided guidance on the application of Granger Causality and provided aid in the conceptual design for Chapters 5 and 6.

Signature:

Daniel Saska

UNIVERSITY OF SUSSEX

DANIEL SASKA

TOWARD AN EXPERIMENTAL AND COMPUTATIONAL APPROACH TO  
CAUSAL ANALYSIS IN BEHAVING ZEBRAFISH LARVAE

SUMMARY

Understanding brain-wide dynamics and their relation to behaviour relies on knowledge of the interactions of the underlying functional regions in the brain. In this work, we aim to demonstrate the applicability and limitations of Granger Causality (GC) as a measure of directed functional connectivity in live zebrafish larvae, offering an alternative to commonly used undirected functional connectivity measures such as correlation. In order to acquire whole-brain datasets, we develop  $\mu$ SPIM: a hardware-agnostic light-sheet microscope control and acquisition toolset which provides functionality focused on functional imaging, providing an open-source alternative to existing light-sheet solutions limited to developmental imaging. Further, we present an independent closed-loop virtual reality solution which provides a flexible extension to existing light-sheet or two photon microscope setups.

In order to demonstrate the applicability of GC to calcium imaging data, we first apply the causal analysis to simulated spiking data generated by an integrate-and-fire model convolved with a calcium filter. We show that the directed functional connectivity reconstructed by GC follows the structural connectivity used to simulate the underlying network both in bi-variate and multi-variate settings. We identify a number of constraints on the performance of the measure in form of sampling rate, recording duration and the number of cells in the network and show that trends in the calcium data result in poor inference which can be mitigated by filtering prior to the application of GC. Next, we show that conditional GC on subsets of neurons can be used to infer directed connectivity between functionally similar neuronal circuits when analysis based on all sources is not viable due to combinatorial and computational constraints. Finally, we show that directed connectivity inferred using GC from calcium data collected in vivo from unstimulated zebrafish larvae displays functional characteristics described in prior research.

# Contents

<b>List of Tables</b>	<b>viii</b>
<b>List of Figures</b>	<b>1</b>
<b>1 General Introduction</b>	<b>1</b>
1.1 Thesis outline and contributions . . . . .	3
<b>2 Background</b>	<b>5</b>
2.1 The Zebrafish Model . . . . .	5
2.2 Acquisition Methods . . . . .	6
2.2.1 Two-photon Microscopy . . . . .	7
2.2.2 Selective Light-Sheet Microscopy . . . . .	7
2.3 Virtual Reality and Behavioural Acquisition . . . . .	8
2.4 Data Processing Methods . . . . .	9
2.5 Present approaches to data analysis in calcium imaging . . . . .	9
2.5.1 Atlas Registration of Calcium Recordings . . . . .	9
2.5.2 Data dimensionality reduction . . . . .	10
2.5.3 Undirected functional connectivity measures . . . . .	10
2.6 Directed functional connectivity measures . . . . .	10
2.6.1 Time-shifted Cross-correlation . . . . .	10
2.6.2 Granger Causality . . . . .	11
2.6.3 Transfer Entropy . . . . .	11
2.6.4 Application of Granger Causality to Calcium Data . . . . .	11
2.6.5 Known limitations of Granger Causality . . . . .	12
<b>3 <math>\mu</math>SPIM Toolset: A software platform for selective plane illumination microscopy</b>	<b>13</b>
3.1 Abstract . . . . .	13

3.2	Introduction . . . . .	14
3.3	Design requirements . . . . .	15
3.4	Principles of operation . . . . .	18
3.5	Computing requirements . . . . .	20
3.6	Calibration . . . . .	22
3.6.1	X-Mirror . . . . .	22
3.6.2	Piezoelectric stage . . . . .	24
3.6.3	Z-Mirror . . . . .	24
3.6.4	X-mirror lag compensation . . . . .	24
3.7	Data acquisition using $\mu$ SPIM Toolset . . . . .	24
3.7.1	Advanced acquisition methods . . . . .	26
3.7.2	Data output format . . . . .	27
3.8	Discussion . . . . .	27
3.9	Materials and Methods . . . . .	30
3.9.1	Optical design of SPIM used in this work . . . . .	30
3.9.2	Characterising lateral and axial resolution of the microscope . . . . .	31
3.9.3	Imaging Brains of Larval Zebrafish . . . . .	31
<b>4</b>	<b>A model-based closed loop virtual reality system based on free-swimming parameters</b> . . . . .	<b>34</b>
4.1	Introduction . . . . .	34
4.2	Methods . . . . .	35
4.2.1	Tail Camera Acquisition . . . . .	37
4.2.2	Tail Tracking . . . . .	37
4.2.3	Swimming Model . . . . .	37
4.2.4	Stimulus Projection . . . . .	40
4.3	Results . . . . .	40
4.3.1	Realistic Swim Model . . . . .	40
4.3.2	Tracking of kinetic variables . . . . .	43
4.3.3	Alignment with other acquisition modalities . . . . .	43
4.3.4	Tail Tracking achieves real-time performance at high acquisition rates . . . . .	44
4.3.5	Support for open and closed loop Optomotor Response (OMR) stimulus . . . . .	44
4.3.6	Highly modular experiment design . . . . .	45
4.3.7	Native multi-threading support . . . . .	46

4.3.8	Tail-movement-based VR can be used to study behaviours in head-fixed zebrafish larvae . . . . .	48
4.4	Discussion . . . . .	50
<b>5</b>	<b>Estimating Functional Connectivity from Calcium Activity <i>in silico</i></b>	<b>54</b>
5.1	Introduction . . . . .	54
5.1.1	Granger Causality . . . . .	55
5.2	Feedback model <i>in silico</i> . . . . .	58
5.2.1	Methods . . . . .	58
5.2.2	Results . . . . .	60
5.2.3	Discussion . . . . .	65
5.3	Multi-cell networks . . . . .	67
5.3.1	Methods . . . . .	67
5.3.2	Results . . . . .	72
5.3.3	Discussion . . . . .	74
5.4	Scalable GC analysis methods . . . . .	77
5.4.1	Methods . . . . .	77
5.4.2	Results . . . . .	80
5.4.3	Discussion . . . . .	81
5.5	Conclusion . . . . .	83
5.5.1	Recording requirements and considerations . . . . .	83
5.5.2	Processing requirements and considerations . . . . .	84
5.5.3	Relation to existing applications of GC <i>in vivo</i> . . . . .	84
5.5.4	Limitations and future directions . . . . .	86
<b>6</b>	<b>Application of Granger Causality to <i>in vivo</i> Calcium Fluorescence Data</b>	<b>87</b>
6.1	Introduction . . . . .	87
6.2	Methods . . . . .	88
6.2.1	Acquisition of light-sheet data . . . . .	88
6.2.2	Motion correction and cell identification . . . . .	90
6.2.3	Automated scheduling of processing tasks . . . . .	90
6.2.4	Atlas registration . . . . .	93
6.2.5	Brain region selection . . . . .	95
6.2.6	Functional measure inference . . . . .	95
6.3	Results . . . . .	98

6.3.1	Inferred directed functional connectivity is bilaterally similar . . . . .	98
6.3.2	Granger Causality provides information different from Correlation . . . . .	99
6.3.3	Cerebellum shows functional efferents consistent with previously found structural connectivity . . . . .	99
6.3.4	Medial Octavolateralis Nucleus exhibits similar functional connectiv- ity to Cerebellum . . . . .	101
6.3.5	Efferent functional connections from the Tectum is greater than af- ferent ones . . . . .	101
6.4	Discussion . . . . .	103
<b>7</b>	<b>General Discussion</b>	<b>106</b>
7.1	New tools to study embodied behaviours using functional light-sheet mi- croscopy . . . . .	106
7.2	Validation of Granger Causality in calcium fluorescence data . . . . .	107
7.2.1	Validation in simulated neuronal network models . . . . .	108
7.2.2	Validation in <i>in vivo</i> light-sheet imaging . . . . .	109
7.3	Applicability and limitations of GC in <i>in vivo</i> calcium recordings . . . . .	110
7.3.1	Granger Causality and Two Photon microscopy . . . . .	110
7.3.2	Granger Causality on volumetric light-sheet recordings . . . . .	110
7.3.3	Acquisition setting and parameters for single plane light-sheet re- cordings . . . . .	111
7.3.4	Signal de-trending and pre-processing . . . . .	112
7.3.5	Photo-bleaching of calcium reporters . . . . .	112
7.3.6	Analysis of connectivity without conditioning . . . . .	113
7.4	Future directions . . . . .	113
7.4.1	The effects of the signal-to-noise ratio to data stationarity trade-off . . . . .	113
7.4.2	Methods to compensate for calcium reporter photo-bleaching . . . . .	113
7.4.3	Comparison of GC with other directed functional connectivity metrics . . . . .	114
7.4.4	Directed functional connectivity in behaving animals . . . . .	114
7.5	Concluding remarks . . . . .	114
	<b>Bibliography</b>	<b>116</b>

# List of Tables

3.1	Parts list for construction of SPIM . . . . .	32
5.1	Neuron parameters used for the integrate-and-fire simulation . . . . .	58
5.2	Parameters used for the simulation of the synaptic connections . . . . .	59
6.1	ANTs registration parameters . . . . .	92
6.2	MPIN atlas brain region abbreviations . . . . .	96

# List of Figures

3.1	Light-Sheet Microscope Implementation . . . . .	16
3.2	Light-Sheet Microscope Hardware Control . . . . .	17
3.3	$\mu$ SPIM Toolset User Interface . . . . .	19
3.4	Imaging larval zebrafish . . . . .	21
3.5	$\mu$ SPIM Toolset Calibration . . . . .	23
3.6	Light-Sheet Laser Masks . . . . .	25
3.7	Acquisition using $\mu$ SPIM Toolset and Output Data Format . . . . .	28
4.1	Closed Loop Virtual Reality Setup . . . . .	36
4.2	Processing Data Acquired from Freely Moving Fish . . . . .	38
4.3	Estimating Swim Model Paramters from Freely Moving Fish . . . . .	41
4.4	Virtual Reality Stimulator Data Output . . . . .	47
4.5	Bout response to gain and drift change . . . . .	50
5.1	Simulated model of two cells . . . . .	61
5.2	Performance of GC inference as a function of recording parameters . . . . .	63
5.3	GC is affected by synaptic delay in slower acquisition rates . . . . .	64
5.4	Simulated model of multiple cells . . . . .	71
5.5	Effect of observational noise in simulated network of cells . . . . .	73
5.6	High-pass filtering improves signal stationarity . . . . .	75
5.7	GC is invariant to filtering in the absence of noise . . . . .	76
5.8	Conditioning methods in large networks . . . . .	79
6.1	Acquisition and processing pipeline of light-sheet imaging data . . . . .	89
6.2	Atlas registration and neuron grouping . . . . .	94
6.3	Functional connectivity estimated <i>in vivo</i> . . . . .	97
6.4	Cross-lateral comparison of Granger Causality results . . . . .	100
6.5	Directed functional connectivity map of larval zebrafish brain . . . . .	102

# Chapter 1

## General Introduction

One of the primary goals of neuroscience is to understand how neural activity gives rise to different behaviours. Fully answering this question requires an understanding of how sensory and motor areas mediate the closed-loop flow of information between the brain, the body and the environment. The larval zebrafish is becoming a promising candidate model within which it is possible to address these questions, owing to its optical accessibility of the brain, relatively simple and available genetic manipulation compared to other species, and repertoire of robust behaviours, many of which are exhibited already only several days post fertilization ([Orger and de Polavieja, 2017](#)). Thus, this animal model provides an unique opportunity for the interrogation of behaviour, brain-wide neural activity and their relationship.

While zebrafish are equipped with visual, auditory, somatosensory and olfactory systems at this young age, behaviours driven by visual stimuli have been of particular interest. The investigated behaviours driven by vision include the optomotor response ([Ahrens et al., 2012](#)), opto-kinetic response ([Portugues et al., 2014](#); [Kubo et al., 2014](#)), escape behaviour ([Barker and Baier, 2015](#)), prey capture ([Bianco et al., 2011](#)) and photo-taxis ([Wolf et al., 2017](#)).

With the advent of selective plane imaging ([Huisken et al., 2004](#)), collection of calcium fluorescence from nearly the entire brain of larval zebrafish brain ([Ahrens et al., 2013b](#)) became available. This allows modern acquisition setups to acquire activity traces for the vast majority of all neurons in the larval brain ([Chen et al., 2018b](#)). Combined with closed loop virtual reality, this allows for identification and understanding of sensorimotor processing pathways. While the technological advancement these experimental methods provide has been extensive, significant technical challenges have prevented wide adoption of these methods.

The first set of challenges stems from the complex signal control necessary for the light-sheet microscope setup to function optimally and fast processing requirements of the closed-loop virtual reality stimulation system which needs to process data in real time. Due to the lack of published open-source solutions that would provide flexible functionality for a wide range of use scenarios, laboratories wishing to adopt these experimental methods must resort to developing custom in-house solutions specific to their problem. This not only affects the adoptability of these methods but also results in findings which are harder to reproduce as authors are not compelled to extensively document their custom solutions or utilize commercial solutions which are proprietary.

Next, the data produced by light-sheet acquisition can range from 100GB upward for every 10 minutes of recording, resulting in handling and processing of such data to be demanding both in terms of storage and processing. While the processing challenges have, to a large extent, been solved by tools such as CaImAn ([Giovannucci et al., 2018](#)) which has been purposed for processing large datasets at scale, handling such datasets is still non-trivial.

Finally, the large amount of information about the brain-wide activity during behaviour (or otherwise) can be challenging to interpret. Causal connections between brain activity and behaviour cannot be in most cases established from calcium data analysis alone, relying on additional tools such as opto-genetic stimulation or targeted cell ablation for the validation. Standard analysis methods such as measures of undirected functional connectivity (e.g. Correlation), however, provide very little information about the interactions between brain regions, offering very little guidance for establishing likely hypotheses, making validation through targeted opto-genetic stimulation or cell ablation problematic.

In this work, we aim to develop and provide flexible open source methods for functional light-sheet microscope control which can be easily adapted and extended for specific research tasks and applications. This is to complement existing solutions such as Open-SPIM ([Pitrone et al., 2013](#)) and Open SPIM microscopy ([Gualda et al., 2013](#)) which focus on providing microscope control for light-sheet systems with focus on developmental imaging at the expense of their applicability to functional light-sheet imaging. Through this, we strive to address the lack of open source functional light-sheet control software, offering a turn-key solution which provides all common functionality for future research efforts requiring functional light-sheet imaging while preserving the extensibility inherent to open-source software.

Secondly, we explore closed-loop virtual reality processing in order to develop an ac-

curate virtual reality model which can be used for the study of behaving fish. While the importance of environmental feedback has been shown (Ahrens et al., 2012), no published studies explore an optimal approach to estimating larval movement from acquired larval tail activity. We aim to address this in our work by exploring the tail dynamics in freely swimming zebrafish larvae and designing an open-source modular virtual reality stimulator solution which can be used in combination with existing 2P or light-sheet functional microscope systems.

Further, we aim to validate Granger Causality (Granger, 1969), a directed functional connectivity measure which has been widely applied to human neuroimaging data (Seth et al., 2015), on calcium fluorescence data. While a number of studies applying Granger Causality (GC) to calcium fluorescence data were conducted in the past (Fallani et al., 2014; Vanwalleggem et al., 2017; Oldfield et al., 2020), all of these studies assume validity of GC on low acquisition rate data contrary to findings in other data modalities (Seth et al., 2013). We provide a proof of principle first in simulated neural network models, exploring effects of acquisition rate as well as other acquisition parameters, and then in spontaneous activity in the brain in absence of closed-loop brain-body interactions *in vivo*. Our technical contributions set a clear path for future works aimed at identifying sensory-motor circuits involved in the control of behaviour.

## 1.1 Thesis outline and contributions

We start by defining the context of this work in Chapter 2 where we discuss past research with particular focus on the acquisition and processing methods available to-date, their primary goals and limitations, outlining the importance of our contributions in further chapters in the context of the wider field of larval zebrafish research. In the second part we then focus on reviewing known applications and limitations of Granger Causality for estimating directed functional connectivity from neural imaging data, providing justification for our exploration of the method in the context of calcium fluorescence imaging.

In Chapter 3, we introduce the  $\mu$ SPIM Toolbox, an open-source software for control of a selective plane light-sheet microscope (SPIM) setup, built around the widely adopted MicroManager platform. We cover the control signal generation process which allows adoption for various functional imaging tasks and outline calibration methods included in  $\mu$ SPIM which provide the user with flexibility in terms of microscope hardware used such that the setup can be adapted to specific usage scenarios and requirements. This work has been published in "Saska, D., Pichler, P., Qian, C., Buckley, C. L., & Lagnado, L. (2021).

$\mu$ SPIM Toolset: A software platform for selective plane illumination microscopy. *Journal of Neuroscience Methods*, 347, 108952.”

In Chapter 4, we introduce a closed-loop virtual reality stimulator solution which uses a projector to deliver stimulus to a head-fixed larval zebrafish sample and visually captures its tail activity to allow for closed loop feedback. We base the functionality transforming tail activity into motion on a physically plausible parameterized model which we fit based on motion captured in freely swimming larvae.

In Chapter 5, we further introduce Granger Causality and establish its applicability to calcium fluorescence data in a number of simulated neural network models where the underlying structural topology is known. We address a number of experimental issues including measurement noise and signal drift and conclude by providing Granger Causality methods which are applicable to large number of sources in a structured neural network, such as that of a larval zebrafish brain.

In Chapter 6, we follow our results from simulated models and we apply Granger Causality to *in vivo* calcium fluorescence data which was collected from larval zebrafish in the absence of stimuli. We relate our findings of the inferred transfer of information between the observed brain regions to past findings of structural connectivity, providing a validation for Granger Causality on *in vivo* data.

We conclude by Chapter 7 which summarizes our findings and offers discussion on the applicability and limitations of Granger Causality in the context of calcium fluorescence imaging data. We also outline the limitations of our results and future research directions which consider the full sensorimotor loop, building on the basis of our work.

## Chapter 2

# Background

Studying the relationship between neural circuits and animal behaviour requires simultaneous access to both the behavioural variables as well as the neural activity in a live animal. This has led to increasing interest in larval zebrafish as an animal model for this purpose as acquisition protocols have been extensively established and the robust behaviour can be controlled through virtual reality stimuli. In this chapter we outline past research in terms of the available acquisition methods and protocols commonly used to establish causal connections between brain dynamics and behaviour. Further, we focus on the processing and analysis tools available and introduce Granger Causality in terms of its present application in neuroscience, outlining the importance of our work presented in the later chapters.

### 2.1 The Zebrafish Model

At early stages, the larval zebrafish offers superior optical accessibility of the brain tissue compared to other animal models as the tissue surrounding the brain is quite transparent (Halloran et al., 2000). This has further been improved over the past years by the elimination of many pigments in the tissue through genetic manipulation (Lister et al., 1999; White et al., 2008; Antinucci and Hindges, 2016), allowing modern acquisition methods the ability to image from nearly the full population of the 100 000 neurons in the brain (Chen et al., 2018b) at 6 days post fertilization (dpf). The well-established genetic manipulation protocols led to many useful advances in the zebrafish model, including genetically coded calcium indicators (Higashijima et al., 2003) which, in combination with well understood animal husbandry (Lawrence, 2011) allows for reliable access to samples with consistent brain imaging properties.

At young age, zebrafish also exhibit a wide repertoire of robust behaviours driven through different sensory pathways.

Visual pathway drives opto-motor reflex which allow the larva to combat incoming water current by integrating the visual flow parallel to the orientation of the larval body (Ahrens et al., 2012), opto-kinetic reflex which aids in stabilizing the eyes and the body of the larva and turning against incoming current through integration of angular optic flow (Portugues et al., 2014; Kubo et al., 2014) and prey capture during which eyes converge on a small moving spots in assistance to hunting and feeding (Bianco et al., 2011).

Auditory pathway signalling is still rather rudimentary in larvae with frequencies ranging from 100 to 1000 Hz resulting in neural responses (Vanwallegem et al., 2017; Privat et al., 2019). Low frequencies (150 to 450 Hz) were additionally found to elicit tail movement over time, suggesting that the auditory system is involved in the temporal integration of sensory input result in in motor behaviour (Privat et al., 2019).

Somatosensory inputs were then found to drive escape behaviours (Douglass et al., 2008) and olfactory system aids in navigation toward life-sustaining fresh-water environment through detection and avoidance of water salinity (Herrera et al., 2021) stimuli.

These behaviours are then generally studied in freely behaving animals (Olive et al., 2016) or in a head-restrained preparation in which imaging of the brain is more accessible (Ahrens et al., 2012) through suspension in agar gel, sometimes supplemented by injection of neuro-muscular blocker (such as  $\alpha$ -bungarotoxin, Chang (1962)) to aid image stability when the investigation of stimulus response is accompanied by brain imaging.

## 2.2 Acquisition Methods

The small size of zebrafish larvae prohibits use of many brain activity acquisition methods available in primates, however it also allows application of methods which are not available in larger animals. Namely, calcium fluorescence imaging where change in properties in fluorescent molecules upon binding with  $\text{Ca}^{2+}$  ions released during neural activity can be observed and serve as a proxy for the underlying neural activity of the individual cells imaged. Originally, this was achieved by injection of fluorescent dye into each sample (Fetcho and O'Malley, 1995; Cox and Fetcho, 1996) which has been replaced by genetically encoded reporters (Higashijima et al., 2003) with single fluorescent protein (GCaMP) sensors (Nakai et al., 2001) being most common. The state-of-the-art sensors continue to improve over time in terms of the dynamics and relative signal-to-noise ratio (Chen et al.,

2013; Dana et al., 2019) and offer better sensitivity and improve the applicability of the acquired data.

### 2.2.1 Two-photon Microscopy

Two-photon (2P) imaging (Denk et al., 1990) is one of the prevalent acquisition methods in larval zebrafish (Renninger and Orger, 2013) (but also other animal models such as mice, Stosiek et al. (2003)) which, rather than green light, utilizes infra-red (IR) light to excite the expressed fluorescent dye in the tissue. Since a single photon of the IR light does not provide enough energy to cause the dye to fluoresce, two (or in some modalities more) photons are needed at the same time to elicit the fluorescence that a single photon of higher wavelength light would. The requirement of simultaneous arrival of multiple photons results in a non-linear excitation probability with a much higher excitation chance at the focus point. This reduces acquisition from out-of-focus tissue as well as significantly limiting visible light pollution which is beneficial when imaging brain activity underlying visually evoked behaviours. The ability to use infra-red light is also advantageous due to its ability to better penetrate live tissue when compared to higher wavelength light (Helmchen and Denk, 2005).

Compared to other methods, two-photon microscopy is, however, a relatively slow acquisition method, particularly when acquiring volumes (Vanwalleghem et al., 2018) even in modern resonance scanner microscopes (Piyawattanametha et al., 2006). The constraint to smaller areas is, however, accompanied with a significantly superior spatial resolution which allows for imaging of individual synapses (Lu et al., 2017).

### 2.2.2 Selective Light-Sheet Microscopy

Selective Light-sheet Imaging (SPIM) has been originally applied to imaging of development of live embryos (Huisken et al., 2004) and has only recently been applied to functional calcium imaging (Ahrens et al., 2012). Unlike two-photon imaging which uses point scanning to form a plane which is then acquired by the camera, SPIM uses a beam of visible light which is expanded using a cylindrical lens or scanned by a mirror galvanometer in order to form an illumination sheet within the tissue sample. By ensuring that the excitation plane coincides with the focal plane of the collection objective, the fluorescence from the tissue sample can then be collected by the camera. While a single frame allows for collection of single illuminated plane at the same depth (along the axis of the collection objective), scanning the illumination plane along the collection axis using a secondary mirror

galvanometer in the illumination path allows for sequential illumination of tissue at different depths. Correct acquisition is enabled by movement of the collection objective (e.g. by a piezo-electric stage) such that the illumination always stays in focus throughout the acquisition. This allows for a volumetric acquisition of nearly the full volume of the larval brain (Chen et al., 2018b). The collection of activity from larger neural populations then offers recordings with comparably lower bias, allowing for identification of sparsely coded or distributed cell types (Feierstein et al., 2015). A modification to the SPIM design by Wolf et al. (2015) allows for two-photon light sheet excitation, which provides improved applicability in the study of visually-driven behaviour due to the reduced light pollution when compared to single-photon light.

Unlike for imaging of developmental processes, where a number of microscope control and acquisition solutions have been published, including OpenSPIM (Pitrone et al., 2013) and the Open SPIN microscopy (Gualda et al., 2013) project, functional imaging is usually accomplished by the use of proprietary or in-house solutions (Ahrens et al., 2012; Reynaud et al., 2015).

## 2.3 Virtual Reality and Behavioural Acquisition

Although imaging of neural activity in freely swimming zebrafish larvae is possible (Kim et al., 2017; Symvoulidis et al., 2017), it is still challenging which is reflected in the number of studies which utilize head-fixed preparations for simultaneous stimulation, calcium imaging and behavioural acquisition (Ahrens et al., 2012; Severi et al., 2014; Semmelhack et al., 2014; Naumann et al., 2016). In such preparations, a stimulus projected on a screen is commonly used to emulate the environment in studies of visually-evoked behaviour (Ahrens et al., 2012; Severi et al., 2014). To then access the swim information, the activity of the tail is acquired either visually through imaging its movement and tracking it over time (Semmelhack et al., 2014; Severi et al., 2014; Portugues et al., 2014) or by electrode acquisition of the neural signal directly from the tail (Ahrens et al., 2012, 2013a).

Studies of behaviours without the opportunity for the fish to interact with its environment (open-loop) are common (Severi et al., 2014; Oldfield et al., 2020), however it has been shown that the feedback of the motor action through the environment is an important input for the behaviour of the fish and its absence leads to behaviour expression unlike that which occurs in a freely-swimming setting (Portugues and Engert, 2011; Ahrens et al., 2012), promoting the use of closed-loop stimulation systems in order to quantify the neural origin of aspects of behaviour which were not apparent in open-loop setting

([Kawashima et al., 2016](#)).

## 2.4 Data Processing Methods

Processing of calcium imaging data can be challenging due to the increasing sizes of datasets, especially in the case of light-sheet recordings where a single 10-minute acquisition can yield over 100GB of imagery. With the volumetric nature of the recordings and the number of sources in each recording being potentially in the tens of thousands, manual labeling is not an option, requiring the implementation of compute-based cell identification methods.

For this purpose, [Pachitariu et al. \(2017\)](#) introduced Suite2P, a toolbox focused on the identification of sources in calcium recordings acquired using two-photon microscopy. While the software has the ability to identify neurons from multiple planes of a volume recording, it does offer tracking the same cells over multiple planes which can be problematic for the analysis of light-sheet recordings which generally have higher depth resolution where this is necessary. [Giovannucci et al. \(2018\)](#) later introduced the CaImAn toolbox which provides a number of tools for the analysis of calcium data, generally better suited for light-sheet microscopy acquisition: By providing 3D non-rigid motion correction and cell extraction based on non-negative matrix factorization ([Pnevmatikakis et al., 2016](#)) it allows for much more flexible identification of a larger number of neuronal sources.

## 2.5 Present approaches to data analysis in calcium imaging

In order to interpret data, many whole-brain imaging publications rely on adapting existing single-neuron analyses, most frequently limited to correlating the activity to the observed behavioral parameters. While early studies had to rely on per-pixel data ([Ahrens et al., 2012](#); [Portugues et al., 2014](#)), more recent publications take advantages of aforementioned processing methods and perform analysis on individual cells ([Chen et al., 2018b](#)).

### 2.5.1 Atlas Registration of Calcium Recordings

In order to aggregate data over multiple recordings and larval samples, it is necessary to map the recording to a common reference. A number of zebrafish atlases have been published, including Z-brain ([Randlett et al., 2015](#)), MPIN ([Kunst et al., 2019](#)) or ViBE-Z ([Ronneberger et al., 2012](#)) which can serve as common registration target as they provide reference volume brains from different zebrafish lines and expression backgrounds.

### 2.5.2 Data dimensionality reduction

With many individual neuronal sources, it can be desirable to identify functional similar populations of neurons (Cunningham and Byron, 2014). A number of different approaches have been used with calcium data. In particular, some studies group neurons based on their spatial location into functionally similar regions either manually (Oldfield et al., 2020) or by mapping on a common atlas (Kunst et al., 2019). An alternative approach based on unsupervised clustering of the sources based on their calcium activity (Naumann et al., 2016; Chen et al., 2018b) can offer lower bias.

### 2.5.3 Undirected functional connectivity measures

A common approach to establishing the relationship between neural activity and behaviour is by correlating the changes in activity of the individual neurons or brain regions to the changes in behaviour or stimulation. This approach offers only limited information on the affinity of the neurons by correlating them to different stimuli (Naumann et al., 2016) or motor activity (Ahrens et al., 2013a). This in itself is not sufficient to determine connectivity between neurons (Orger, 2016), requiring additional methods such as tracing of neuronal projections (Hildebrand et al., 2017; Kunst et al., 2019), opto-genetic stimulation of cells (Kaur et al., 2017) or targeted ablation of cells (Roeser and Baier, 2003) or dendritic connections (Naumann et al., 2016) in order to confirm hypotheses about the functional interactions within the brain that give rise to the behaviour studied.

## 2.6 Directed functional connectivity measures

The flow of information in the brain and directed functional interactions between its part during behaviour can offer more insight than simple undirected functional measures. A number of measures of directed functional connectivity have been formulated in order to quantify information flow and causality from data.

### 2.6.1 Time-shifted Cross-correlation

One of such measures is cross-correlation which quantifies causality through the computation of cross-correlogram based on Pearson's correlation coefficient between two time series where the source time series is shifted by a set number of time lags with respect to the target time series. In this way, the cross-correlation can provide information about precedence, however it does not take into account the past time of the target time series

such that a cross-correlation between a source and completely deterministic target process could still indicate high directed causal connection.

### 2.6.2 Granger Causality

This issue has been addressed by [Granger \(1969\)](#), who devised Granger Causality (GC), a statistical measure of a causal relationship between two processes such that the cause must always precede its effect and the causal process must provide improvement in prediction of the effective process. This avoids the issue of possible high causality measures where the target process is (at least in part) deterministic. In practice, this concept of 'improved prediction' is implemented in terms of the improvement in variance of the prediction error of a vector auto-regressive (VAR) model based on the past of the target process by inclusion of the past history of the causal process. Later on, the original GC measure has been extended to the case of inferring the directed functional connectivity between two multi-dimensional (multivariate) time series ([Geweke, 1982](#)) and later to a case which conditions on other processes such that indirect causal connections can be effectively ruled out ([Geweke, 1984](#)).

### 2.6.3 Transfer Entropy

Transfer Entropy is a non-parametric generalization of a Granger Causality proposed by [Schreiber \(2000\)](#). It can be defined as mutual information between the current value of the target process and the past of the source process conditional on the past of the target process. Transfer Entropy and Granger Causality have been shown to be equivalent for Gaussian processes ([Barnett et al., 2009](#)). Transfer Entropy can be a useful alternative to GC in case of non-linear processes, however it requires significantly more data to estimate accurately [Pereda et al. \(2005\)](#).

### 2.6.4 Application of Granger Causality to Calcium Data

A number of studies have attempted to apply Granger Causality to calcium imaging data. [Stetter et al. \(2012\)](#) show unconditional Granger Causality in context of their proposed TE-based measure of causality on small networks showing some but not optimal performance for structural connectivity reconstruction.

No study has systematically demonstrated the limitations of applicability of Granger Causality to calcium recordings in large neuronal populations such as those found across the brain of larval zebrafish. Some studies applying Granger Causality to calcium data

from this setting have been published however: [Fallani et al. \(2014\)](#) apply Granger Causality to neurons in the tail of zebrafish embryos indicating an ipsi-lateral directional flow of information between the neurons in the direction away from the brain. [Vanwallegem et al. \(2017\)](#) apply GC to volumetric calcium recording data in presence of auditory stimuli, presenting neurologically plausible results but do not offer significant interpretation of their findings. [Oldfield et al. \(2020\)](#) investigate prey capture in naive and experienced zebrafish larvae and attempt to use Granger Causality to quantify the difference between the two groups of fish.

### 2.6.5 Known limitations of Granger Causality

Granger Causality has been extensively studied in the context of EEG and MEG recordings in humans ([Seth et al., 2015](#)) and thus many of its limitations are already known. In particular, GC assumes the input data to be stationary which can often be achieved through high-pass filtering ([Florin et al., 2010](#); [Barnett and Seth, 2011](#)). Furthermore, the sampling interval of the data has been shown to have a profound impact on Granger Causality inference as for acquisition rates significantly higher than the causal timescale, Granger Causality ends up fitting noise, reducing the statistical power of the inferred result and much slower acquisition speeds lead to identification of spurious and missed connections as the it becomes impossible to determine the time precedence of the events in the time series [Barnett and Seth \(2017\)](#). The latter has been a source of controversy over the applicability of Granger causality to fMRI data where the acquisition speeds are much slower (order of seconds) than the underlying synaptic delays (order of milliseconds) ([Seth et al., 2015](#)). Further studies have shown that for fMRI BOLD data, is in fact invariant to the slow convolution of the underlying neural activity but not significant down-sampling [Seth et al. \(2013\)](#). This is encouraging for our application on calcium fluorescence data as while acquisition methods such as light-sheet microscopy allow recording at high acquisition speeds (100Hz) the calcium reporter dynamics are slow [Chen et al. \(2013\)](#) compared to the time scale of the underlying neural processes.

## Chapter 3

# $\mu$ SPIM Toolset: A software platform for selective plane illumination microscopy

### 3.1 Abstract

Selective Plane Illumination Microscopy (SPIM) is a fluorescence imaging technique that allows volumetric imaging at high spatio-temporal resolution to monitor neural activity in live organisms such as larval zebrafish. A major challenge in the construction of a custom SPIM microscope using a scanned laser beam is the control and synchronization of the various hardware components.

We present an open-source software,  $\mu$ SPIM Toolset, built around the widely adopted MicroManager platform, that provides control and acquisition functionality for a SPIM. A key advantage of  $\mu$ SPIM Toolset is a series of calibration procedures that optimize acquisition for a given set-up, making it relatively independent of the optical design of the microscope or the hardware used to build it.

$\mu$ SPIM Toolset allows imaging of calcium activity throughout the brain of larval zebrafish at rates of 100 planes per second with single cell resolution.

Several designs of SPIM have been published but are focused on imaging of developmental processes using a slower setup with a moving stage and therefore have limited use for functional imaging. In comparison,  $\mu$ SPIM Toolset uses a scanned beam to allow imaging at higher acquisition frequencies while minimizing disturbance of the sample.

The  $\mu$ SPIM Toolset provides a flexible solution for the control of SPIM microscopes and demonstrated its utility for brain-wide imaging of neural activity in larval zebrafish.

## 3.2 Introduction

Selective Plane Illumination Microscopy (SPIM) is a powerful method for 4D imaging of biological samples at a high spatio-temporal resolution (Power and Huisken, 2017; Ahrens et al., 2013b). This is achieved by excitation of fluorescent structural or functional reporters expressed in the sample by a few micron thick light sheet and simultaneous recording by an orthogonally positioned high-resolution camera focused on the plane of excitation. Rapid movement of the sheet through the tissue then allows for volumetric recordings with cellular resolution. While SPIM imaging has been originally used for the study of developmental processes in a number of animal models including embryos of *C. elegans* (Rieckher et al., 2015), *Drosophila* (Keller et al., 2010) and zebrafish (Kobitski et al., 2015; Wan et al., 2019), more recently it has also found its use in functional imaging as a complement to already established imaging methods such as two photon imaging. By providing the ability to image much larger volumes of tissues while maintaining temporal and spatial resolution, SPIM provides the ability to investigate the interaction of much larger neuronal populations as illustrated by imaging of the whole nervous systems in *Drosophila* embryos and larvae (Chhetri et al., 2015; Lemon et al., 2015), *C. elegans* (Ardiel et al., 2017) and whole-brain imaging in zebrafish (Ahrens et al., 2012, 2013b; Weisenburger and Vaziri, 2018).

While all implementations of SPIM share the common design including an illumination arm to create a 2D plane of illumination and an orthogonal collection arm that is forming the image onto the camera, some variations on this basic design have been developed with different applications in mind (Keller and Ahrens, 2015). These can be divided based on two main aspects of the design: firstly, the method of light sheet formation and secondly, the movement of the light sheet relative to the sample. The original approach creates a stationary light sheet using a cylindrical lens and then translates or rotates the sample with a moving stage. This implementation has been particularly useful in studies of development (Huisken and Stainier, 2009), but it is too slow to monitor the activity of neurons using, for instance, genetically-encoded calcium indicators. Applications in functional neuroscience therefore favour a configuration in which the sample is kept stationary while the light sheet is created using a fast scanning mirror that moves the light beam across a plane at least once per imaging frame, with a secondary mirror moving the beam in the z-dimension (Fig. 3.2A). This method has allowed "brain-wide" imaging of neural activity in live zebrafish with single neuron resolution and acquisition frequencies of  $\sim 1$  Hz (Ahrens et al., 2012, 2013b; Weisenburger and Vaziri, 2018).

Several SPIM solutions have been published but probably the most accessible in terms of both hardware and software are the OpenSPIM (Pitrone et al., 2013) and Open SPIM microscopy (Gualda et al., 2013) projects. Both these implementations are focused on imaging of developmental processes using the slower "moving sample" configuration and therefore have limited use if the aim is to image neural activity through volumes of the brain. The "moving light-sheet" configuration is more complex to control because it requires synchronization of several hardware components with millisecond temporal accuracy. For instance, movements of the imaging objective in the z-dimension must be synchronized with movement of the light-sheet to stay focused on the plane of illumination, often involving 50–100 planes per second. Perhaps for this reason, open software tools for the control of SPIMs in functional imaging experiments are not easily available and published research utilizing functional imaging generally does not include openly published and documented microscope control solutions. We aim to fill this gap by providing a control software solution,  $\mu$ SPIM Toolset, which adopts an open software approach for control of a SPIM microscope in which the light-sheet is scanned through a stationary sample. We achieve this by building  $\mu$ SPIM Toolset around Micro-Manager (Edelstein et al., 2010), an open-source platform widely used for control of microscopes, to provide a comprehensive user interface (Fig. 3.3) with a smooth learning curve.  $\mu$ SPIM Toolset has been designed to be used to control a range of custom-built microscopes, for which it is calibrated using semi-automated procedures. We demonstrate the utility of  $\mu$ SPIM Toolset for whole-brain imaging in larval zebrafish.

### 3.3 Design requirements

We start by listing the basic objectives of software controlling a SPIM microscope in relation to the microscope we have constructed (Fig. 3.1, see Methods). 1. The illumination arm requires the control of the laser light-source, using either an internal and/or external shutter. Dual-colour imaging requires two independent laser sources to be combined. Depending on the laser model, power may be controlled on long time-scales using USB or RS232 interfaces and modulated on short time-scales by analogue signals, while shuttering requires digital signals. 2. The beam has to be scanned in the x and z dimensions using fast galvanometer mirrors driven through analogue inputs. 3. The collection arm must follow the beam and collect images that are in focus at different z-positions, which is achieved by using a piezo-electric mount for the objective, again controlled through an analogue signal. Finally, 4. The camera acquisition must be accurately timed to collect

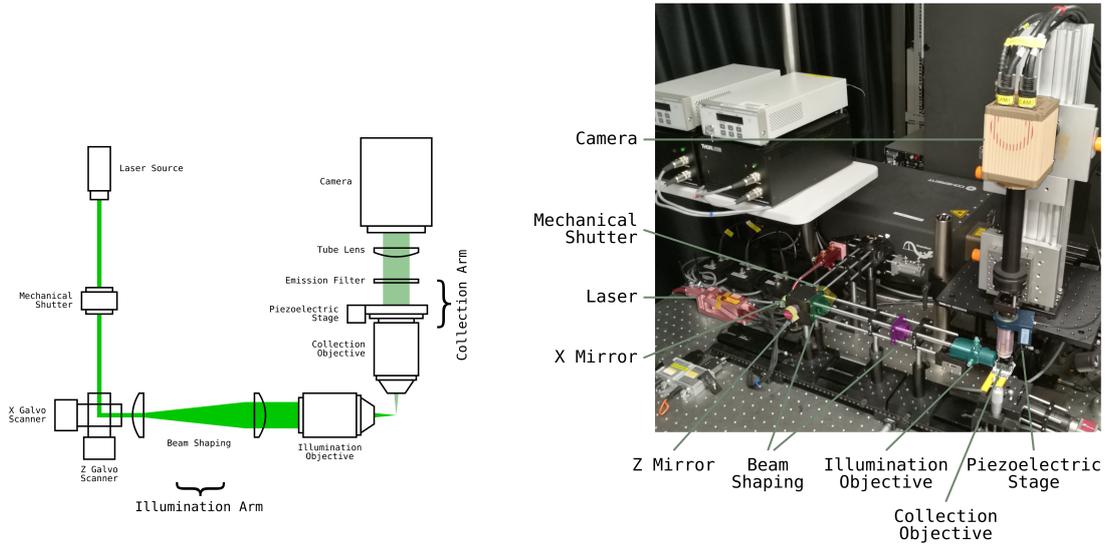


Figure 3.1: **Light-Sheet Microscope Implementation:** **A.** Diagram outlining light-sheet microscope with two mirror galvanometers and a stationary sample. The light sheet is created from a laser beam by the 'X Galvo scanning mirror'. The 'Z Galvo scanning mirror' then moves the sheet through the sample to create volumetric excitation. This is synchronized with a Piezoelectric stage that moves the imaging objective so that the excitation plane always coincides with the imaging plane. **B.** A photo of an example light-sheet microscope setup with one light-sheet path. The main components are colour coded.

frames at each  $z$  position, which in turn requires accurate calibration of the  $z$  mirror so that the  $z$  position of the light-sheet is set from the command signal. The software should then allow control of several features of the camera, including frame integration time, pixel binning and gain. Furthermore, it should support saving and retrieving images and movies, zooming, defining ROIs, quick and easy changes in imaging parameters and the ability to define various imaging protocols using different sequences of laser illumination.

We built  $\mu$ SPIM Toolset around the open-source software Micro-Manager ([Edelstein et al., 2010](#)), which is based on ImageJ ([Schindelin et al., 2015](#)), because it immediately offers the ability to integrate different hardware from a wide range of manufacturers through specific plugins, including a range of cameras, light-sources and shutters required for a SPIM. It provides an integrated environment for image acquisition and a very wide range of post-acquisition processing capabilities through ImageJ. Micro-Manager does not, however, provide hardware triggering, analogue control and monitoring with the precision required for volumetric imaging of neural activity. For this purpose, we used a National Instruments DAC card (NI PCIe-6738) and wrote an executable ' $\mu$ SPIM Control' to allow interaction between Micromanager and hardware through its own user interface.  $\mu$ SPIM Control synchronizes internal laser shutters,  $x$  and  $z$  mir-

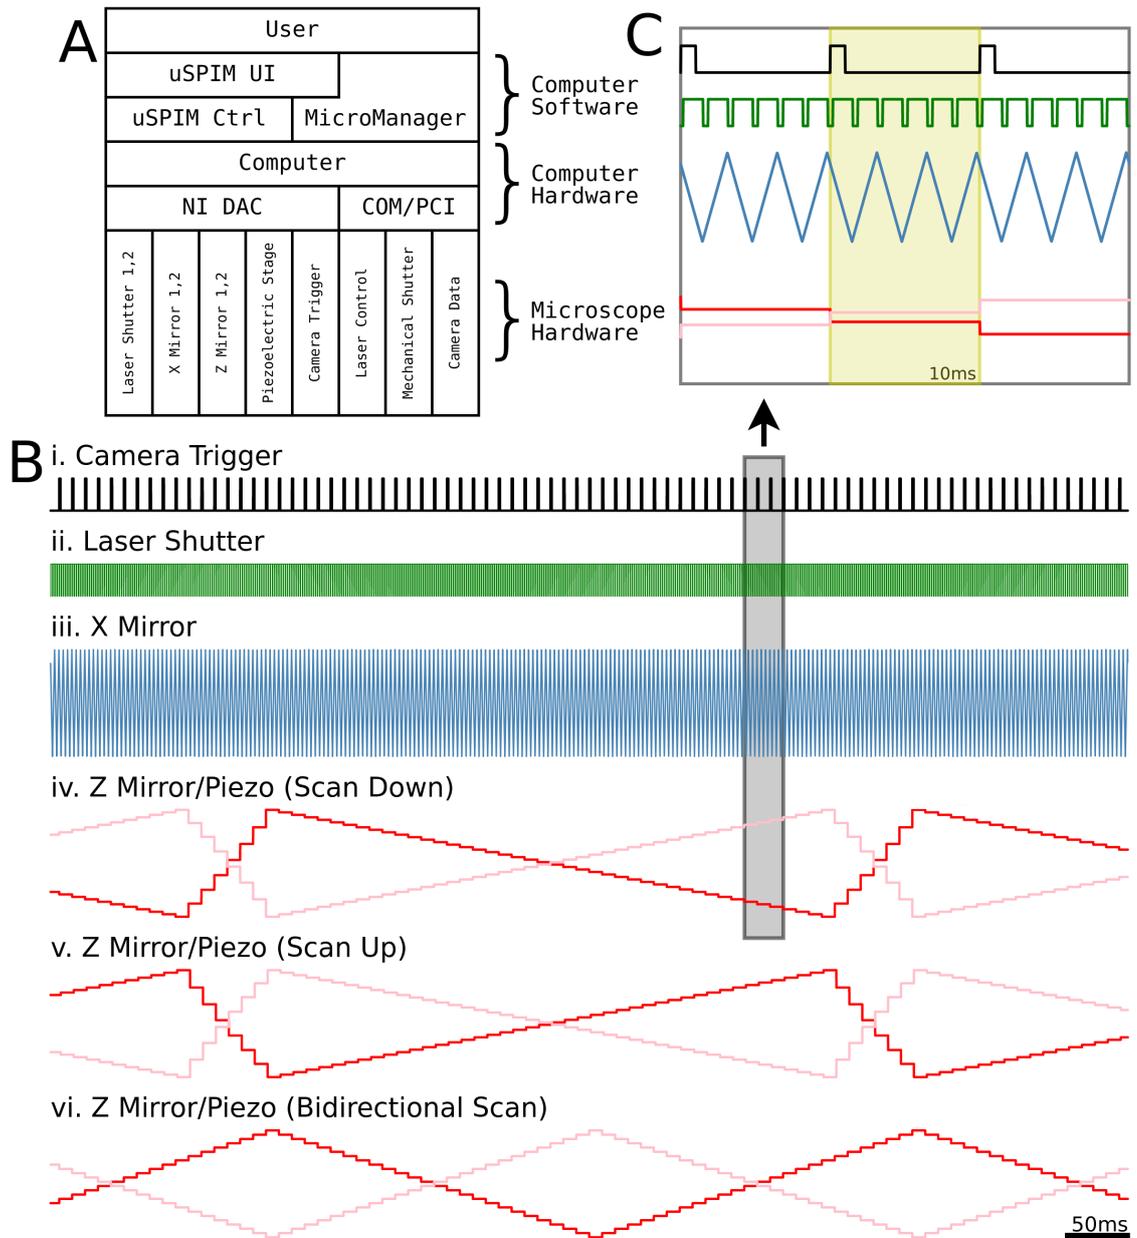


Figure 3.2: **Light-Sheet Microscope Hardware Control:** **A.** Interaction of different components of the  $\mu$ SPIM Toolset-based setup in a typical acquisition setup.  $\mu$ SPIM Toolset provides control and synchronization of hardware through NI DAC with MicroManager controlling the camera and mechanical shutters through PCI and COM ports. **B.** Traces of the command signals for a volume acquisition with 43 planes showing camera trigger (i.), laser shutter (ii.), X mirror signal (iii.), Z mirror (red) and Piezo (pink) signals for Scan Down (iv.), Scan Up (v.) and Bidirectional (vi.) acquisition modes generate by the  $\mu$ SPIM control software. **C** Magnification of boxed region in B with yellow region showing a single plane signal. Laser shutter signal is a result of a recording with Edge Masks enabled.

rors, the objective piezo stage and all camera triggering, leaving laser power, external mechanical shutters and camera operating properties under control of Micro-Manager (Fig. 3.2A).  $\mu$ SPIM Toolset therefore provides both accessibility and versatility, retaining the user’s ability to select hardware best suited for the particular setting, as long as it is supported through a Micro-Manager plug-in (The current range can be viewed at <https://micro-manager.org/wiki/Device%20Support>).

### 3.4 Principles of operation

$\mu$ SPIM Toolset consists of two main components: a Java plugin for MicroManager which facilitates all user interaction and configuration, and a C++ control executable which interacts with the National Instruments board to produce hardware control signals (Fig. 3.2). These support the control of two scanning X mirrors for light sheet formation, two scanning Z mirrors and one piezoelectric stage for Z motion of the collection objective, two laser shutters for laser masking and one trigger for the camera acquisition control and synchronization with other hardware. After an appropriate setup, the Java plugin facilitates all necessary tools for calibration and control of the hardware, providing the users with an intuitive interface without requiring the need for specialized computer knowledge or programming, while allowing flexibility provided by MicroManager. Fig. 3.2 illustrates the control diagram, outlining the interaction between the user and the light-sheet microscope hardware (Fig. 3.2A) as well as the temporal sequence of the control signals (Fig. 3.2B and C).

A typical user interface for data acquisition includes the  $\mu$ SPIM Control window (Fig. 3.3A) and the standard MicroManager windows (Fig. 3.3B-D). The primary Microman-ager window (Fig. 3.3C) provides access to all functionality related to the acquisition, such as live mode, binning, ROI selection or exposure time. The  $\mu$ SPIM plugin (Fig. 3.3A) provides access to the variables used for light-sheet formation, such as light-sheet width, volume depth or number of planes as well as providing a set of calibration procedures described in the Calibration section below. During scan mode, the image can be viewed in the classical MicroManger live window (Fig. 3.3D) which also allows the acquisition of snapshots. While acquisition can be manually started through the Micro-Manager interface, the  $\mu$ SPIM plugin interface provides an acquisition routine executing a synchronized start of acquisition and illumination, making the process of acquiring data (namely volumes) simpler and more robust.

Taking full advantage of the stationary sample and fast acquisition speeds, we used

A

1. Find Plane 2. Design Mask 1 3. Design Mask 2 4. Acquire Data Calibration Configuration

Number of volumes: 30000 Light Sheet Width: 250.0 300.0 um  
 Scan mode: Scan Down Framerate: 97.92314579210255 fps  
 Number of planes: 1 and 0 return  
 Z Start 100.0 um -10 -1 +1 +10  
 Z End 50.0 um -10 -1 +1 +10

Acquisition will produce 30000 frames (30000 volumes consisting of 1 planes and 0 for return) and take 5.1060min@97.92fps.

Advanced  
 Light Adaptation Phase:  Duration: 2 min  
 Sequential Acquisition:  Add Remove Move Up Move Down

Type	Mode	Volumes	Planes	Return	Z Start	Z End	Duration	Acquire?

Config: C:\Users\ds376\normalmode\_default.json  
 Load Config Save Config Start Acquisition Pause

```
ScalePifoc=40.000000, CyclesPerFrame=3
5, Masks1="", Masks2="", GalvoDelay1=0.000000, GalvoDelay2=0.000000, NumVolumes=10000000.00000
ScalePifoc=40.000000, CyclesPerFrame=3
```

ControlEngine link: ACTIVE  
 Light Sheet Realtime:  
 Direction: Down Vol Start: 100.00  
 Position: 100.00 Vol End: 100.00  
 Last update 0s ago

Acquisition:  
**READY**

B

ImageJ

File Edit Image Process Analyze Plugins Window Help

x=368, y=1009, value=254

D

Snap/Live Window (not yet saved) (75%)  
 1128x1108 pixels, 16-bit, 2.4MB

x=368, y=1009, value=255 Pos0, 12.93s, Default  
 Snap Stop ... Album FPS: 93.2 (display 34.2)

100µm

C

Micro-Manager 1.4.22 - C:\SPIMConfig4\TheBetterSPIM2.cfg

File Tools Plugins Help

Camera settings Exposure [ms] 9.9977 Binning 1x1 Shutter AllShutters Auto shutter  Close

Configuration settings Group Preset CameraTrigger Internal ScanMode Normal System Startup VisibleLaserFront On VisibleLaserSide On

Image info (from camera): 1128 X 1108 X 2, Intensity range: 16 bits, 0nm/pix

Contrast Metadata Comments  
 Scale Bar Top-Left White Sync channels  Slow hist  
 Display mode: Grayscale Autostretch ignore % 2 Log hist

Full Auto  
 Hist range: Camera Depth  
 Min: 63 Max: 631 Mean: 192 Std Dev: 88

Figure 3.3: **μSPIM Toolset User Interface:** Following a common MicroManager design, the control interface is separated into several components: μSPIM Toolset-provided plugin window with control over the light-sheet generation including calibration and acquisition routines shown in **A.**, basic ImageJ tools shown in **B.** and MicroManager interface providing control over the acquisition hardware shown in **C.** and live view of the camera shown in **D.**. The shown user interface has been captured during the acquisition from a 7 dpf larval zebrafish from the *Tg(elavl3:H2B-GCaMP6f)* transgenic line.

$\mu$ SPIM Toolset to monitor the activity of neurons across most of the brain of larval zebrafish expressing the nuclear localised Ca<sup>2+</sup> reporter GCaMP6f panneuronally (*Tg(elavl3:H2B-GCaMP6f)*), Fig. 3.4). Larvae were embedded so that the laser excited the sample from the side (arrows in Fig. 3.4A). The width of the sheet spanned 500  $\mu$ m and covered the entire hindbrain, the cerebellum and parts of the optic tectum. Single planes were recorded at a frequency of 98 Hz. The volume was set to cover a distance of  $z = 100$   $\mu$ m and consisted of 50 planes with a spacing of 2  $\mu$ m. The piezo returned to its original position within 10 frames yielding an overall imaging frequency of 1.6 volumes per second. Fig. 3.4C shows three representative sections at depths of 20, 30 and 40  $\mu$ m, respectively. The lateral and axial resolution of the microscope were  $0.7 \pm 0.1$   $\mu$ m and  $5.4 \pm 0.5$   $\mu$ m, respectively (mean  $\pm$  sd, full width at half maximum, see Methods) and the nuclei of individual cells are clearly visible in the cerebellum, the anterior and the posterior hindbrain, respectively (Fig. 3.4D) and their activity over time is shown in Fig. 3.4E.

### 3.5 Computing requirements

$\mu$ SPIM Toolset relies only on minimal third-party software for its operation. This includes MicroManager for integration of acquisition hardware (laser control, light path shutter and camera) and National Instruments drivers, in order to interface with the National instruments DAC card. Naturally, drivers and accompanying software must be installed for the selected hardware (such as lasers or camera) in order to ensure MicroManager can communicate with those parts properly. Due to the continuous high data throughput, it is recommended to use a secondary computer for stimulation and/or other intensive tasks to avoid potential performance issues caused by the multiple tasks competing for computational resources or ensuring that this is not the case when a single machine is used.

The computer specifications necessary for optimal function and acquisition of the setup are highly dependent on the acquisition requirements. For large field of view and high acquisition speeds, the acquisition produces large datasets, making the process highly storage-dependent. State-of-the-art acquisition cameras are able to reach 4 megapixel resolution at 100 Hz and 16 bit depth, producing a theoretical throughput of 838.9 MB/s. This poses a write speed requirement higher than that which can be satisfied by a regular SATA3 storage interface which is limited to a theoretical maximum of 600 MB/s. It is therefore highly recommended to utilize a faster storage solution, such as fast PCI-based solid state disk (SSD) storage (e.g. Intel Optane 905p) or hardware SSD RAID. The

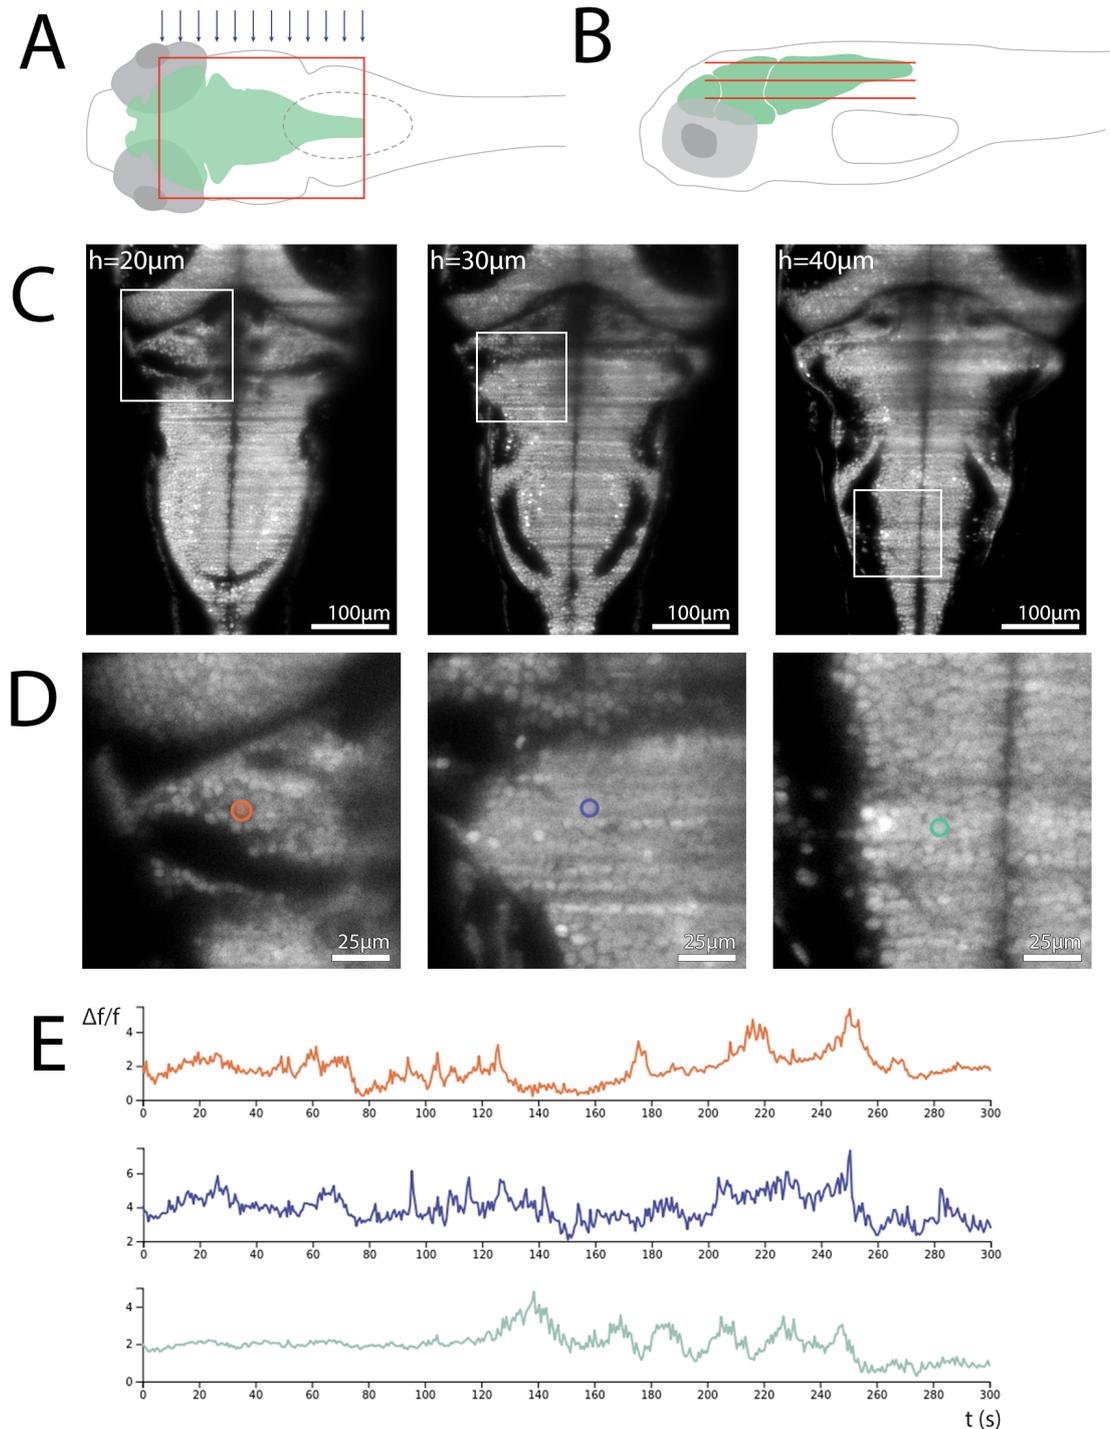


Figure 3.4: **Imaging larval zebrafish:** **A.** Schematic of a single light sheet covering the hindbrain of a larval zebrafish from above as well as a number of light sheets (constituting a volume) from the side in **B.** (sheet size and spacing are not to scale). **C.** Transgenic zebrafish expressing the calcium reporter GCaMP6f in the cell nuclei of all neurons (*Tg(elavl3:H2B-GCaMP6f)*) were embedded in agarose and positioned so that the laser entered the brain from the side (blue arrows in A). Representative sections of the volume taken at 20, 30 and 40  $\mu\text{m}$  depths at a frequency of 1.6 volumes/sec and an integration time of 10.2 ms per section. The step size was 2  $\mu\text{m}$ , hence representative sections are 5 sections apart. The laser was set to 1.8 mW. **D.** magnified view of boxed areas in C. Single cell nuclei are clearly visible and three examples are highlighted. **E.** Excerpts of activity traces of single cells highlighted in D,  $\Delta f/f = (f - \min(F))/\min(F)$  where  $F$  is full activity trace.

acquisition is not highly CPU dependent and neither  $\mu$ SPIM Toolset nor MicroManager take a significant advantage of parallel processing. A state-of-the-art consumer line 4- or 6-core CPU is sufficient for optimal performance.

$\mu$ SPIM Toolset supports control of two independently calibrated laser paths. Each path consists of a visible light laser, two scanning mirrors and one laser path shutter (Fig. 3.1). The laser path shutter must be supported by MicroManager and is used to block light between periods of acquisition and thus reduce photo-bleaching of the sample. One of the two scanning mirrors is used for light sheet formation and the second is used for Z movement. The laser should provide digital or analogue shuttering capability. All other devices are controlled using an analogue signal produced by the National Instrument DAC adapter and should support inputs in appropriate voltage ranges (commonly  $\pm 10$  V).  $\mu$ SPIM Toolset allows for incomplete configurations e.g. a single light path with unmodulated laser and a single sheet-generating scanning mirror would provide minimal functionality, however most of the features of the  $\mu$ SPIM Toolset would be unavailable. To support the full functionality, the National Instruments DAC card must support at least 8 analogue outputs with no need for analogue or digital inputs as the control is fully feed-forward and does not require any feedback from the hardware.

## 3.6 Calibration

The spatial scales of the signals are greatly dependent on the chosen hardware and thus require the user to perform a set of calibration procedures to ensure appropriate alignment of the hardware prior to the first use of the software for acquisition. Subsequently, the calibration needs to be done only when hardware components are changed and routinely when it is suspected the software has gone out of calibration. Here and in Fig. 3.5 we outline the principles of the initial calibration of the different components sufficient to achieve an acquisition-ready configuration.

### 3.6.1 X-Mirror

Proper X-plane mirror calibration is essential for the correct assessment of the size of the illuminated region when starting acquisition. To calibrate the X mirror, it is necessary to know the size of the field of view (or pixel size). This can be calculated as the effective camera image sensor size divided by the magnification of the collection objective or by measuring a standard stage calibration slide under the microscope. Using this measurement, the X Mirror scale can be adjusted such that a sheet of a particular width matches

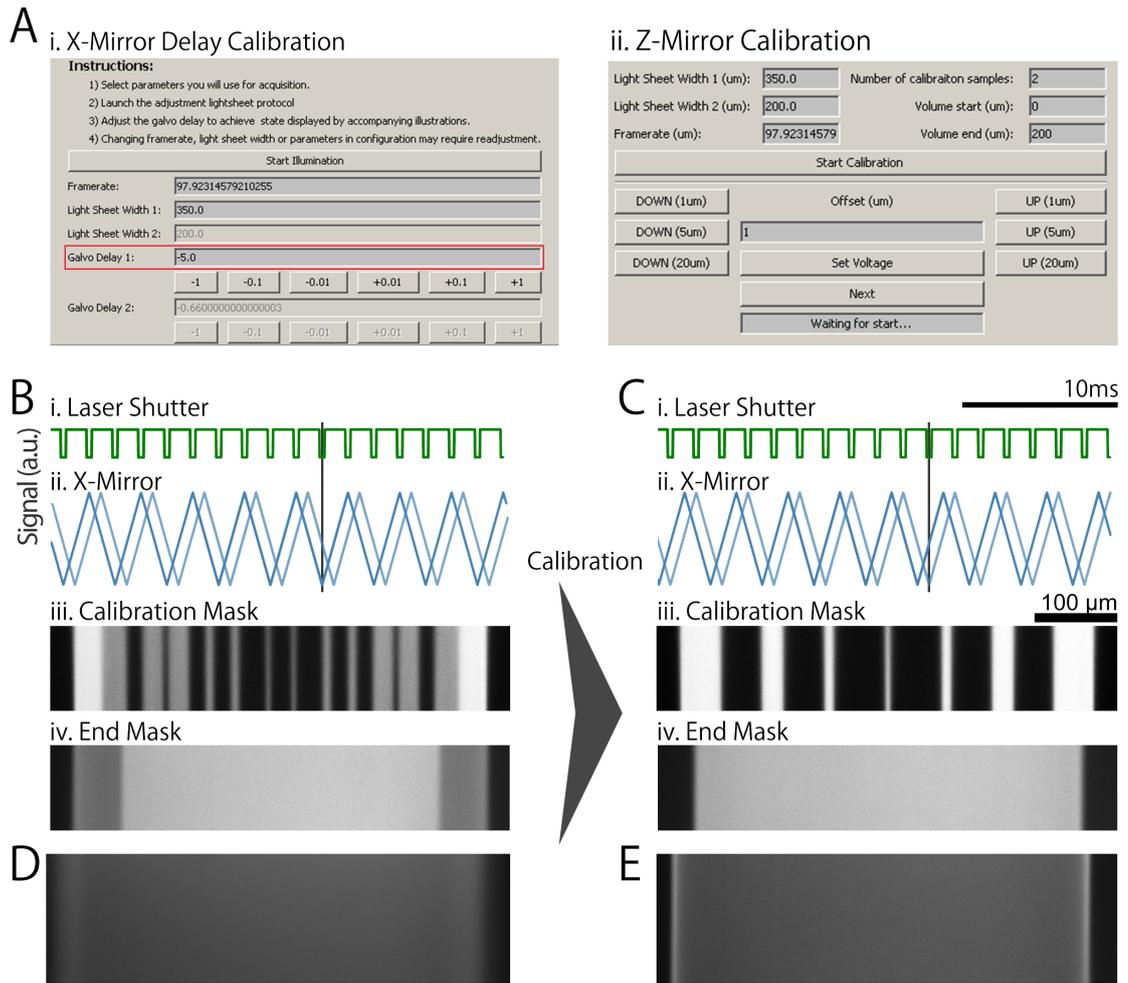


Figure 3.5:  **$\mu$ SPIM Toolset Calibration:**  $\mu$ SPIM Toolset provides the user with calibration tools which can be used to assess and correct the performance of the individual microscope elements. **A. i.** and **ii.** show the calibration interface for X mirror with respect to the laser shutter when using masks and for Z mirror, respectively. **B.** The movement of the X mirror galvanometer lags behind the supplied signal, resulting in artefacts when using laser masks. **(i.)** shows the laser shutter signal (green) and **(ii.)** shows the X mirror signal (blue) and actual mirror movement (light blue). Black bar highlights the misalignment between the shutter signal and mirror movement when the X mirror signal is not calibrated and no delay is introduced. **(iii.)** shows a mask used for calibration, emphasizing the difference between uncalibrated and calibrated system. **(iv.)** shows the corresponding light sheet produced with poor calibration for the mask signal in **(i.)**. **C.** is analogous to **(B)**, showing a well calibrated system where the laser shutter **(i.)** is synchronized X mirror movement **(ii.)**, producing overlapping masks shown in **(iii.)** and **(iv.)**. **D.** Uncalibrated Z-Mirror results in loss of focus through the volume. **E.** After calibration, the light sheet is always in focus throughout the volume scan. All figures illustrating the light sheet in both uncalibrated and calibrated scenarios were acquired by scanning the laser through a fluorescein solution.

the corresponding number of the pixel. In the case of two illumination paths, this procedure should be repeated for both paths separately. This calibration should be carried out such that the x-mirror moves around its central position when energized (i.e. with zero volts driving signal) and the illumination arm should be adjusted such that this central position is centred on the field of view.

### 3.6.2 Piezoelectric stage

Piezoelectric stages are generally factory-calibrated to move a predefined distance by set voltage difference (e.g. 1  $\mu\text{m}$  per 40 mV for the PIFOC model) to which the Piezo Scale setting should be set. The unit is mV/ $\mu\text{m}$  and for the previous example this setting would be set to 40.0.  $\mu\text{SPIM}$  Toolset also provides a Piezo Offset setting which can be used to adjust voltage offset in software. For optimal operation, it is suggested to adjust the offset such that a plane at 0.0  $\mu\text{m}$  is near the top of the available range. Setting the offset could be beneficial, for instance, when imaging at multiple depths in a sample so that when 0.0  $\mu\text{m}$  is positioned at the top of the sample it serves as a reference while allowing the full range of the piezoelectric stage to be utilized.

### 3.6.3 Z-Mirror

Next, it is necessary to calibrate the relative movement of the Z-plane mirror(s) with respect to the movement of the now-calibrated Piezoelectric stage. This can be done using a calibration procedure provided by  $\mu\text{SPIM}$  Toolset which steps the stage through a user-defined number of levels across a volume of choice and allows the user to adjust the Z mirror voltage for each level such that the laser sheet is in focus. Finally, a Z Mirror scale is generated which can then be used to adjust settings.

### 3.6.4 X-mirror lag compensation

The X Mirror movement lags the signal which can be problematic when using laser masks (discussed in Section 3.7.1). To correct for this,  $\mu\text{SPIM}$  Toolset provides a simple calibration process during which a distinct mask is displayed and the mirror lag is adjusted by the user until the masks perfectly overlap (Fig. 3.5).

## 3.7 Data acquisition using $\mu\text{SPIM}$ Toolset

For the acquisition of a single plane, the X mirror galvanometer has to be scanned at least once across the full range of the desired light sheet width and back to its original

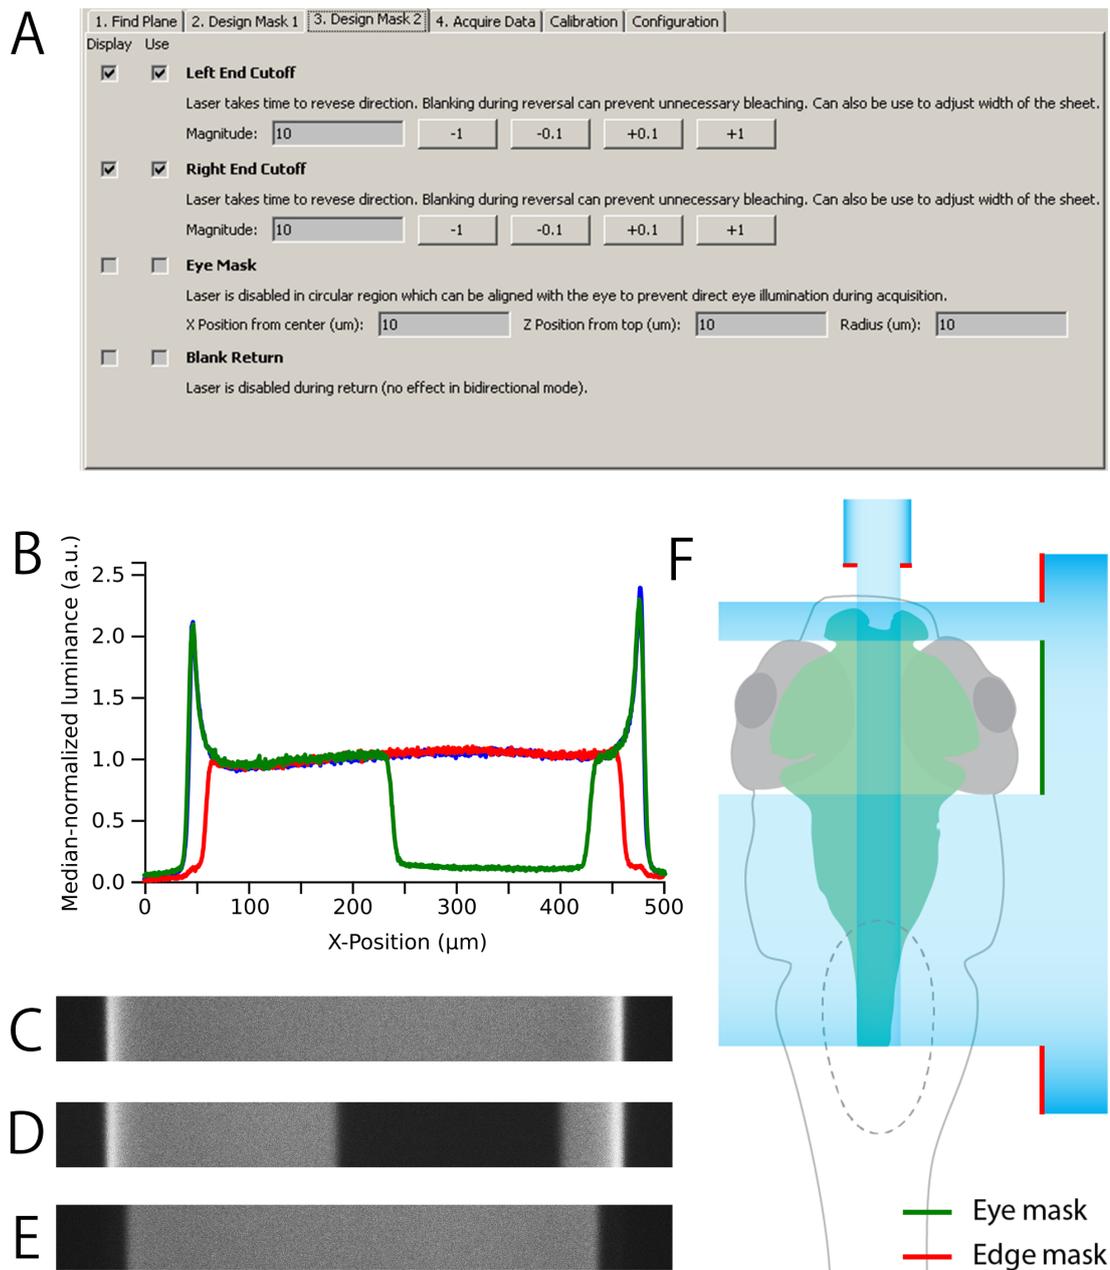


Figure 3.6: **Light-Sheet Laser Masks:** **A.** User Interface used to define laser masks allowing the user to display and use a number of masks for recordings: Edge masks help eliminate excess light during laser return as visible in **D** (with Edge masks enabled) compared to **C** (without any masks), Eye Mask allowing to turn the laser off in a specific region (such as eye of larval zebrafish) and Blank Return which turns the laser off during return (flyback), avoiding unnecessary photobleaching of the sample. **B.** The median-normalized luminance intensity profiles of a light sheet in Fluorescein with no mask applied (blue, shown in **C.**), eye mask (red, shown in **E.**) and edge masks (green, shown in **D.**). **F.** An illustration of how the 'edge' and 'eye' masks could be used in the context of imaging the brain of a larval zebrafish while minimizing illumination of the retina.

position (Fig. 3.2, showing 3 scans per frame). To ensure uniform illumination between frames,  $\mu$ SPIM Toolset takes control over the camera using a rising edge trigger signal synchronized with the X mirror (Fig. 3.2). Generation of a volume recording is provided by a synchronous movement of the Z Mirror and the Piezoelectric Stage between frames, acquiring consecutive frames from different heights in the samples (Fig. 3.2).

$\mu$ SPIM Toolset implements three volume acquisition modes: Scan Down, Scan Up and Bidirectional mode. In the Scan Down mode, the volume frames are acquired starting from the top, moving down through the volume and then finally returning to the top of the volume to prepare for the acquisition of the next volume. Since the stage moves a relatively heavy objective, instant movement from the bottom to the top of the volume is not possible and fast movement can introduce oscillations in the stage's position, reducing the quality of the recording. To eliminate this issue, a flyback period is introduced at the end of volume acquisition (shown in Fig. 3.2B v.), slowing down the return of the stage at the expense of several frames. The Scan Up mode offers similar functionality to that of Scan Down mode, however in the opposite direction. Bidirectional mode eliminates the need for the flyback period by alternating between Scan Down and Scan Up modes on consecutive volume acquisitions (Fig. 3.2B vi.) at the expense of unequal time period between frame acquisitions.

To allow for better distribution of sample illumination or illumination of occluded regions,  $\mu$ SPIM Toolset also supports the control of two separate paths, allowing for two light sheets of different widths.

### 3.7.1 Advanced acquisition methods

To allow for better distribution of sample illumination or illumination of occluded regions,  $\mu$ SPIM Toolset also supports the control of two separate paths, allowing for two light sheets of different widths.

#### Laser Masks

Removing parts of the light sheet can be beneficial in some acquisition settings.  $\mu$ SPIM Toolset utilizes fast software shuttering to provide two main masks which may be beneficial for in vivo imaging: Edge Masks, removing the ends of the light sheet and thus reducing the regions which are over-illuminated due to slow reversion of the direction of the X mirror (Fig. 3.6) and Eye Mask which can be used to remove illumination from a selected region in the light sheet (Fig. 3.6), allowing for acquisition where a specific region must not be

illuminated, such as the eye (Fig. 3.6F). Both masks are easily adjusted (Fig. 3.6A) and provide further flexibility in designing illumination sheets.

### Light Adaptation

Zebrafish larvae have been shown to react to sudden light changes (Burgess and Granato, 2007), interfering with behaviours of interest. This proves problematic in visible light recordings as the light is commonly turned off between recordings to prevent photodamage to the sample, resulting in a sudden light change during the start of the recording. We provide a simple light adaptation functionality which places a light sheet in a position near the top of the sample prior to and in between recordings, where photodamage is of less concern and then moves the sheet into the defined position when the recording starts.

### Multi-recording sequence

Similarly, to the adaptation period, sudden light changes can be reduced by consecutive acquisition of a number of recordings without the need for user input in between the recordings. This eliminates the concern of photodamage during non-acquisition time periods, allowing the laser to be turned on for the whole sequence of the recordings.

#### 3.7.2 Data output format

The image data output follows the MicroManager format with acquired image sequences being split into 4GB files accompanied by a metadata file (Metadata format can be found in the official MicroManager documentation: [https://micro-manager.org/wiki/Files\\_and\\_Metadata](https://micro-manager.org/wiki/Files_and_Metadata)). To accommodate for the extra information regarding the light sheet acquisition,  $\mu$ SPIM Toolset provides a separate metadata file which carries the information about the volume acquired as well as information about the experiment and sample provided by the user (Fig. 3.7).

## 3.8 Discussion

We have presented the  $\mu$ SPIM Toolset which provides a flexible software solution for the control of SPIM microscopes and demonstrated its utility for brain-wide imaging of neural activity in larval zebrafish (Ahrens et al., 2012). In contrast to other open source control solutions for light sheet microscopes, such as OpenSPIM (Pitrone et al., 2013) and Open SPIM microscopy (Gualda et al., 2013),  $\mu$ SPIM Toolset focuses on a microscope implementation better suited for functional imaging (as opposed to developmental imaging),

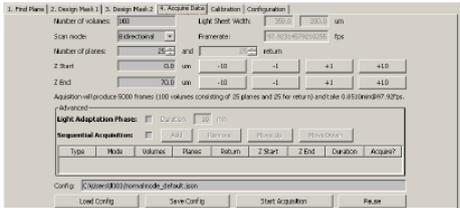
**A**

```

1 {
2   "version": "1.0.1",
3   "id": "5cb72daac47c0b1578c69192",
4   "datetime": "2019-04-17T13:44:31.215",
5   "sample": {
6     "dpf": 6,
7     "id": 147,
8     "expression": "HuC:H2B-GCaMP6f",
9     "description": "Sample Description"
10  },
11  "recording": {
12    "volume": {
13      "start": 0.0,
14      "end": 100.0,
15      "planes": 50,
16      "flyback": 10,
17      "n": 1000
18    },
19    "name": "RecordingName",
20    "description": "Recording Description",
21    "fps": 97.92314579210255,
22    "frames": [
23      {
24        "frameId": 0,
25        "file": "2019..._Pos0.ome.tif"
26      }, ...
27    ],
28    "scanMode": "scanDown",
29    "files": [ "2019..._Pos0.ome.tif", ... ]
30  },
31  "experiment": {
32    "description": "Experiment Description",
33    "tags": [ "plane", ... ],
34    "type": "Plane"
35  }
36 }

```

**B**



**C**

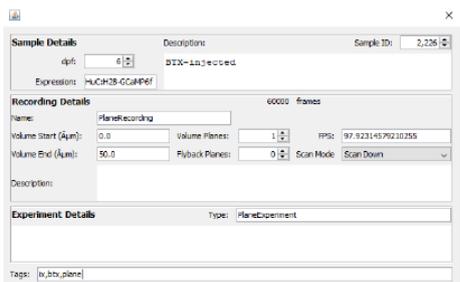


Figure 3.7: **Acquisition using  $\mu$ SPIM Toolset and Output Data Format: A.**  $\mu$ SPIM Toolset generates a separate metadata file, supplementing metadata supplied by the MicroManager platform. The generated data contains information detailing the settings of the acquisition (such as the height of the recorded plane or number of planes in the recorded volume) based on the parameters set by the user in **B.**. This information can be supplemented by the user to contain information about the sample imaged, comments about the experiment protocol as well as the recording itself to aid cataloguing of the data in acquisition information dialog shown in **C.**

filling a gap in available software solutions rather than competing with the existing implementations. With a range of in-built calibration protocols,  $\mu$ SPIM Toolset allows the user to select hardware based on the needs of the given application. The MicroManager platform on which  $\mu$ SPIM Toolset is built provides support for a wide range of acquisition hardware which will provide users with a range of options in customizing their instrument. Comprehensive documentation and the open source nature of the toolbox allows the user to adapt the software to advanced application and non-standard use-cases, overcoming the current limitations of the software.

While our framework offers great flexibility in the choice of hardware used, and is modular allowing incomplete implementations,  $\mu$ SPIM Toolset fundamentally depends on the National Instruments DAC and other digital-analogue converters are currently not supported. Adding support for other DAC devices is supported by the structure of the source code of our control software, but the modifications necessary may be quite extensive.

As previously outlined, the focus of  $\mu$ SPIM Toolset is primarily on the hardware flexibility aimed at a particular application (functional light sheet imaging) and thus some more complex functionality may not be supported out of the box. One notable example would be the use of multiple lasers with various wavelengths which is required when imaging with multiple fluorescent reporters or markers:  $\mu$ SPIM Toolset is currently limited to providing support for two independent lasers. We have not yet explored the possibility of replacing one or either of these by a multi-colour light engine, but envisage that this will be possible given that micromanager supports a number of such devices. Again, advanced users may choose to modify our solution to support this functionality if required.

Being built around MicroManager,  $\mu$ SPIM Toolset inherits all of its limitations. For instance, while the list of acquisition hardware supported by the MicroManager platform is extensive, it is possible that some less common devices are not supported by the platform and thus cannot be used with our toolbox. Another notable limitation is the lack of support for the “rolling shutter” mode of acquisition in which the exposure of each line of pixels on the sCMOS sensor is delayed to coincide with the time at which those pixels are parfocal with the laser beam passing through the sample (also called electronic confocal slit detection (Hu et al., 2017)). Because only that single line of pixels is activated, much background fluorescence caused by scattered illumination of neighbouring areas is rejected, thereby improving resolution. We did not implement this mode of acquisition as a standard feature of  $\mu$ SPIM because it results in lower frame rates compared to the usual “global shutter” mode of acquisition in which all pixels are activated simultaneously.

While our framework does not support some use cases in its current implementation, it provides a comprehensive solution to the control of light sheet microscopes with support for a wide range of both control and acquisition hardware while retaining a gentle learning curve supported by the well documented calibration and acquisition protocols. For this reason, we hope the  $\mu$ SPIM Toolset will be adopted in both standard and custom light sheet imaging use cases thanks to the well-documented open source code base. Both documentation and source code are available for free from the  $\mu$ SPIM Toolset git repository (<https://uspim.org>).

## 3.9 Materials and Methods

### 3.9.1 Optical design of SPIM used in this work

A partial parts list for the SPIM we constructed is provided in Table 3.1. The illumination arm (Fig. 3.1) was fed by one or two lasers through optical fibres (kineflex). When using both red and green fluorophores, a long-pass dichroic mirror (Thorlabs DMLP505R) combined the beams of the blue and yellow lasers, which then passed through a mechanical shutter. The beam was reflected by a pair of X and Z Galvo scanning mirrors with 1 kHz bandwidth at deflection angles  $\pm 0.2^\circ$ . The beam was then expanded by a factor of 2.5 using a pair of achromatic lenses ( $f = 50$  mm and 125 mm) The optimum beam parameters were calculated using calctool ([http://www.calctool.org/CALC/phys/optics/f\\_NA](http://www.calctool.org/CALC/phys/optics/f_NA)). The illumination lens was  $f = 40$  mm achromat. The whole of the illumination arm, beginning with the combining mirror, was mounted on an optical rail which was itself mounted on two translation stages (Thorlabs XR25C/M), one of which was slaved to the other. This arrangement allowed centering of the illumination beam on the field of view when the x-mirror was set to its position of zero offset, which is especially important for calibration of x-mirror displacement described above. With the rapid scanning of the X Galvo mirror, the laser line will form an illumination plane over the integration time of a single frame. The thickness of the beam waist was  $\approx 7$   $\mu\text{m}$ , with a Rayleigh length of 440  $\mu\text{m}$ . We therefore achieved a relatively uniform plane of illumination across half of the total field of view, which was 810  $\mu\text{m}$  wide. The specimen stage was custom-designed according to application and manufactured using a 3D printer. The stage was attached to an x-y-z translation assembly (Thorlabs PT3) for positioning of the specimen. For detection, a 16X/0.8NA Nikon CFI LWD Plan Fluorite objective was placed perpendicularly to the illumination plane to collect the emitted fluorescence signal. The excitation light was rejected by the

emission filter and then a tube lens of 200 mm focal length (AC508-200-A-ML, Thorlabs, Inc) used to project an image onto the sensor of the Hamamatsu Flash4.0 sCMOS camera (13 x 13 mm CCD, so that each pixel imaged an area of  $0.16 \mu\text{m}^2$  using a 16x objective). The objective lens was mounted on a piezoelectric stage (P-721 PIFOC High-Precision Objective Scanner, Physik Instrumente Ltd) and its movement synchronized with the Z Galvo scanner to make sure the illumination plane is always in the imaging focal plane. The objective and PIFOC scanner were themselves mounted on a manual stage for finding the focal plane before starting acquisition.

### 3.9.2 Characterising lateral and axial resolution of the microscope

To characterise the lateral (x/y) and axial (z) optical resolution (defined as the minimum distance the diffraction images of two points can approach each other and still be resolved is the resolution limit) of our SPIM we used fluorescent sub-diffraction size beads (0.1  $\mu\text{m}$  diameter “fluospheres”, Invitrogen), suspended in 1% low melting point agarose (Sigma-Aldrich). Using a 200  $\mu\text{m}$  wide light sheet and a step-size of 0.5  $\mu\text{m}$  several volumes of the suspended beads were acquired. 22 non-overlapping beads were used to determine lateral and axial resolution. Axial resolution was determined by fitting a Gaussian function to the z- profile running through the point of highest intensity (the centre of the bead). The position of the peak of that function was used to determine the imaging plane running through the centre of the bead and a 2D Gaussian where  $\sigma(X) = \sigma(Z)$  was fit to that plane. Using this procedure, we determined a lateral resolution of  $0.7 \pm 0.1 \mu\text{m}$  (mean  $\pm$  sd, full width at half maximum, n = 22 beads) and an axial resolution of  $5.4 \pm 0.5 \mu\text{m}$ .

### 3.9.3 Imaging Brains of Larval Zebrafish

All procedures were in accordance with the UK Animal Act 1986 and were approved by the Home Office and the University of Sussex Ethical Review Committee. Transgenic zebrafish larvae of undetermined sex expressing the nuclear localised H2B-GCaMP6f calcium reporter panneuronally (Tg(elavl3:H2B-GCaMP6f), [Chen et al. \(2013\)](#); [Dunn et al. \(2016b\)](#)), were imaged at 6-7 days-post- fertilisation (dpf) as a part of an ongoing investigation into Optomotor Response in larval zebrafish. All lines were maintained in the nacre background to limit pigmentation. Larvae were paralysed by positioning them sideways in a small slit of PDMS (Sylgard184, Dow Corning) on a coverslip ([Pichler and Lagnado, 2019](#)) and injecting 0.25mM  $\alpha$ -Bungarotoxin (Tocris Bioscience) into their heart. They were then embedded dorsal side up in 2% low-melting-point agarose (Biogene) in

<b>Mechanical Parts</b>	<b>Component (Manufacturer)</b>
Optical table	1.4 x 1.8 x 0.2 m clean top on Micro-g Pneumatic and Rigid Legs (TMC Vibration Control)
Rail systems	Precision 100 mm Dovetail Optical Rails with PRC carriers (Newport)
Translation stages	XR25C/M (Thorlabs)
<b>Illumination arm</b>	
Lasers	LuxX+ 488 nm 60 mw (Omicron) Jive 561 nm 100 mW (Cobolt)
Optical Fibre and laser launch	Kineflex Fibre system (0.7 mm), 400-640nm, 2m, FC/APC connectors (Qioptiq)
Mechanical Shutter	SH05/M shutter and SC10 controller (Thorlabs)
Beam expander 2.5x	f = 50 mm achromat (AC254-050-A, Thorlabs) and f = 125 mm achromat (Edmund Optic)
Dual-axis galvanometer scanner	GVS102 - 2D Galvo System (ThorLabs)
Illumination lens	f = 40 mm achromat (#49-354-INK, Edmund Optics)
<b>Collection arm</b>	
Camera	ORCA-flash4.0 v2
Laser cleanup notch filters	488/10, and 561/10 (Chroma)
Detection objective	16X/0.8NA CFI LWD Plan Fluorite (Nikon)
Piezo	PIFOC E-665 amplifier/controller + P-725 (400 $\mu$ m travel range from Physike Instrumente)
Emission filter	535/35 (xf3007, Omega)
Tube lens	f= 200 mm achromat (AC508-200, Thorlabs)
Mounting of collection arm	X-95 mounting system and carriers (Linos)
<b>Computing</b>	
Custom built PC incorporating Intel®	
Xeon® processor E5-2600	
Frame Grabber	Firebird Camera Link (2xCLM-2PE8) (Active Silicon)
DAC card	PCIe-6738 (National Instruments)

Table 3.1: **Parts list for construction of SPIM**

E2 medium (Brand et al., 2002) on a 22 mm x 22 mm coverslip and placed into a custom glass-walled 3D printed chamber. After filling the chamber with E2 medium, the majority of the agarose that fell between the laser source and the larva's head was removed with a scalpel to maintain sufficient stability of the sample while minimising the distance the laser is travelling through the agarose. The chamber was positioned underneath the objective so that upon turning on the laser, it would not hit the head, and the eyes in particular, but rather the tail of the fish. Once the imaging and excitation planes were successfully aligned, the larva's brain and in particular the desired brain region was approached. Finally, experimental parameters such as sheet widths, frame rate, number of volumes, step-size and adaptation time were defined before starting acquisition. Volumes were acquired using z-steps of 2  $\mu\text{m}$ , which was sufficient to identify ROIs corresponding to individual neuronal cell bodies (typically 5-8  $\mu\text{m}$  in diameter) without an unduly high data acquisition load. Once acquired, data was reduced to time-series of fluorescence signals for each ROI corresponding to a single neuron. Operation of  $\mu\text{SPIM}$  Toolset was tested using more than 10 zebrafish larvae as part of another ongoing investigation.

## Chapter 4

# A model-based closed loop virtual reality system based on free-swimming parameters

### 4.1 Introduction

Neural processes giving rise to behaviour have been of increasing interest in the larval zebrafish model thanks to the optical accessibility of brain, availability of genetic tools and the number of robust behaviours the larvae exhibit (Orger and de Polavieja, 2017). Despite this, a number of challenges exist when acquiring the neural activity data from behaving animals. It has been extensively shown that environmental feedback is necessary for the animal's behaviour to accurately reflect that which would have been expressed in its natural setting (Ahrens et al., 2012; Portugues and Engert, 2011; Severi et al., 2014; Markov et al., 2020). This makes freely moving setting favourable for behavioural studies, however it proves to be difficult when neural activity is of interest as the brain is not fixed, making it hard to be imaged (but see Kim et al. (2017) and Marques et al. (2020)). Moreover, the physical environment offers limited options for perturbations and control of the sensory feedback, narrowing the opportunities of exploration for the embodiment of the animal in the real world and its effect on behaviour. For these reasons, many studies have adopted a fixed animal virtual reality approach for the investigation of animal behaviour, especially in conjunction with neural imaging, such as for the purposes of opto-motor reflex (Ahrens et al., 2012) or prey capture (Bianco et al., 2011; Semmelhack et al., 2014) assays.

In many studies, an open-loop virtual reality (VR) environment has been sufficient (Severi

et al., 2014; Bianco et al., 2011; Semmelhack et al., 2014), however as the importance of environmental feedback became more apparent (Ahrens et al., 2012), many studies have focused on the effects of feedback in a head-restrained closed-loop virtual reality setting. This modality has a number of advantages in comparison to free swimming, mainly in terms of accessibility for simultaneous imaging of the brain with 2P (Naumann et al., 2016) or light sheet (Vladimirov et al., 2014) microscopy and the possibility of perturbation to the feedback (Ahrens et al., 2012).

Existing studies utilizing virtual reality with recording activity rely either on electrode recording from the tail (Vladimirov et al., 2014) or visual capture of tail movement to estimate the movement of the fish (Severi et al., 2014) to construct the feedback signal. Here we aim to provide an accurate feedback to the fish based on the recording of tail movement since the interaction between the larva and its environment is of a particular interest.

We explore the movement in freely behaving larvae and formulate a realistic feedback model based on the observed behaviour. This is advantageous as the derived model is both physically plausible and accurately reflects the properties of the swimming motion in freely swimming fish. This model is then implemented in a head-fixed setting by recording the activity of the tail and delivering the stimulus to the sample by an independent stimulus/acquisition arm combining both the stimulation and the necessary tail tracking, allowing the VR stimulator to be incorporated into existing acquisition setups, irrespective of present setup.

## 4.2 Methods

The basic operation of a closed loop stimulator relies on a continuous and consistent execution and interaction of multiple elements, including camera acquisition of the tail activity from the sample, tail tracking, integration of the swim model and subsequent presentation of the computed stimulus back to the fish larva (Fig. 4.1A). This results in a higher complexity when compared to open loop stimulation as inconsistency in any of the components will result in inconsistent feedback to the fish, potentially invalidating the results of the experiment. Here we show the implementation of the individual components as well as their interaction to produce the closed loop stimulus.

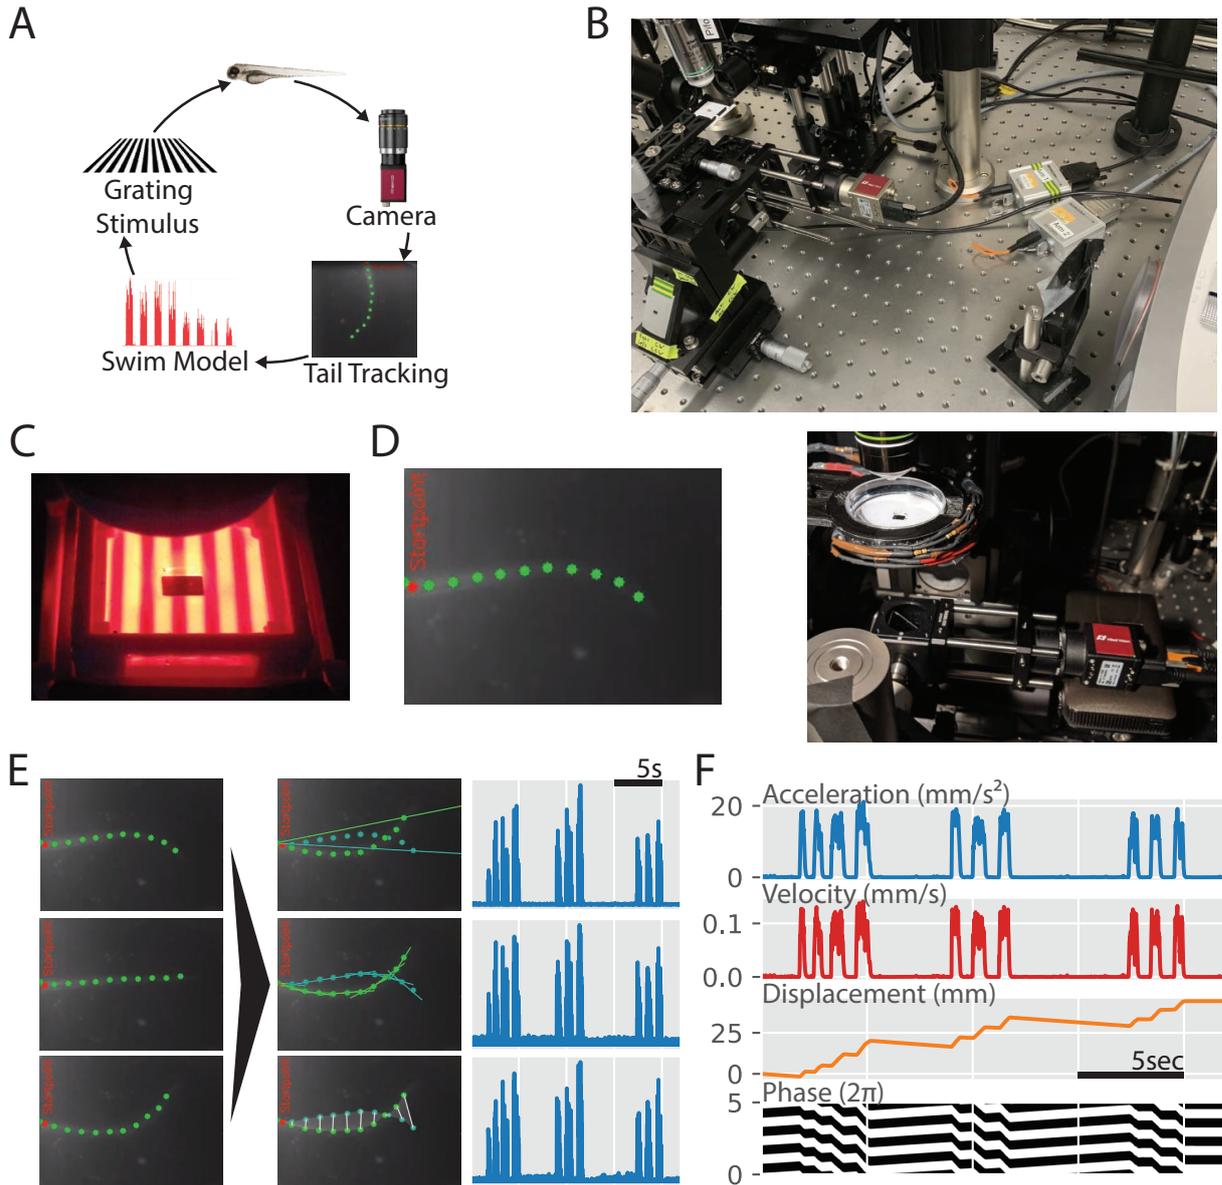


Figure 4.1: **Closed Loop Virtual Reality Setup:**

**A.** Diagram of the closed-loop virtual reality setup based on camera acquisition of larval tail movement. **B.** Photos of an example stimulation setups. Top: VR setup used with light-sheet microscopy with IR illumination from above. Bottom: VR setup used with two-photon imaging, illumination provided by IR diodes positioned in a ring around the sample. **C.** Photo of the stimulation chamber showing projection screen with acquisition window for tail movement acquisition. **D.** Tail image acquired by the camera fitted with points tracking the tail in the frame. **E.** Tail movement is tracked over time (left) and is converted to variables approximating the tail activity based on tail angle change (top), angle between each pair of consecutive points fitted to the tail (center) and area covered by the tail in consecutive frames acquired(bottom). Traces on the right show the respective measures of tail movement on the y-axis (a.u.) tracked over time. **F.** Tail activity measures shown in E are converted into acceleration which is converted into velocity through a simple integration model. Velocity is then integrated over time into displacement. This model is then used to infer the phase of the stimulus.

### 4.2.1 Tail Camera Acquisition

We facilitate the acquisition of the tail movement by a camera mounted on an arm positioned below the sample (Fig. 4.1B). This allows seamless integration with 2P or light sheet acquisition as the sample is not occluded from the top nor from the sides. Since the stimulus is also projected from the bottom of the sample (Fig. 4.1B), a small window in the projection screen is required to expose the tail for acquisition (Fig. 4.1C). For the acquisition itself, we chose to use Allied Vision Mako U-130B as it offers high acquisition speeds ( $>200\text{Hz}$ ) sufficient to capture short, high-frequency, tail bouts in sufficient detail and is small enough to be easily mounted on a horizontal railing without additional support. To avoid light pollution from and to the acquisition and from the varying light stimulus, an 850/10 nm band-pass filter has been included in the VR acquisition path and an external IR light source (850nm) has been positioned above the sample (Fig. 4.1B) in case of the light sheet microscope setup and in a ring configuration in the 2P microscope setup (Fig. 4.1B, bottom). Green light band-pass filter has been included in the acquisition arm of the 2P and light sheet setup allowing avoidance of light pollution from the tail illumination and thus preserving the quality of the neural activity acquisition.

### 4.2.2 Tail Tracking

A number of tracking solutions have already been formulated and used in many past investigations of larval behaviour (Portugues and Engert, 2011; Portugues et al., 2015; Semmelhack et al., 2014). We chose to use a tail tracking algorithm formulated by Semmelhack et al. (2014) as it offers both sufficient accuracy and high processing speed of up to 3ms per frame (300Hz), making it ideal for a real-time application with high acquisition speeds. This algorithm outputs a present number of points equally spaced along the tracked tail which is illustrated in Fig. 4.1D.

### 4.2.3 Swimming Model

As a basis for the swimming model, we use a basic integration model of motion where the fish has control over its acceleration through some measure of effort derived from the movement of the tail:

$$\begin{aligned}\frac{da}{dt} &= -\frac{a}{\tau_a} + \mathcal{N} * T \\ \frac{dv}{dt} &= -\frac{v}{\tau_v} + a \\ \frac{dd}{dt} &= v\end{aligned}\tag{4.1}$$

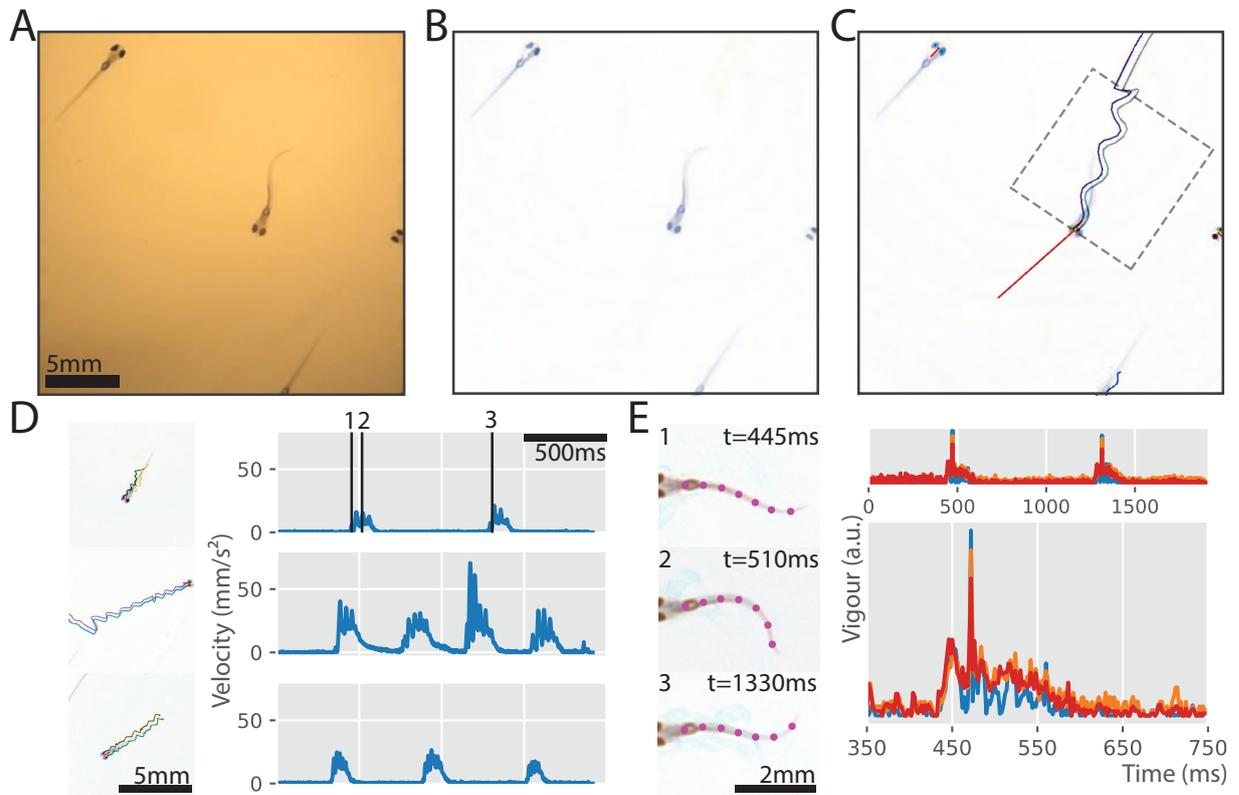


Figure 4.2: **Processing Data Acquired from Freely Moving Fish:**

**A.** Image excerpt acquired by a high speed camera at 500 frames per second. **B.** Image from **A.** with background subtracted. **C.** Tracking parameters shown on a processed frame shown in **B.** Purple and green traces show the position of larva eyes over time, red line indicates the velocity vector and dashed rectangle outlines an eye-aligned excerpt used for tail fitting shown in **E.** **D.** Larval fish can be tracked over the period of two seconds, three examples are shown with the tracked position (left) and forward velocity (right). **E.** An rectangle aligned with the eyes can be extracted from the recordings, emulating input from a head-embedded sample for tail fitting (left, showing tail at times 1, 2 and 3 indicated in **D.**). The tail activity variables can then be estimated for simple tail angle change (blue), pair-wise tail angle change (orange) and area covered (red). Figure on the right tracks the variables for the first trace shown in **D.** over the full recording (top) and first bout (bottom).

where  $T$  is a measure of tail movement (swimming effort),  $\mathcal{N}$  is a *normalizer* variable used to normalize the tail effort across fish,  $a$  is acceleration in mm/s<sup>2</sup>,  $v$  is velocity in mm/s and  $d$  is displacement in mm.  $\tau_a$  and  $\tau_v$  are decay constants for acceleration and velocity, respectively. This representation of the motion model allows us to capture the instantaneous acceleration, velocity and position of the larval sample at any given point in time and to log it for further analysis (Fig. 4.1F).

We formulated 3 models to approximate the swim effort  $T$  as a function of the points fitted along the tail (Fig. 4.1E). The first model is based on the deflection angle of the tail in the successive image acquisitions. This is computed as the angle between a vector from the first to the last tail point in the current and the previous acquired frame:

$$T_\alpha \propto \tan^{-1} \left( \frac{y_t^N - y_t^1}{x_t^N - x_t^1} \right) - \tan^{-1} \left( \frac{y_{t-1}^N - y_{t-1}^1}{x_{t-1}^N - x_{t-1}^1} \right) \quad (4.2)$$

where  $x_t^n$  and  $y_t^n$  are the  $x$  and  $y$  coordinates of a  $n$ -th tail point at a time  $t$  and  $N$  is the number of tail points captured.

The second model is based on the derivative of the tail, approximated as the change in angle of a vector formed by a pair of two consecutive tail points summed over all pairs fitted:

$$T_d \propto \sum_{i=2}^N \tan^{-1} \left( \frac{y_t^i - y_t^{i-1}}{x_t^i - x_t^{i-1}} \right) - \tan^{-1} \left( \frac{y_{t-1}^i - y_{t-1}^{i-1}}{x_{t-1}^i - x_{t-1}^{i-1}} \right) \quad (4.3)$$

The third model is loosely based on the area covered by the tail and is formulated in terms of the average distance covered by each tail point over time:

$$T_A \propto \frac{1}{N} \sum_{i=1}^N \sqrt{(x_t^i - x_{t-1}^i)^2 + (y_t^i - y_{t-1}^i)^2} \quad (4.4)$$

## Model Parameter Fitting

We evaluate all three models in freely swimming zebrafish larvae where both the tail movement and corresponding change in position are observable. To do this, we recorded 7 dpf larvae under a microscope for short periods of time with a 500Hz high-speed camera (Fig. 4.2A). Subtracting the average background of the recording then yielded a good basis for tracking and analysis (Fig. 4.2B).

To extract the orientation, position and other kinetic variables, we used blob detection to track the eyes of the zebrafish over time (Fig. 4.2C,D). An eye-aligned square region was then used to produce a pseudo-head-fixed view (Fig. 4.2C, dotted rectangle) which could then be used for tail tracking using the algorithm outlined in Section 4.2.2 (Fig. 4.2E, left). *Swimming effort* could then be estimated using Eq. 4.2, 4.3 and 4.4.

Using the estimated swimming effort, we can use Eq. 4.1 to model the velocity of the fish where  $\mathcal{N}$ ,  $\tau_a$  and  $\tau_v$  are model variables that need to be fitted. Knowing the velocity of the fish over time (Fig. 4.2D, right), these variables can be estimated by minimizing the mean squared error between the model and real velocity. Fig. 4.3 A, B and C show results for fitted models using swimming effort based on tail angle, tail derivative and tail area, respectively. The example fitted trace on the left shows the model fitted on 500Hz acquisition data in light colour and the model applied on to sub-sampled recordings with 250Hz, 100Hz and 50Hz rate in darker colours. The figure on the right then details the change in the mean squared error over the different acquisition rates, indicating the robustness of the model to sub-optimal or varying frame rate.

During the parameter fitting, we found  $\tau_a$  to tend towards 0 as the model needed to produce a fast rise in velocity at the start of a swimming bout. We therefore fixed  $\tau_a = 0.001$  which is both fast and still possible to reliably integrate in real time.

#### 4.2.4 Stimulus Projection

The stimulus is projected by the bottom-mounted VR arm (see Section 4.2.1) on a 3D-printed translucent screen by an Optoma Pico Projector PK320 or BenQ W121OST. We found the 3D-printed screen to offer sufficient stimulus brightness and contrast (Fig. 4.1C is representative) while being resistant to liquid spills and simple to produce. It is easily replaceable by a paper- or film-based projection screen if necessary.

### 4.3 Results

#### 4.3.1 Realistic Swim Model

To evaluate the performance of the different *Swimming effort* metrics outlined in Section 4.2.3, we investigated both the fitted variable values as well as model metrics over a range of acquisition speeds.

Over the recorded samples (N=33), we plotted the value of  $\tau_v$  for each model fitted (Fig. 4.3D). We found no significant difference in the values of  $\tau_v$  between the models based on tail area covered (Eq. 4.4),  $0.021 \pm 0.0065$  (mean  $\pm$  1 STD), and tail derivative (Eq. 4.3),  $0.020 \pm 0.0061$ . The  $\tau_v$  values for models based on tail angle (Eq. 4.2),  $0.028 \pm 0.0081$ , differed significantly from both the tail area and derivative models ( $p < 0.0001$ , Wilcoxon

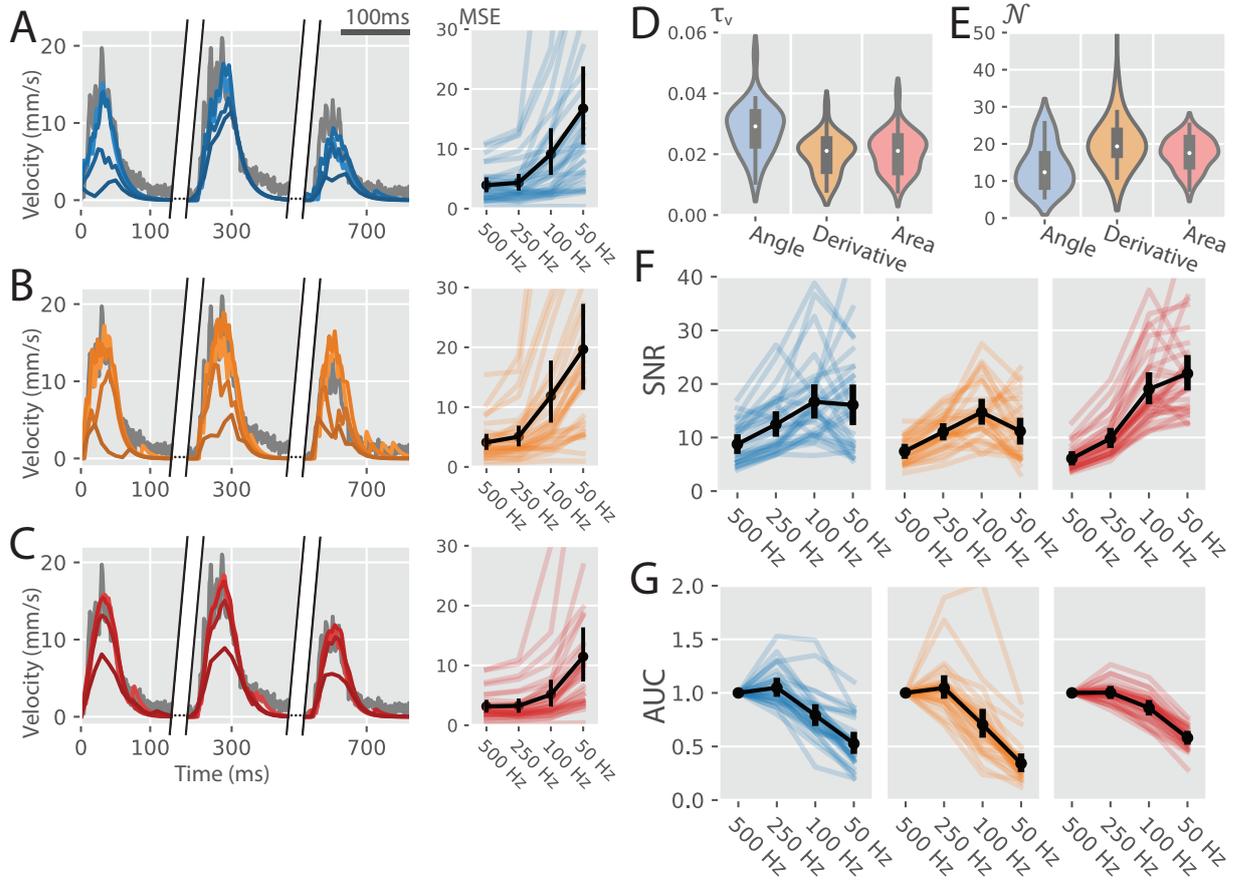


Figure 4.3: **Estimating Swim Model Parameters from Freely Moving Fish:**

**A.** Example fitting results (right) for a swimming velocity trace shown in gray (periods of no movements were omitted), fitted with a model based on tail angle change (light colour). Darker colours show the same model applied with acquisition rates of 250Hz, 100Hz and 50Hz. The mean squared error (MSE) for the fitted models applied with varied acquisition rates is shown on the right. Each coloured trace represents a model fitted on a single larval sample while the black trace represents the mean with error bars displaying the standard deviation. **B.** Results for fitting using the pair-wise angle change presented in the same format as F. **C.** Results for fitting using the area covered presented in the same format as F. **D.** Violin plots showing the parameter value distribution for  $\tau_v$  for tail angle change (left, blue), pair-wise tail angle change (center, orange) and area covered (red, right). **E.** Violin plots showing the parameter value distribution for  $\mathcal{N}$  for tail angle change (left, blue), pair-wise tail angle change (center, orange) and area covered (red, right). **F.** Signal-to-Noise ratio (SNR) at different acquisition rates for the tail activity variables; tail angle change (left, blue), pair-wise tail angle change (center, orange) and area covered (red, right). **G.** Area Under a Curve (AUC) at different acquisition rates for the tail activity variables; tail angle change (left, blue), pair-wise tail angle change (center, orange) and area covered (red, right).

signed-rank test<sup>1</sup>). All of the  $\tau_v$  for the different models were on the same order of magnitude as would be expected for a variable closely related to the kinetic model rather than the representation of swim vigour.

An important parameter of a model is its proneness to the noise from tail tracking. The tail tracking algorithm used does not take time history of the tail into account, operating on individual frames. This benefits from a lower bias and a faster processing speed due to the relative simplicity of the algorithm. A disadvantage of this approach, however, is a higher noise where some tail points may move by one or more pixels even if the tail itself is not moving. This can be very problematic in a low illumination setting, such as that of a setup with combined neural activity acquisition. To capture this, we approximated signal-to-noise ratio as the ratio of the average swim vigour,  $T$ , during bouts (tail moving) and outside bouts (tail not moving) over different acquisition rates (Fig. 4.3F). At 500Hz acquisition rate, the SNR for angle-based measure, 8.8, was significantly higher ( $p < 0.01$ , Wilcoxon signed-rank test) than both derivative-based measure, 7.4 and significantly worse performing area-based measure ( $p < 0.01$ , Wilcoxon signed-rank test) as small variation in the tail point position results in a small angle change in case of the angle angle-based measure whereas derivative- and area-based models accumulate the noise from the movement of all of the tail points. With slower acquisition rates, the change in the observed variables during tail movement becomes larger between subsequent frames whereas the noise remains constant. All three measures then benefit from the slower acquisition rates up to 100Hz where angle- and derivative-based measure perform comparatively (16.6 and 14.7 respectively) and area-based measure offers a significantly better signal-to-noise (19.0). The larger improvement in SNR for area measure can be attributed to a greater increase in the signal due to lesser susceptibility to underestimating the tail movement as the tail reverses direction. This is even more apparent at the 50Hz acquisition rate where only the area-based shows significant improvement in the SNR whereas angle-based measure shows no improvement and derivative-based measure performs worse than at 100Hz.

Zebrafish larvae are hypothesized to integrate their movement with respect to their environment over time (Portugues and Engert, 2011). This is not expressed in mean squared error as even with a significant MSE, the mean velocity could remain constant over time. We therefore looked at the area under the curve (AUC) for the velocity output

---

<sup>1</sup>In all instances, two-sided Wilcoxon signed-rank test was used to characterize the difference between the models. Some data did not appear to be normally distributed, limiting the applicability of regular paired t-test.

by the model over time (Fig. 4.3G). An ideal model would retain constant AUC, however we observed a significant decrease past 250Hz acquisition speed. The decrease was less prominent in the case of area-based measure down to 86.0%, followed by angle-based to 79.2% and finally derivative-based to 70.3%.

While basing the swimming vigour on the angle of the tail can yield superior results for high acquisition speeds where it benefits from higher signal-to-noise ratio, the use of area is significantly superior for slower acquisition speeds which are more viable for real-time closed-loop stimulation. This is mainly due to better retention of the information (signal) in the lower sampling rate setting, leading to better SNR and lesser underestimation of the distance swam by the fish (velocity AUC). The smaller variance of this effect with decreasing frame rate also makes it easier to be compensated for by increasing the normalizer parameter  $\mathcal{N}$ .

### 4.3.2 Tracking of kinetic variables

To provide a basis for a detailed analysis of the behavioural data, we provide granular swimming data logging at regular intervals (tested up to 200 data points per second). Along with the basic timestamped kinetic variables produced by the swim model, including acceleration, velocity and displacement (Fig. 4.4E ii-iv.), the stimulator also saves exact tail point position and change (Fig. 4.4E i.) as well as the state of the other components of the stimulator such as the visual stimulus state and the current frame count of accompanying acquisition (if any, explained below). All experiment-variables are saved as well, including the time since start of the current experiment, experiment parameters, experiment mode and the duration of individual stimulation modes (see Section 4.3.6).

### 4.3.3 Alignment with other acquisition modalities

The VR stimulator was designed to be used in parallel with other acquisition modalities such as 2P or light sheet microscopy acquisition of neural activity. The alignment of behavioural and neural activity is particularly important for many studies, allowing the combination of the data through correlation and other related analysis. To facilitate this, we implemented triggering and frame counting through a National Instruments acquisition board which is able to take signals from a different signal-generating device to establish common timing truth. This can be achieved through sending a single triggering signal at the start of the acquisition to the VR stimulator to attain simultaneous start of both recordings, using the relative time to reference events throughout both recordings. Altern-

atively, if the neural acquisition setup utilizes camera triggering to request the acquisition of images from the camera, the same signal can be supplied to the VR stimulator to obtain perfect frame-accurate alignment of the acquired behavioural data.

#### 4.3.4 Tail Tracking achieves real-time performance at high acquisition rates

We evaluated the performance of the tail fitting algorithm as well as the swim model computation to ensure that high data throughput can be achieved in real-time (Fig. 4.4C). Each execution of the swim model loop consistently finished under 1ms with 99.999th percentile of 50.4 $\mu$ s including processing of the tail points into *swim effort*, making it more than sufficient for handling integrative models on the millisecond time-scale. The tail fitting algorithm processed each frame consistently under 10ms with 99.999th percentile of 8.44ms, offering sufficient speed for processing tail acquisition with up to 100Hz speed in real-time.

#### 4.3.5 Support for open and closed loop Optomotor Response (OMR) stimulus

The main focus of the VR stimulator is to provide functionality focused on the optomotor response behaviour. We provide closed loop models outlined in Section 4.2.3 deriving the swimming activity from tail angle change, tail derivative change or area covered by the tail of the larva (Eq. 4.2-4.4) with default parameters based on our investigation (Section 4.3.1). However, the significant variance of the normalizer parameter,  $\mathcal{N}$  (Fig. 4.3E) suggests some variability of the fish larvae in terms of their physical properties, resulting in a significantly faster acceleration for some fish compared to others. To attenuate this issue and offer means to normalize across fish, we provide an on-line fitting protocol that can be carried out prior to the experiment on each new fish sample which alters the normalizer parameter such that a given fish perfectly compensates for a drift of a desired speed (by default, 10mm/s) similar to closed loop VRs developed in other studies (Portugues and Engert, 2011; Marques et al., 2020).

Closed loop stimulation suffers from an inherent limitation in terms of latency arising from the camera, computation and projector components. We measured the latency of our solution by measuring the time duration between presentation of a light flash and its subsequent appearance on the projector screen by a high-speed camera. With the hardware used, the setup achieved low latency of 22 ms which can be primarily attributed

to the high speed projector. In comparison, commonly used projectors can have latency of 65 ms or higher as the intended usecase of media consumption does not rely on fast response times. While this is a significant difference, it is important to note that the latency is still significant on the timescale of freely swimming bouts the zebrafish larvae exhibit (Fig. 4.4A). This could lead to behavioural differences between the freely swimming and head-fixed preparation settings.

Some experiments may benefit from the option to switch between closed and open loop modes for purposes of comparison, or do not require closed loop functionality at all. We provide an open-loop swim model that takes only the stimulus movement into account, disallowing the interaction of the fish larva with it's virtual environment. As such, this mode does not require the tail fitting or the camera input and can thus be operated without it if only stimulation is of interest. Naturally, this would mean that none of the behavioural data can be collected. By default, if the camera input is provided, the stimulator saves the tail movement data and a virtual acceleration and velocity based on the tail movement of the fish even if this is not then converted into visual feedback. This then allows for more comprehensive comparison between closed and open loop behaviour.

### 4.3.6 Highly modular experiment design

The VR stimulator has been designed with modularity in mind. To support this, it provides an experiment model comprising of simple stimulus blocks that can be assembled to produce any customized protocols (Fig. 4.4B) as well as extended for custom functionality that is not supported out of the box. A single execution consists of a battery of experiments where each experiment is parametrized by the user, this could be high and low gain in a gain adaptation experiment or a set of velocities in a drift experiment. Each experiment then consists of sub-protocols, such that each sub-protocol is parametrized by a single value, for example high or low gain in a gain adaptation experiment or a single velocity in a drift experiment. The sub-protocol then comprises of a fixed stimulus sequence which is generally the same across the sub-protocols in a given experiment. This could be a set length drifting grating at a set velocity followed by a set length rest period. The user is then able to create experiments based on pre-set sub-protocols, define their own or even define custom stimuli as needed.

As outlined in Section 4.3.5, open and closed loop swim models are predefined, however the user has the ability to define a completely custom swim model as well as use different swim models in different experiments or multiple swim models in different parts of a single

experiment.

Similarly, the predefined grating stimulus can be exchanged for a user-defined one, allowing the VR setup to be used outside of the default OMR paradigm.

### 4.3.7 Native multi-threading support

The integrative swim model requires an uninterrupted, consistent performance for optimal operation. Having significant delays between the model updates will result in less accurate approximation of the movement of the larva and it is therefore important for long-running tasks to be completed asynchronously. The VR stimulator therefore splits tasks into a number of separate threads that are mostly independent and only communicate results of their workloads as needed (Fig. 4.4D):

First the main thread (the user thread) is primarily controlled by the user and is responsible for initialization of other processes such as the camera interface, visual stimulator and the swim model. During the experiment, it handles the instrumentation of the other threads such as changes of swim model parameters or stimulus speed as well as communication between individual sub-processes.

The swim model thread solely handles the conversion of tail point data into swim effort and the integration of the swim model. The output kinetic variables are written into a shared memory where they can be accessed by the main thread. In this design, the tail point data is only written by the main thread and only consumed by the integrator thread and vice-versa for the kinetic variables, allowing uninterrupted operation of the integrator without the need for extensive memory-locking mechanisms.

The tail acquisition thread awaits incoming frames captured by the camera and handles the fitting of the tail points of each frame. This functionality itself is completed asynchronously, allowing the tail fitting of the previous frame to take place while the next frame is being captured by the camera. Processed tail points are sent to the main thread which writes them directly into the shared memory, ready for consumption by the swim model integration thread.

The stimulus thread primarily handles the rendering of the stimulus. Each execution of the main loop supplies the stimulus thread with the most recent position of the fish relative to its environment, potentially sending hundreds of updates per second. Since modern projection devices generally only support output of up to 60-120 frames per second, prior to rendering the next frame, all outdated updates are discarded and only the most recent position is used. The rendering is handled by PsychoPy (Peirce, 2009) and the thread

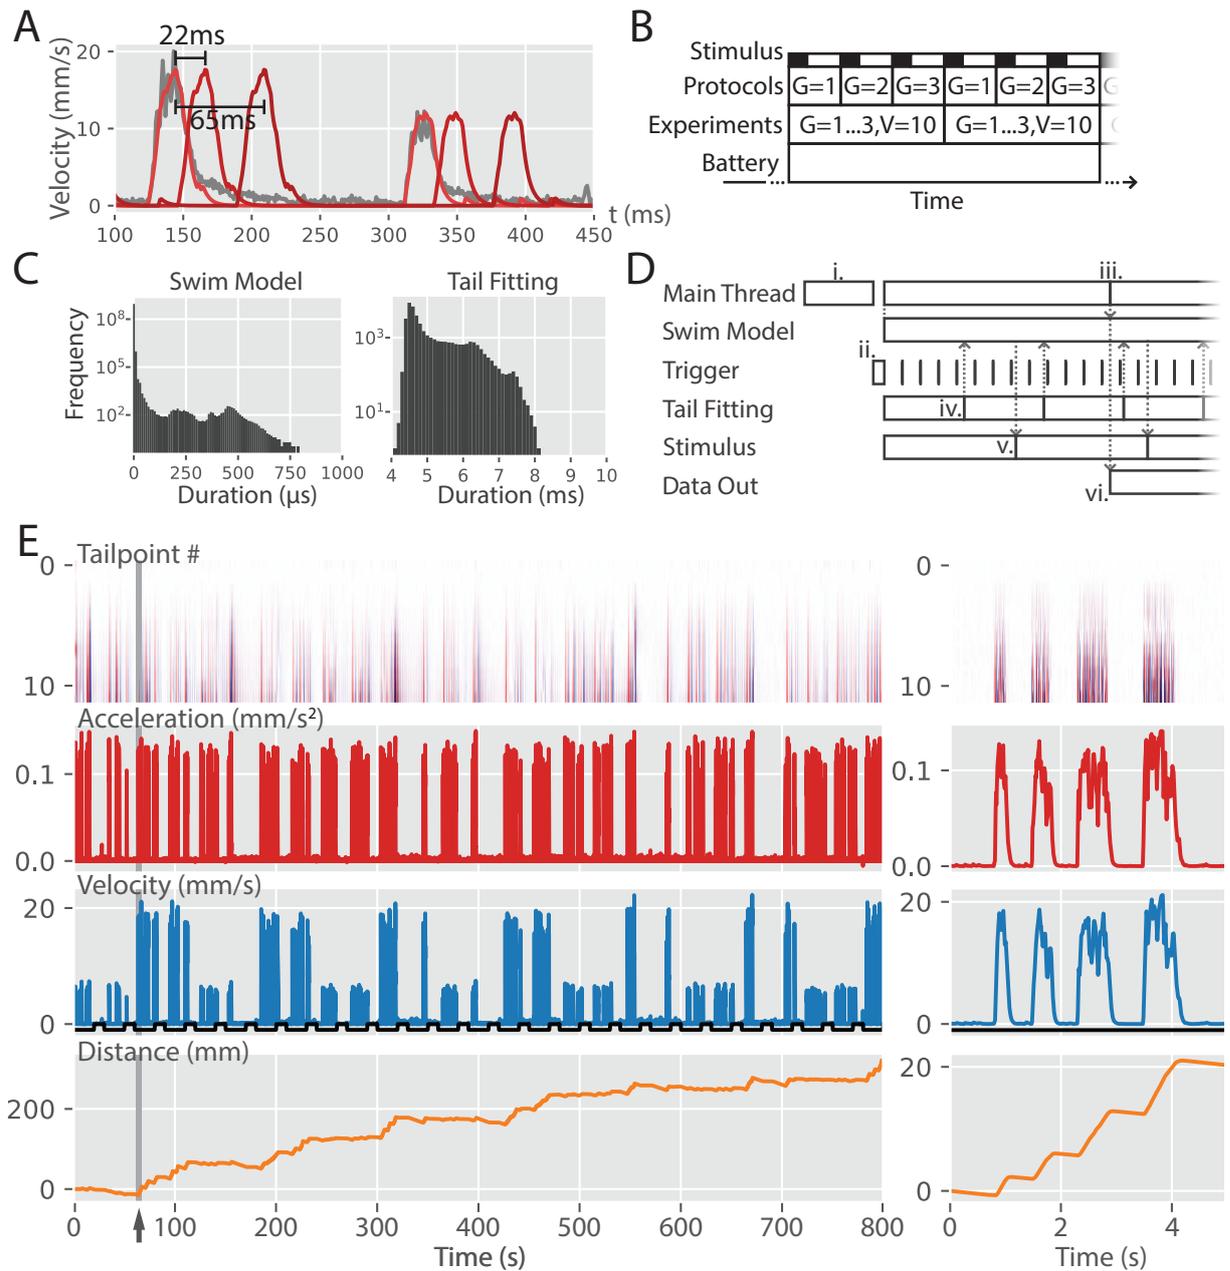


Figure 4.4: **Virtual Reality Stimulator Data Output:** **A.** Illustration of the effect of stimulator latency on the perceived velocity by the larval sample with 22ms and 65ms delay. **B.** The stimulator is structured with experiments in mind: stimulus sequences are grouped into simple protocols which are repeated within one experiment with varying parameters. Batteries of experiments can be run against any given animal sample in order to execute a number of experiments in succession without need for user input. **C.** The distribution of compute time required to complete computation of swim model (left) and tail fitting (right). **D.** The VR Stimulator utilizes multi-threading to ensure optimal performance. The main (user) thread prepares the stimulation settings (i.) and optionally awaits trigger from other acquisition devices to synchronize the stimulus with the acquisition (ii.) before setting the swim model parameters and starting the stimulation. The swim model parameters are changed by the main thread again only when required by the experimental setting (iii.).

Figure 4.4: **D. cont.** The swim model regularly receives information about the state of the tail as the frames acquired by tail camera are processed by the fitting algorithm (iv.) and updates the stimulation thread with the virtual position of the sample such that the stimulus presented accurately reflects the computed position (v.). The stimulation data is saved asynchronously such that it does not interfere with the stimulation performance (vi.). **E.** Detailed data of arbitrary stimulation duration is produced automatically with any acquisition. 13 minute excerpt (left) shows the tail point change over time, acceleration, velocity (sample in blue, grating drift in black) and position over time. Same parameters are shown again in detail over 5 seconds on the right.

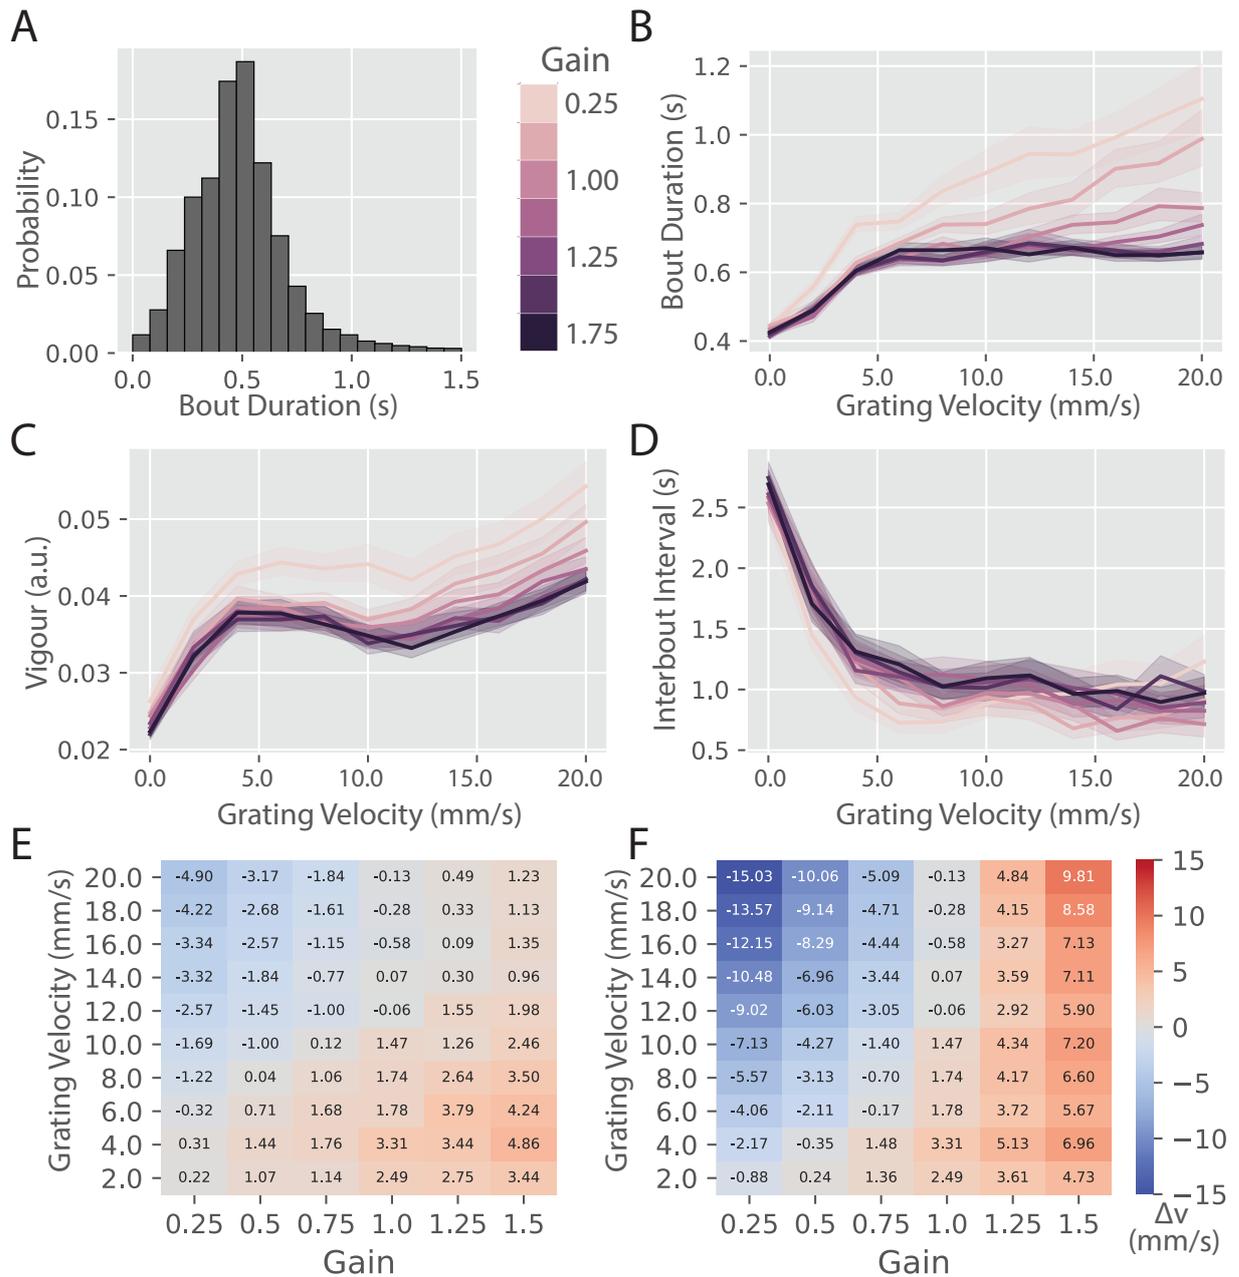
can be replicated multiple times with each instance handling one output. By default, one thread is created for the stimulator projector and a second one for a monitoring window displayed along with current state of the stimulation to the user.

Finally, the data logging thread handles writing the data output of the acquisition. Receiving data from main thread, it writes the data into a JSON file during the execution of the other processes, minimizing the impact on the other components of the stimulator.

### 4.3.8 Tail-movement-based VR can be used to study behaviours in head-fixed zebrafish larvae

We used our Closed-loop VR stimulator to investigate the optomotor behaviour in head-fixed larval zebrafish and their response to varying drift speed and gain. We defined gain as a multiplicative factor to the normalizer parameter,  $\mathcal{N}$ , such that it controls the effort needed to produce forward movement. Lower gain values decrease  $\mathcal{N}$  such that a greater tail movement vigour is needed to produce set acceleration and vice-versa, high gain values increase  $\mathcal{N}$  such that lesser tail movement (and thus effort) is needed to produce the same acceleration. To normalize across fish samples and recordings, we fit the normalizer  $\mathcal{N}$  at the start of each recording such that each larva achieves near-perfect drift compensation for drift speed of  $10\text{mm s}^{-1}$  and gain of 1.0 (see Section. 4.3.5).

We quantified the bout duration across all trials. With mean of 0.53s (Fig. 4.5A), the bout duration in the head-fixed larvae is longer than in freely swimming fish, this follows the results of related studies on head-restrained open-loop behaviour done by Severi et al. (2014) where the bout duration ranged from 270 ms to 370 ms. In the our closed-loop results, the fish exhibited a significant grating speed dependant increase in bout duration for all gains up to grating velocity of  $5\text{mm s}^{-1}$  (Fig. 4.5B) which contrasts the results of prior open-loop studies and closely follows results in freely swimming fish, albeit with



substantially longer bout durations (Severi et al., 2014). This difference between open and closed loop could be attributed to the tendency of larval zebrafish to attenuate their response to the OMR stimulus in the presence of no visual feedback (Portugues and Engert, 2011). While it is unclear why head-restrained fish exhibit longer bouts, it is possible that the feedback loop latency is at least one of the causes. Other explanations could include the lack of lateral line and inner ear (acceleration) input in the head-fixed preparation.

The fish also exhibited gain adaptation in the form of longer bout durations (Fig. 4.5B) as well as higher mean vigour over the duration of a given bout (Fig. 4.5C) with decreasing values of gain. The interval between bouts was not significantly affected by gain and only decreased with increasing grating speed (Fig. 4.5D).

Figure 4.5: **Bout response to gain and drift change:**

**A.** Bout duration probability density distribution for all gain-velocity combinations,  $n=57164$  bouts. **B, C, D.** Change of Bout Duration, Average Bout Vigour and Interbout Interval (respectively) on varying gain (hue, indicated by scale bar) and drift velocity (x axis),  $n=57164$  bouts. **E.** Difference between grating velocity and the average velocity achieved by the fish as a function of grating velocity and swim gain showing imperfect but significant gain adaptation (compare with F.). **F.** Difference between grating velocity and the hypothetical mean velocity that the fish would achieve if the swim vigour used at gain 1.0 was used for all lower and higher gains.

In order to validate the ability of the fish to stabilize its position in space against the simulated water flow in the closed loop virtual reality, we examined the difference between the velocity of the fish samples and the velocity of the displayed grating for different combinations of swim gains and grating velocities. We observed progressively worse flow stabilization performance with further deviation from the gain-velocity setting used for normalization between fish samples (gain of 1.0 and velocity of  $10\text{mm s}^{-1}$ ) with highest deviation from perfect compensation at high velocities and low gains where fish swam on average  $4.90\text{mm s}^{-1}$  (or 24.5%) slower than the grating speed of  $20\text{mm s}^{-1}$  and low velocities and high gains where fish swam on average  $3.44\text{mm s}^{-1}$  (or 172%) faster than the grating speed of  $2\text{mm s}^{-1}$  (Fig. 4.5E). Although imperfect, these optic flow stabilization results show significant signs of gain adaptation which is apparent in contrast with the hypothetical swim behaviour where vigour would be fixed across all gains (Fig. 4.5F). In such setting, the under- and over-compensation would range from  $-15.03\text{mm s}^{-1}$  to  $4.73\text{mm s}^{-1}$  as compared observed values of  $-4.90\text{mm s}^{-1}$  to  $3.44\text{mm s}^{-1}$ .

## 4.4 Discussion

With extensive evidence that embodiment and environmental feedback play a vital role in the formulation of animal behaviour both in zebrafish (Ahrens et al., 2012; Markov et al., 2020) and other animal models (Keller et al., 2012), the study of animal behaviour is shifting from simple feed-forward stimulation to feed-back stimulation paradigms to capture aspects of behaviour that previously could not have been explored. In conjunction with the acquisition of brain activity, such setups provide a very powerful tool for understanding of how neural activity gives rise to behaviour and how it adapts to the environment.

We have presented a Closed-loop Virtual Reality stimulator for larval zebrafish which provides a flexible, high performance, solution for for feedback-enabled stimulation of head-

fixed larvae. We facilitate this functionality through a single combined acquisition and stimulation arm, allowing for easy integration into many existing light sheet or two-photon microscope setups.

With the primary focus on opto-motor behaviour, we have established a kinetic model approximating the swim behaviour based on tail activity captured in freely swimming zebrafish larvae. As such, we included the model in our stimulator where it can be used alongside a grating stimulus to study OMR in a closed-loop setting. We demonstrate this capability by showing the gain adaptation of the swimming behaviour, identifying the bout duration and vigour as the two main variables changing with gain. Our results regarding the bout duration and inter-bout interval closely follow results previously shown in open-loop studies ([Severi et al., 2014](#)), suggesting that the small window in the projection screen used to capture the tail activity has very little, if any, impact on the fish behaviour.

The modular design of our stimulator offers innate flexibility when building and executing new experiments as well as the capability to integrate novel stimuli, swim models and acquisition hardware. The individual components of our stimulator are mostly independent, communicating only vital information as needed, allowing for multi-threaded implementation assuring high and consistent performance capable of real-time execution as demonstrated by our benchmarks. As such, the VR stimulator offers a good solution for not only studies of optomotor response but also serves as a good basis for the investigation of other behavioural paradigms by offering extensible infrastructure providing common functionality such as tail acquisition, continuous stimulus updating and data logging capability.

It should be noted that our VR stimulator solution still suffers from many inherent limitations of the hardware constraints imposed by the head-fixed preparation. The latency of the closed feedback loop still constitutes a significant portion of the bout duration, possibly affecting the anatomy of the bouts and swimming behaviour. Since our primary goal was to develop an extensible, open source, state-of-the-art stimulator which can be used in conjunction with calcium imaging, we did not significantly improve on the performance of the current stimulator solutions reported in recent literature nor have we investigated to which extent the feedback latency affects the behaviour as compared to other limitations of the head-restrained setup such as the lack of accurate information from the lateral line. As such, we identify this as one of the future directions that can be pursued in order to distinguish the major factors contributing to the behaviour observed in the freely swimming fish as well as to improve on the performance of the current head-restrained closed-loop

stimulator designs.

A significant limitation of a projector-based stimulation is the latency of the widely accessible projectors. Large amount of the consumer-grade hardware available operates at latencies around 65ms. We measured the latency of our closed loop to be 22ms with the BenQ W121OST projector. While this is a significant improvement, it could still pose a challenge for fast swim bouts as the tail movement does not overlap with the generated forward motion (Fig. 4.4A).

Another aspect that needs to be considered when using our solution based on tail movement activity in conjunction with brain imaging is the negative effect the tail movement may have on the quality of the calcium recording. The movement of the tail may result in movement or distortion of the head to a small but significant extent, requiring additional care and preprocessing when handling the resulting data. In this aspect, an electrode-based tail acquisition in paralysed larvae is superior, however significantly more challenging in preparation. Tail movement based solutions have also been successfully used in the past along with calcium acquisition (Severi et al., 2014), demonstrating their applicability despite the potential sample movement.

We have demonstrated that despite these limitations, our head-fixed closed loop preparation is suitable for the study of gain adaptation. However, our results show that the both gain adaptation and optic flow compensation of the OMR behaviour is imperfect as we observed fish to swim too slow at higher grating velocities and lower swim gains and too fast at slower grating velocities and faster swim gains. It is unclear whether this is result of aforementioned limitations or innate to fish larvae of this age and a study of freely swimming larvae would be necessary to gain further understanding of the true robustness of the OMR behaviour and function.

A number of future directions can be considered for improving and extending our open-source solution. Namely, we only provide a basic functionality for optomotor behaviour experiments with forward motion. Extension of our swim model to include turning behaviour could not only offer more opportunities in terms of behaviours studied but also improve the understanding of OMR as many head restrained studies assume forward motion without accounting for the swim effort expended on turning rather than propulsion. The causes for differences between freely swimming and head-fixed preparations are also poorly understood and further investigation may yield potential solutions for improving the virtual environment such that behaviours that were previously not expressed in the head restrained setting can be investigated with calcium imaging.

With the limited amount of closed loop virtual reality solutions for head-restricted zebrafish larvae available, our solution offers sufficient functionality for the study of optomotor behaviour and serves as a good starting point for a virtual reality for investigation of other behavioural paradigms. The single arm design with relatively low dependence on the hardware used offers high ease of implementation, including integration into existing setups used for calcium (or other) imaging, including time-scale synchronization between the recordings. All source code and documentation is available from our git repository ([https://bitbucket.org/ezfish/ez\\_spim\\_vr/src/master/](https://bitbucket.org/ezfish/ez_spim_vr/src/master/)).

## Chapter 5

# Estimating Functional Connectivity from Calcium Activity *in silico*

### 5.1 Introduction

Establishing relationships between brain activity and behaviour is one of the fundamental goals in neuroscience. Since the neuronal activity is distributed across regions throughout the whole brain, imaging methods able to simultaneously record from the whole brain are needed to gain an understanding of what role the contributing regions of activation have in the particular behaviour. In human brains, a number of methods exist to interrogate activity from spatially separated regions, such as fMRI, EEG or Electrocorticography (Neil, 1993; Lehmann, 1984; Miller et al., 2007).

Similar efforts have been applied to animal models, namely mice, larval zebrafish and others where calcium imaging is available to achieve single-cell resolution activity recordings at scale. In particular, study of larval zebrafish has yielded a large number of results (Ahrens et al., 2012; Portugues et al., 2014) on the behaviours available in the repertoire with the aid of two photon and selective plane imaging microscopy (SPIM) which allows calcium imaging at a cell resolution from nearly the full brain of the larval fish. The results found based on the SPIM data are generally supported by further interrogation of the brain structural connectivity by ablating or directly stimulating certain spatial regions of the brain hypothesised to play a role in the given behaviour (Gahtan et al., 2005; Roeser and Baier, 2003; Portugues et al., 2013). As such methods are generally not available in the human brain, causal connectivity is commonly first formulated based on directed

functional connectivity measures, such as Granger Causality (Seth et al., 2015). This provides not only an alternative to simple correlation-based study of the neural origin of behaviour but also offers a means to understand the information flow in the brain over the underlying structural connectome under different conditions.

Despite the additional information offered by directed functional connectivity measures over undirected methods (such as correlation or mutual information), the understanding of its applicability to calcium data is very limited. A study by Stetter et al. (2012) into simulated neuronal cultures provides some theoretical understanding of the applicability and limitations of directed functional connectivity measures to calcium imaging data in networks consisting of individual cells. Granger Causality (GC) has then been applied by Fallani et al. (2014) in the spinal cord of zebrafish embryos to a small subset of neurons along the tail. While both of these studies provide valuable pointers towards the applicability of Granger Causality to calcium imaging data from neural activity in live animals, it remains unclear how to apply GC efficiently to a large number of sources (neurons) across the full brain. Unlike the datasets explored by the two studies where the number of sources is generally less than a hundred, a recording of larval zebrafish brain can yield up to tens of thousands of individual neurons (Vanwallegheem et al., 2020). This limits the applicability of fully conditional Granger Causality due to both computational and combinatorial constraints.

To address this issue, we explore applicability and limitations of GC applied to calcium imaging data with particular focus on the practical limitations imposed by both the acquisition method and the size of the neural population imaged. First, a limited simulated system of two cells with feed-forward and feed-back connections where ground truth connectivity is known is examined to establish basic constraints that need to be satisfied in order to be able to establish GC effectively. Secondly, a small simulated network of 5 to 50 cells is examined for the effect of the number of components as well as acquisition artifacts such as measurement noise and signal drift on the results of the inference. Lastly, a larger simulated network is used to establish an implementation of Granger Causality which would be applicable to large *in vivo* datasets.

### 5.1.1 Granger Causality

Originating in the field of economics, Granger Causality (GC, Granger (1969)) is a statistical measure of causality (directed functional connectivity) between two processes under the assumption that 1) cause precedes its effect and 2) knowledge of the causal process

improves the predictions of the effective process when compared to taking only the history of the effective process into account.

In a multivariate setting, the second premise then extends to all other processes such that the cause must contain information about the effect that is unique and present in no other process (Granger, 2004). Estimation of conditional GC in this case then must take indirect connections into account. This has been made possible later on by extensions to the original unconditional GC which have been formulated by Geweke (Geweke, 1982, 1984) and thus modern multi-variate implementations of GC often refer to it as Granger-Geweke Causality.

Granger causality also has an information-theoretic interpretation as the transfer of information between two stochastic processes, through its equivalence with transfer entropy (Schreiber, 2000; Paluš et al., 2001). This equivalence is exact for Gaussian processes (Barnett et al., 2009), and more generally asymptotic for Markovian processes (Barnett and Bossomaier, 2012).

## VAR modelling

Commonly, Granger Causality estimation is based on the comparison of the error term variance from a vector autoregressive (VAR) model of the observed processes (Geweke, 1982, 1984). This can be formulated as

$$U_t = \sum_{k=1}^p A_k \cdot U_{t-k} + \epsilon_t \quad (5.1)$$

where  $U_t$  is vector of  $n$  stochastic processes at time  $t$  and  $A_k$  is a  $n \times n$  matrix of regression coefficients for the  $k$ -th lag and  $\epsilon_t$  is  $n$ -dimensional process of residuals.  $p$  is then a meta-parameter representing the number of lags in the VAR model (Barnett and Seth, 2014).

The lag order of the VAR model,  $p$ , is selected based on the individual characteristics of the underlying data. Generally, higher order models will be able to fit the data better, however due to the larger parameter set of  $A$  for the higher lag order, more data would be needed to estimate the parameter values with consistent statistical power compared to models with lower order. For a dataset of a constant size, the performance of the parameter estimation (and in turn GC estimation) would suffer with both too-large and too-small model order due to the bias-variance trade-off of the model order (Stokes and Purdon, 2017). It is therefore desirable to select the smallest possible lag order for which the VAR model accurately represents our data. To quantify this, a number of criteria have been formulated, with most common being Akaike information criterion (AIC, Akaike

(1973), Bayesian information criterion (BIC, Schwarz et al. (1978) and Hannan–Quinn information criterion (HQC, Hannan and Quinn (1979)). For our analysis, we used HQC to quantify the representation of our data by the VAR model as we have found it to provide a good compromise between BIC and AIC which generally tended to underestimate and overestimate the model order, respectively.

To then estimate the VAR model order based on the information criterion selected, it is necessary to estimate the unknown parameters of the VAR model. We considered two linear regression fitting algorithms: ordinary least squares (OLS) and the Levinson, Wittle, Wiggins, Robinson (LWR) (Levinson, 1946; Whittle, 1963; Wiggins and Robinson, 1965) algorithm (Morf et al., 1978). Unlike OLS which needs to be re-computed for every model order, LWR is computed recursively (Barnett and Seth, 2014), supporting its wide adoption for its stability (Pereda et al., 2005). This is favourable since more consistent results for VAR models with subsequent model orders would contribute to determining the lag number after which the extent to which the VAR model reflects the underlying data no longer improves. For this reason we also used LWR in our analysis of calcium imaging data. We limited our estimation of the model order to a maximum of 30 lags as generally the estimated model order did not exceed 20 for the majority of our models.

## 5.2 Feedback model *in silico*

To first demonstrate the applicability of Granger causality to calcium imaging data, we investigated a simple simulated feedback model of two cells connected by a reciprocal connection (Fig. 5.1A). Based on simulation tools developed by Stetter et al. (2012), we have formulated a spiking model simulation based on the NEST Simulator (Fardet et al., 2020) and a modified calcium convolution model which generates calcium traces based on the simulated spiking data.

### 5.2.1 Methods

#### Simulation of spiking data

The two neurons were modeled using the integrate-and-fire model *iaf\_psc\_alpha* implemented in the NEST simulator (parameters summarised in Table 5.1) with synapses based on the Tsodyks synapse model (Tsodyks et al., 2000) *tsodyks\_synapse* implemented in the NEST simulator (parameters summarised in Table 5.2).

Parameter	Value	Unit	Description
E_L	0.0	mV	Resting membrane potential
C_m	1.0	pF	Capacity of the membrane
tau_m	20	ms	Membrane time constant
t_ref	2	ms	Duration of refractory period
V_th	20	mV	Spike threshold
V_reset	-70	mV	Reset potential of the membrane
tau_syn_ex	2	ms	Rise time of the excitatory synaptic alpha function
tau_syn_in	2	ms	Rise time of the inhibitory synaptic alpha function
I_e	0.0	pA	Constant input current

Table 5.1: **Neuron parameters used for the integrate-and-fire simulation:** Neurons were simulated using the *iaf\_psc\_alpha* model implemented in the NEST Simulator. Details are documented at [https://nest-simulator.readthedocs.io/en/nest-2.20.1/models/neurons.html#classnest\\_1\\_1iaf\\_psc\\_alpha](https://nest-simulator.readthedocs.io/en/nest-2.20.1/models/neurons.html#classnest_1_1iaf_psc_alpha).

#### Calcium convolution

We have extended the Stetter et al. (2012) implementation of the calcium convolution model formulated by Vogelstein et al. (2009) to include the rise time of the calcium which

Parameter	Value	Unit	Description
U	0.3	a.u.	Parameter determining the increase in u with each spike [0,1]
tau_psc	3.0	ms	Time constant of synaptic current
tau_fac	0.0	ms	Time constant of synaptic current
tau_rec	500.0	ms	Time constant for depression

Table 5.2: **Parameters used for the simulation of the synaptic connections** Synapses were simulated using the *tsodyks\_synapse* model implemented in the NEST Simulator. Details are documented at [https://nest-test.readthedocs.io/en/pynestapi\\_test/models/synapses/tsodyks\\_synapse.html](https://nest-test.readthedocs.io/en/pynestapi_test/models/synapses/tsodyks_synapse.html).

was omitted in the original model. Even for fast calcium reporters, such GCaMP6f, the time to peak can be as much as 150 ms (Chen et al., 2013). It is unclear whether this has any measurable effect on the GC estimation process. However, the sampling rate used here for the GC estimation is an order of magnitude faster than the rise time, justifying its inclusion into the model.

To replicate the properties of the H2B-GCaMP6f calcium reporter, we used a mean lifetime  $\tau_{decay} = 2885$  ms equivalent to a half decay time of about 2 seconds (Kawashima et al., 2016) and rise constant of  $\tau_{rise} = 80$  ms (Chen et al., 2013).

### The two cell feedback model

We constructed a two cell model such that the first neuron,  $C_1$  has a feed-forward excitatory synaptic connection to the second neuron,  $C_2$  with synaptic strength  $W = 10$  (in pA).  $C_2$  then has a feedback excitatory synaptic connection to  $C_1$  with synaptic strength of  $w_{fb} \times W$  where  $w_{fb}$  is multiplicative factor between 0 and 1 such that the strength of the feedback connection is always defined as a fraction of the feed-forward strength, allowing for the inspection of Granger causality as a function of the relative strength of the two synaptic connections. We have set the synaptic delay to 1.5 ms which reflects values found in natural chemical synapses (Eccles, 2013). The two neurons then each receive an individual external input stimulus from a Poisson spiking generator with mean time to spike,  $\lambda_P$  in seconds (Fig. 5.1A).

### Granger Causality inference

To estimate Granger Causality, we used the Multivariate Grange Causality (MVGCM) Toolbox (Barnett and Seth, 2014). Unlike in multivariate situations, in this case with two vari-

ables, indirect associations do not need to be taken into account as no common sources outside the two cells exist. The Poisson spike generator is the only input the cells receive (apart from connections to each other) and it outputs a unique and independent spike train into each neuron.

For the purposes of this section, we estimate the Granger causality on unprocessed calcium traces captured from the two cells with the only parameters varied being the sampling rate and the sampling duration.

## 5.2.2 Results

### Functional connectivity relies on external excitation

An input a given neuron receives can have a profound impact on the activity it then exhibits. To evaluate the effect of the external input both on the activity of the two cells and the resulting directed functional connectivity, we varied the mean time between spikes produced by the Poisson generator,  $\lambda_P$ , between 1 and 1/4 such that the resulting spiking rate is either 1, 2, 3 or 4 spikes per second. As one would expect, the activity of the cell increased from very sparse ( $\lambda_P = 1$ ) to very dense ( $\lambda_P = 1/4$ ) with increasing spiking rate (Fig 5.1B). For further sections,  $\lambda_P = 1/3$  was used as it exhibited strong but still realistic activity with periods of both high activity and silence.

To evaluate the transmission of information between the two cells with the different external stimulus frequencies, we computed Granger Causality in a model with base synaptic strength  $W = 10$  pA and feedback multiplier  $w_{fb}$  between 0.0 and 1.0 (Fig. 5.1C). With increasing spiking rate of the external stimulus, the directed functional connectivity measured as Granger Causality better reflected the underlying anatomical connectivity defined by the relative strengths of the synapses. While more frequent activity in the cells resulting from more frequent external stimuli aids the accuracy of the GC inference, the improvement in the ability to discern lower feedback strengths with increasing external spiking rate likely results from the higher amount of time the neuron spends in an excited state, increasing the likelihood of co-occurrence of the input from the other cell and the external input, exceeding the spiking threshold. With sparse external input the likelihood of a neuron reaching the spiking threshold upon receiving input from the other neuron is much lower, thus reducing the directed functional connectivity between the two cells.

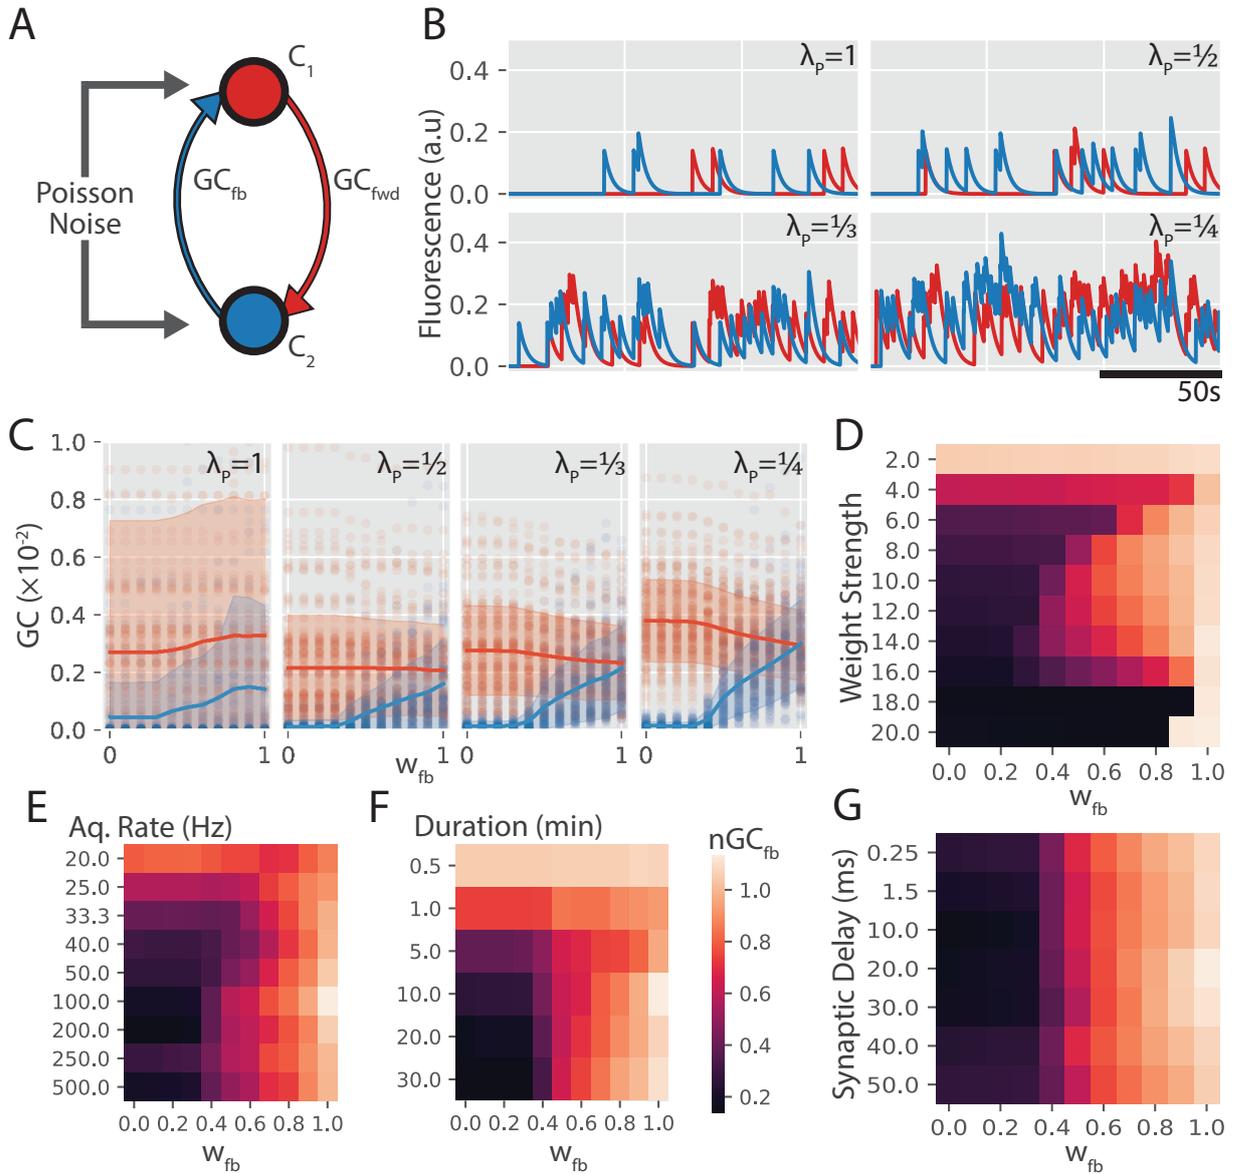


Figure 5.1: **Simulated model of two cells:**

**A.** Diagram illustrating the setup of the two cell model. Two cells ( $C_1$  and  $C_2$ ) receive Poisson input and are connected by directional connections of different strengths. **B.** Simulated calcium traces of  $C_1$  (red) and  $C_2$  (blue) for different values of the input Poisson distribution mean,  $\lambda_P$  (inverse of the spiking rate) with feedback weight  $w_{fb} = 0.1$ . **C.** Granger Causality value for feed-forward connection,  $GC_{fwd}$  (red) and feed-back connection,  $GC_{fb}$  (blue) varying with the weight of the feedback connection,  $w_{fb}$ . One plot is shown for each value of  $\lambda_P$  investigated. **D, E, F and G** Normalized Granger causality ratio  $nGC_{fb}$  (color) depending on the feed-back weight multiplier  $w_{fb}$  for varying values of base weight strength  $W$  (in D.), sampling (acquisition) rate (in E.), recording duration (in F.) and synaptic delay (in G.).

### Information transmission decreases with both low and high connectivity strengths

The synaptic connection weight  $W$  defines the overall connection strength between the two cells and thus has a direct impact on the information transferred between the cells. In order to evaluate our choice of  $W$ , we explored the effect of the feedback multiplier  $w_{fb}$  between 0.0 and 1.0 under a range of base synaptic strengths  $W$  between 2.0 and 20.0. To simplify the interpretation of our results, we present the feedback GC value as a factor of the overall connectivity:  $nGC_{fb} = GC_{fb} \times 2 / (GC_{fwd} + GC_{fb})$  (Fig. 5.1C). At very low connection strengths, we observed very low difference in the directed functional connectivity as a function of the changing feedback weight. This is the result of inability to reliably produce a response to an incoming stimulus even if the cell is already excited by the external input. With increasing  $W$ , the GC estimate better represents the underlying relative strength of the two connections. However, for high values of  $W$ , this ceases to be the case as the neurons start to fire synchronously.

To formalize this, we quantified the change in performance by measuring the consistency of increase in the GC value with increasing  $w_{fb}$  for each individual value of  $W$  as the length of the longest chain of consecutive increases in  $nGC_{fb}$ . This measure was devised under the premise that if  $GC$  was an accurate measure of the synaptic strength,  $nGC_{fb}$  in a neuron pair with higher  $w_{fb}$  would always be higher than that in a pair with lower  $w_{fb}$ , resulting in a monotonic increase of  $nGC_{fb}$  as a function of  $w_{fb}$ . Our measure may be preferable to simply summing the number of increases observed as it puts more importance on the coherence through the sequentiality requirement. We have not found any significant difference for  $W = 14.0$  and  $W = 16.0$ , however both decreasing and increasing values of  $W$  resulted in a significant progressive decrease in the ability of the directed functional connectivity to track the relative strength of the underlying synaptic connectivity (Fig. 5.2A).

### Estimation of directed functional connectivity relies on fast acquisition rates

Unlike undirected functional connectivity measures such as mutual information or correlation, directed measures rely on the ability to properly identify an accurate sequence of events in order to accurately quantify directional transfer of information. If the sampling frequency is too slow, more specifically significantly slower than the synaptic delay, it becomes not only difficult to detect significant true Granger causalities, but spurious causalities also become more frequent. Equally, over-sampling results in modelling high-frequency noise and may therefore yield adverse effects on the statistical inference of GC as

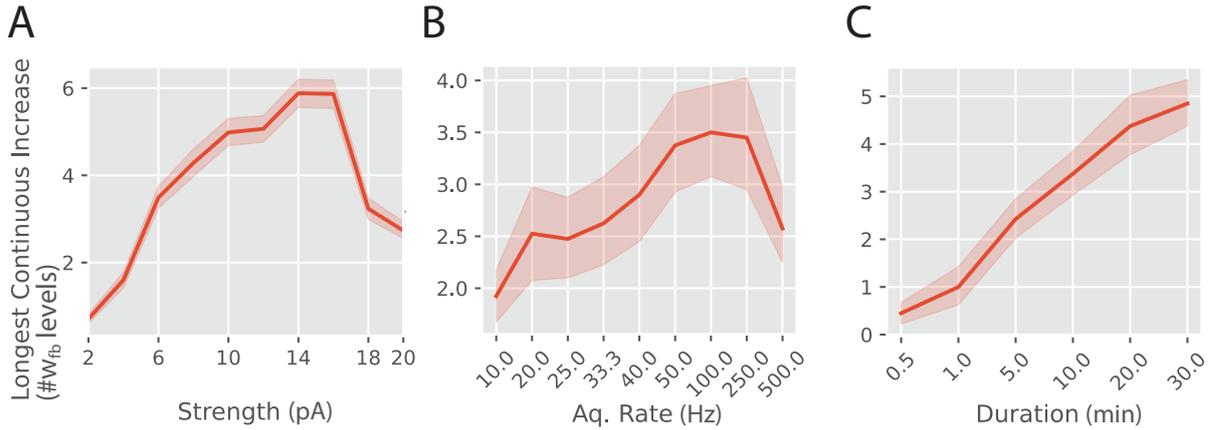


Figure 5.2: **Performance of GC inference as a function of recording parameters:** Average longest continuous increase of estimated GC value with increasing synaptic weight as a function of base synaptic weight  $W$  (in **A.**), sampling rate (in **B.**) and recording duration (in **C.**). The average longest continuous increase is denoted in weight levels examined from  $w_{fb} = 0.0$  to  $1.0$  at a interval of  $0.1$ .

well (Barnett and Seth, 2017), making the proper choice of acquisition rate very important.

Using the methodology described in the previous sub-section, we explored a range of acquisition (sampling) rates from 20 Hz to 500 Hz (Fig. 5.1E). Slower acquisition rates resulted in poor performance with gradual improvement until 50 Hz sampling rate after which there was no significant improvement until 250Hz (Fig. 5.2B). An acquisition rate of 500 Hz resulted in a significantly poorer performance when compared to 250Hz. This follows theoretical results by Barnett and Seth (2017) which indicate that sampling interval length near the transmission delay duration yields the best conditions for the GC inference. The optimal range found falls within the capabilities of state-of-the-art light-sheet microscopes which generally allow acquisition at up to 100 Hz. We therefore continue to use 100 Hz sampling rate for our sub-sequent simulations as it is both viable for current acquisition hardware and offers some margin before deterioration in performance occurring at sub-50 Hz sampling speed.

### Common recording durations are sufficient for estimating Granger Causality

While genetically encoded calcium indicators generally allow for multi-minute recordings (Cho et al., 2017), they still suffer from photo-bleaching with increasing duration and illumination power. This can become an issue as significant amount of bleaching can result in a non-stationary signal, negatively impacting the ability to estimate GC accurately. We have therefore investigated a range of recording durations from 30 seconds to 30 minutes

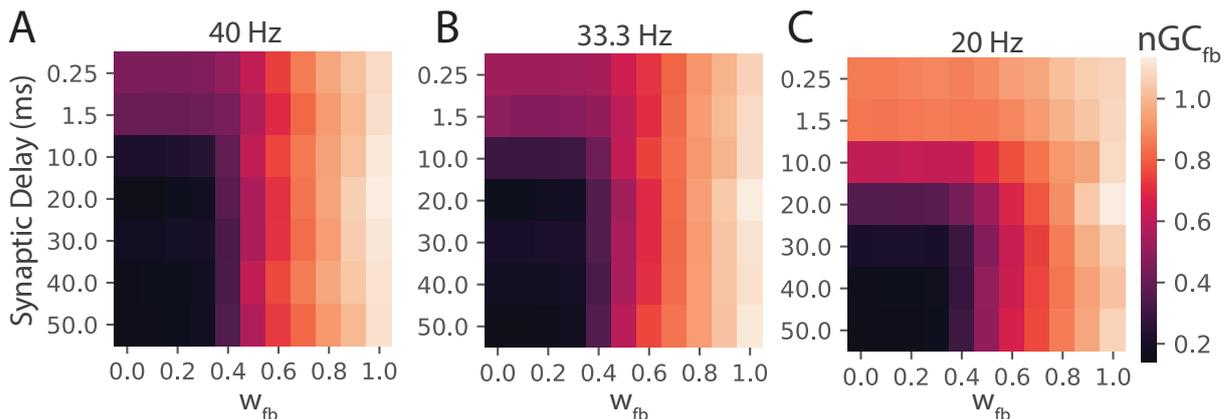


Figure 5.3: **GC is affected by synaptic delay in slower acquisition rates:** All figures follow the format of Fig. 5.1G where the synaptic delay effect is shown for an acquisition rate of 100Hz. Slower acquisition speeds show increasingly poorer performance for short synaptic delays in slower acquisition rates of 40 Hz (in **A.**), 33.3 Hz (in **B.**) and 20 Hz (in **C.**).

(Fig. 5.1F) to assess the viability of producing an adequate *in vivo* calcium recording. Longer recording durations resulted in better performance (Fig. 5.2C) with sub-5 minute recordings providing very little ability to infer GC accurately. We use 10 minute simulated recording duration for our other analyses as it is easily achievable *in vivo* and here we have shown it to be sufficient for GC inference.

### Synaptic delay only affects GC estimation at slower sampling rates

As outlined above, the feasibility of determining the correct sequence of events is vital for correctly estimating directed functional connectivity measures. With synaptic transmission speeds being generally faster (Eccles, 2013) than the shortest exposure durations of 10 ms available in modern calcium imaging systems (see Chapter 3), we identified this as one of the possible factors that could limit the applicability of GC to activity traces recorded from individual cells at this timescale.

We have adjusted the synaptic delay in our simulation, ranging from 0.25 ms (faster than the speed of an average chemical synapse) to 50.0ms (slower than a realistic rate of acquisition). At the sampling rate of 100 Hz, we have not seen any significant delay-dependant change in our results (Fig. 5.1G), suggesting that at the intended acquisition rate, the synaptic delay is unlikely to be a limiting factor. We have found for this to not be the case for slower acquisition speed (sub-40 Hz) where longer synaptic delay resulted in a significant improvement of the estimated GC value (Fig. 5.3).

### 5.2.3 Discussion

In this section, we have outlined a simple bivariate model of two integrate-and-fire neurons convolved with a realistic calcium filter. We demonstrate the importance of external input as well as synaptic strength when considering inference of Granger Causality. The chosen base synaptic weight of  $W = 10$  falls within a range for which GC performs particularly well. While we cannot guarantee that this is the case in real neuronal networks, past research has shown that neural systems often operate around criticality which is hypothesised to maximize information processing power, neither in a state with low relations between the individual neurons nor a state in which the activity of neurons is overly synchronized (Hesse and Gross, 2014), represented by low and high synaptic connectivity in our model, respectively.

Since the calcium reporter dynamics are much slower than the duration of the synaptic delay, it could be conceivably contended that the slow calcium convolution could have an adverse effect on the inference of Granger Causality. The results we present here directly contradict this hypothesis, suggesting that inference is possible as long as a sufficiently high sampling rate is used. Similar results have been shown for fMRI BOLD signals where Granger Causality analysis was shown to be invariant to the slow hemodynamic convolution (Seth et al., 2013).

We have shown that at the acquisition speed of 100Hz, the inference of GC is robust to varying synaptic delay, providing substantial evidence that establishing the time order of events sufficient for the computation of GC is possible in calcium data despite the sampling interval of 10 ms being an order of magnitude longer than the synaptic delay of 1.5 ms used in our models. We have not investigated this aspect of the model further, so the exact details of how the directed causality is computed at 100 Hz with connection delay an order of magnitude smaller, are subject to further investigation. This could be one of, and possibly combination of, following factors: 1) the membrane potential dynamics are sufficiently slow, including the depolarization after spiking and 2) since a resulting spiking activity is often the product of multiple inputs (in this case the external input and input from the other cell), it is possible that activity of one cell increases the probability of spiking in the other cell for several (tens) of milliseconds, such that external input is then more likely to bring the second cell over the spiking threshold, resulting in a delayed spiking event several (tens) of milliseconds later.

Lastly we show that for 10 minute recordings, accurate inference of GC is possible, suggesting that it is also possible to apply GC *in vivo*, provided that it can scale to a large

number of sources which is discussed in the following two sections.

## 5.3 Multi-cell networks

Having confirmed the applicability of GC to calcium fluorescence data in the two cell model, we extended our model to multiple cells. Unlike the previous bivariate setting, a multivariate setting gives rise to additional challenges, namely where indirect associations can result in incorrect inference of directed functional connectivity measures. In a simple example where neuron A has synaptic connections to both neuron B with a shorter delay and neuron C with a longer delay, it may appear as though neuron B has a synaptic connection to neuron C as Bs activity always precedes that of C even though that is not the case here as A is the common cause. We therefore use pairwise-conditional Granger Causality to control for the indirect connections. With increasing number of sources in a system, however, both the computational and combinatorial complexity increase rapidly, posing an additional constraint on the size of network analysed or the recording duration required. We therefore base all our analyses in this section on a range of network sizes up to the number of sources for which GC can be reliably computed.

### 5.3.1 Methods

#### Simulation of the spiking neuron model

We utilized our spiking data and calcium convolution model outlined in Section 5.2.1 with only minor modifications: Firstly, we only allowed connections to be either of full strength  $W$  or none at all (connection not present) with connection probability of 0.3. Secondly, we decreased the mean rate of external input from the Poisson generator to 2 spikes per second ( $\lambda_P = 1/2$ ) as each neuron now receives multiple inputs from other neurons in the model (as compared to one input in the case of two cell model). Lastly, as the number of synaptic connections in our networks grows quadratically whereas the number of neurons grows linearly, we scaled down the synaptic strength  $W$  for larger networks according to the mean number of incoming connections per neuron:

$$W = 20 \times \left( \frac{N_{synapses}}{2 \times N_{neurons}} \right)^{-0.85} + 4.0 \quad (5.2)$$

The exact parameter values have been chosen empirically such that for the explored model sizes and structures the activity remained near-critical i.e. such that neurons were able to incite activity in downstream neurons without the simulated activity in the network being overly synchronous.

### Observational noise

Under normal conditions, the recorded calcium fluorescence signal is always noisy with the signal-to-noise (SNR) ratio being the product of illumination power, calcium reporter expression and the acquisition equipment used. We use additive white Gaussian noise to reproduce this in our simulation. To achieve an output fluorescence signal with a specific SNR, we first compute the power  $P_X$  of the fluorescence trace  $X$  resulting from the calcium convolution (described in Sec. 5.2.1) as the sum of the absolute squares of the individual samples  $x_i$  in the signal divided by the signal length  $N$ :

$$P_X = \frac{1}{N} \sum_{i=1}^N |x_i|^2 \quad (5.3)$$

We then obtain the power of the Gaussian noise  $P_Z$  based on the desired SNR value,  $\eta$  as

$$P_Z = \frac{P_X}{\eta} \quad (5.4)$$

Note that we use linear  $\eta$ . If needed, this can be converted to decibels using  $\eta_{dB} = 10 \times \log_{10}(\eta)$ . Knowing the mean power of the Gaussian signal to be equal to its variance,  $P_Z = \sigma^2$ , we can then obtain the noise  $Z$  of a desired power from a Gaussian distribution  $\mathcal{N}(\mu, \sigma)$ :

$$Z \sim \mathcal{N}(0, \sqrt{P_Z}) \quad (5.5)$$

To produce the noisy calcium fluorescence trace  $Y$ , the generated Gaussian white noise is then added to the original signal:

$$Y = X + Z \quad (5.6)$$

### Signal trend modelling

Calcium traces recorded *in vivo* from larval zebrafish may contain significant trends or drifts resulting from a number of factors. Broadly, these can be categorized into two groups. Firstly, factors related to sub-optimal recording performance such as decay caused by photo-bleaching of the calcium reporter over time of the recording or changes in the region of interest (in this case the neuron body) which cannot be corrected by motion correction during post-processing. Secondly, slow time-scale changes is the activity of the recorded neurons resulting from long-term phenomena such as adaptation to or integration of a stimulus over time. While the latter category may be interesting from a neuroscientific point of view, both of these categories may be equally detrimental to the accurate estimation of Granger Causality as the stationarity of the signal is significantly

reduced with increasing presence of slow timescale trends in the signal. As such, we include slow timescale trends in our calcium fluorescence model to investigate the effects on the inferred directed functional connectivity.

In order to produce realistic signal trends, we fit a polynomial of 10-th degree ( $P(t)$ ) to raw calcium fluorescence traces  $Y_{raw}$  of neurons in *in vivo* recordings from larval zebrafish (protocols of acquisition and processing can be found in Chapter 6) using the *polyfit* function provided by the NumPy package (Harris et al., 2020). We then subtract the fitted polynomial from the raw trace to obtain 'de-trended' activity  $Y_{dtr}$ :

$$Y_{dtr}(t) = Y_{raw}(t) - P(t) \quad (5.7)$$

where  $Y(t)$  would denote an observation of calcium fluorescence value at time  $t$ . We then estimate the base fluorescence as 8th percentile of the de-trended calcium trace recording, similar to when calculating  $\Delta f/f$  of a calcium recording:

$$f_{dtr} = Perc_{8\%}(Y_{dtr}) \quad (5.8)$$

which can be used along with the base fluorescence value of our simulated calcium signal to achieve common scale between the polynomial trend and our simulated calcium trace. Assuming our simulated calcium data  $Y$  is trend-less, we can calculate base fluorescence as we did for the de-trended *in vivo* recording,

$$f = P_{8\%}(Y) \quad (5.9)$$

and then add the rescaled polynomial trend to the simulated calcium signal to obtain its 'trendy' version:

$$Y_{tr}(t) = Y(t) + \frac{f}{f_{dtr}} \times P(t) \quad (5.10)$$

A random unique trend is then added to each simulated trace  $Y$  by estimating pairs of  $P(t)$  and  $f_{dtr}$  for a large number of *in vivo*-recorded neurons. We limited neurons from the same simulated models to be paired with trends extracted from a single random recording since trends extracted from the same recording are more likely to be similar than trends collected across multiple different recordings. This ensures that the addition of trends does not produce effects larger than those that would realistically be observed in an *in vivo* recording.

## Signal processing

In order to improve signal fidelity of the recorded neural data, it is common to incorporate extensive preprocessing in order to de-noise and de-trend the recorded data as a priori

knowledge about the recording modality can be used to identify and reject artifacts from the signal (Pnevmatikakis et al., 2016). However, it has been shown that pre-processing can have a profound impact on the time ordering and information content of the signal which can lead to missing existing or identifying spurious causalities when utilizing directed functionality measures such as Granger Causality (Florin et al., 2010). This may extend to many calcium data-specific methods as well, as their application for pre-processing data prior to causal analysis has not been previously investigated and is thus poorly understood.

Signal filtering, which has extensively been used in junction with GC in EEG and MEG studies (Hesse et al., 2003; Wang et al., 2007; Kanal et al., 2009), has been explored both analytically (Barnett and Seth, 2011) and empirically (Florin et al., 2010) with similar conclusions that a notch filter can be used to remove line noise and zero-phase high-pass filtering is optimal for slow time-scale artifact removal. Barnett and Seth (2011) point out that GC is theoretically invariant to the application of arbitrary invertible multivariate filters provided the data is stationary and can thus be applied to non-stationary data in order to improve its stationarity without causing connections to be missed or spurious connections to be inferred. Florin et al. (2010) has further shown that both the type and order of the high-pass filter used has an effect on the performance of the subsequent GC analysis on simulated and empirical data. A Butterworth filter of low orders has been found to perform particularly well.

Following these studies, we employ a second order zero-phase high-pass Butterworth filter with varying cutoff frequency to investigate its efficiency at removing signal drift from the simulated calcium traces. We do not include a notch filter since calcium fluorescence acquisition does not suffer from line noise to a significant extent as other acquisition methods (such as EEG) do.

### **Evaluating directed functional connectivity results**

A major benefit of *in silico* models is the availability of the ground truth about the synaptic connections between the individual sources in the network. While it must be noted that the notion of directed functional connectivity is not synonymous with physical connectivity, the two are closely linked and we may, therefore, use the similarity between the estimated directed functional connectivity and the underlying synaptic connectivity as a measure of performance for the extent of this investigation.

To do so, we rank the inferred Granger Causality values in ascending order of their magnitude, selecting a threshold from the set of the values. Based on this threshold,

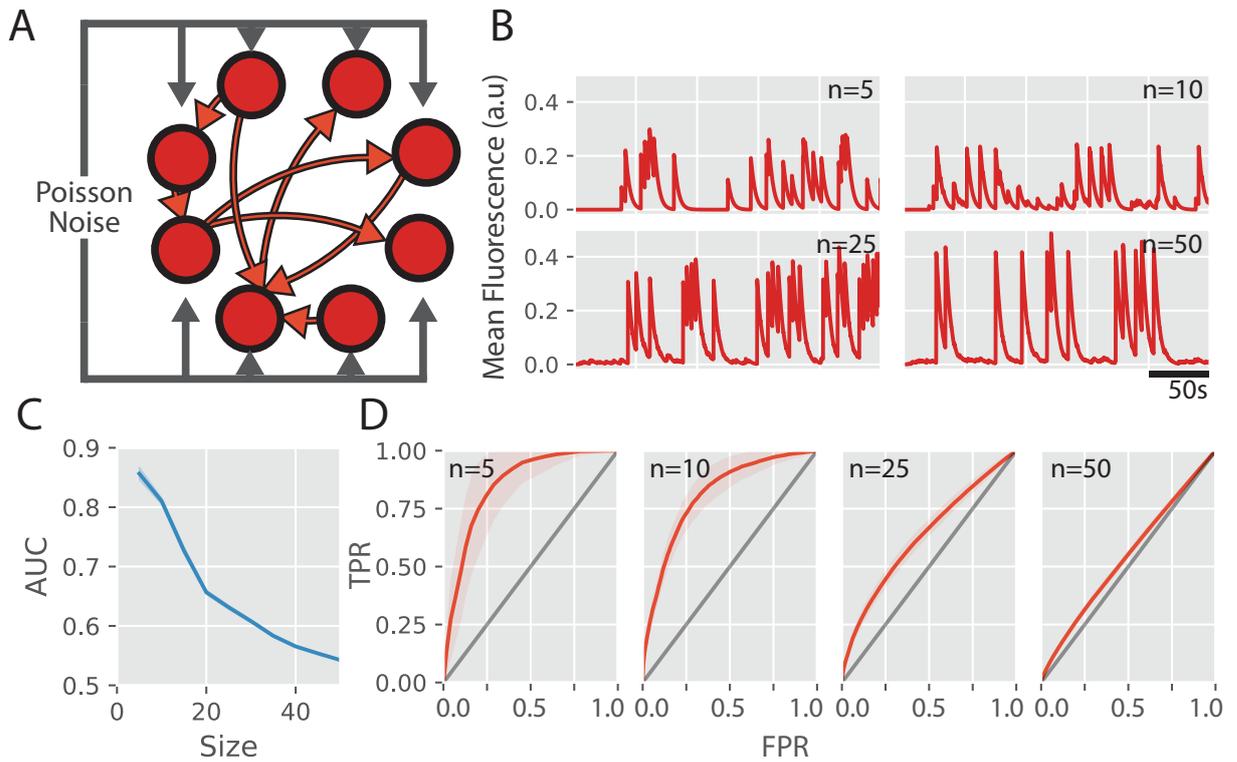


Figure 5.4: **Simulated model of multiple cells:**

**A.** Diagram illustrating the setup of the multiple cell model. 5 to 50 cells receive Poisson input and are connected by random directional connections to each other. **B.** Mean fluorescence trace of models with 5, 10, 25, and 50 cells. **C.** Average Area Under the ROC (Receiver Operating Characteristic) Curve decreases with the size of the model outlining better performance for smaller models and worse for models with large number of sources. **D** Mean Receiver Operating Characteristic (ROC) curve for synaptic connections predicted based on the estimated GC value for models with 5, 10 25 and 50 neurons (red) compared with chance level (gray).

we then label all directed edges in the network with corresponding GC value above the threshold as connections and all directed edges with GC value below the threshold as non-connections, effectively basing a binary classifier with a variable discrimination threshold on the computed Granger Causality values. The performance of said classifier can then be evaluated using a standard Receiver Operating Characteristic (ROC) curve which is created by plotting true positive rate (TPR) against false positive rate (FPR) at various threshold settings (Metz, 1978).

Area Under the ROC Curve (AUC) can then be used to summarize the overall performance across the full range of thresholds (Ling et al., 2003). Caution must be exercised, however, as the AUC measure on its own may be misleading since it occludes some aspects of the classifier’s performance (Lobo et al., 2008) with the limitation most applicable to our study being the inclusion performance in for thresholds which we know to be of rare occurrence e.g. the thresholds assuming full or close-to-full connectivity in our case, which is unlikely in real neural ensembles (Perin et al., 2011). For this reason, we include ROC results wherever possible and discuss performance within each simulation in detail.

### 5.3.2 Results

#### **The complexity of identifying directed functional connectivity increases with network size**

In order to identify combinatorial limitations given the nature of calcium fluorescence data and given our constraints of 10 minute recordings, we computed Granger Causality for simulated networks with sizes ranging from 5 neurons to 50 neurons. Although the activity in the simulated network remained largely invariant with size, we observed an increase in mean amplitude of the fluorescence signal (Fig. 5.4B) which is expected with increasing number of synaptic connections (Brewer et al., 2008).

The ability of the inferred GC values to discern underlying synaptic connectivity decreases sharply with the network size (Fig. 5.4C) which is expected given the recording duration is constant causing substantial decrease in the statistical power of the model to condition out indirect connections (Barnett et al., 2020). This effect was further exaggerated by our choice of dense connectivity (30%). It is important to note, however, that TPR improved faster over the range of lower discrimination thresholds (which assumes sparser connectivity) compared to TPR increase for higher threshold values (Fig. 5.4D, see also following sub-section).

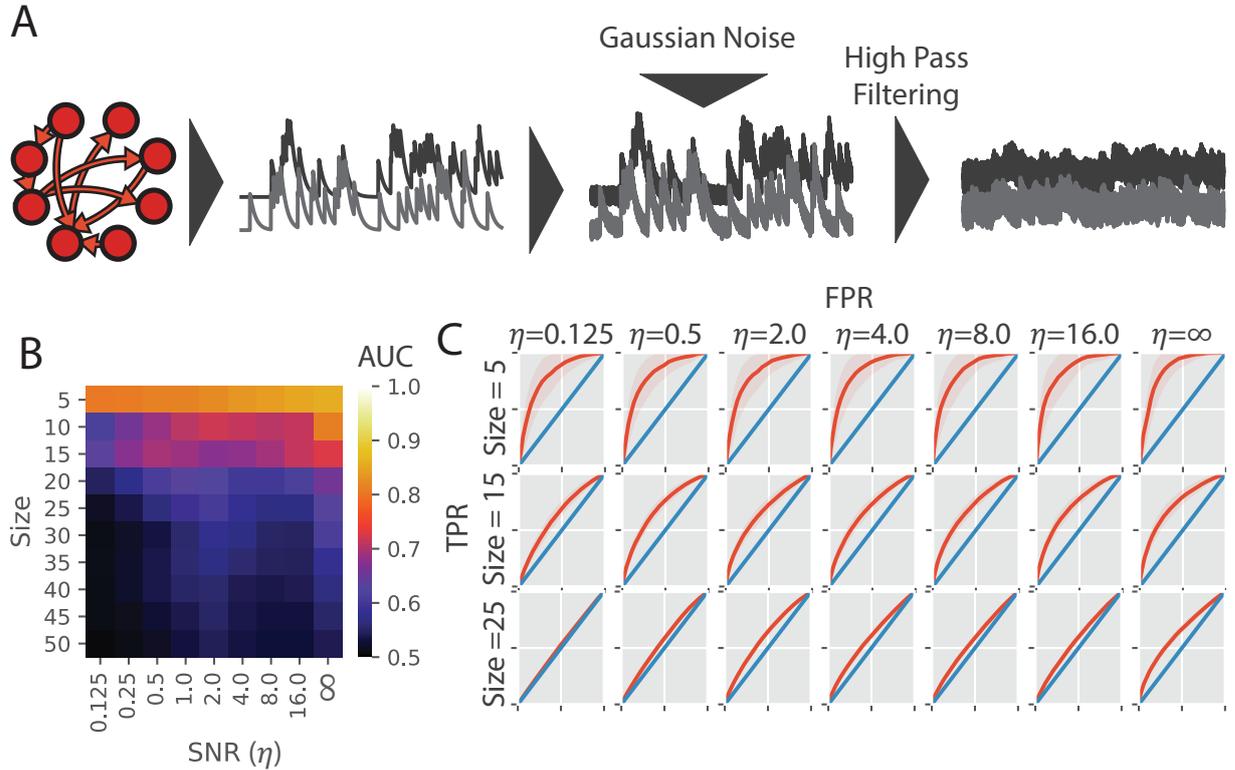


Figure 5.5: **Effect of observational noise in simulated networks of cells:**

**A.** Diagram outlining additional processing of the simulated calcium signal. A drift and Gaussian noise is added to the generated calcium trace to simulate signal drift and observational noise analogous to that which is present in *in vivo* recordings. The resulting signal is optionally processed with a high-pass filter to remove the drift from the calcium trace and improve stationarity of the signal (further shown in Fig. 5.6) **B.** Area under the ROC curve (AUC) as a function of network size and the signal to noise (SNR) ratio of the calcium traces combined with Gaussian noise of different power.  $\eta = \infty$  denotes the original calcium trace with no added noise. **C.** ROC curves for the different values of SNR ( $\eta$ ) for networks of 5, 15 and 25 neurons.

### Observational noise occludes the identification of causal connections

To investigate the effects of observational noise, we applied varying levels of noise to achieve signal-to-noise ratios ranging from 16.0 to 0.125 using method described in Section 5.3.1, without adding signal drift or filtering. Similar to results from past studies (Nalatore et al., 2007), measurement noise has been found to have a negative impact on the directed functional connectivity estimate which gets more pronounced with the relative amplitude of the noise (Fig. 5.5B).

For networks of larger sizes (15 neurons and larger), the AUC decreased with decreasing SNR non-monotonically with apparent increase in AUC around  $\eta = 2.0$  (Fig. 5.5B). We linked this to model order estimate which was increased upon the addition of noise and then steadily decreased again as the noise became more prevalent than the signal. However, this apparent increase in AUC results from improved TPR for higher discrimination thresholds (compare  $\eta = 2.0$  and  $\eta = \infty$  for size of 25 in Fig. 5.5C) offering lesser improvement when networks with sparse connections are considered.

### High-pass filtering is efficient at improving stationarity in calcium imaging data

Preprocessing the calcium signal with no trend in presence of measurement noise s.t.  $\eta = 4$  (Fig.5.6A) resulted in a significant decrease in performance when frequencies of 0.25Hz and faster were filtered in networks with 10 and more neurons (Fig.5.6B). We conclude that this is a result of some of the causal events being dampened by the filtering process and then being occluded by the high-frequency Gaussian noise which remains unaffected by the filtering. Akin to analytical analysis of filtering by Barnett and Seth (2011), we observed no significant filtering-dependent decrease in performance in noise-less simulations (Fig. 5.7).

When the simulated calcium traces are combined with the non-stationary trends extracted from calcium fluorescence recordings (Fig.5.6C), the ability to infer causal connections is significantly reduced (Fig.5.6D,E), often close to pure chance (AUC=0.5). Zero phase high-pass filtering fully restores the ability to infer functional connections using GC (Fig.5.6E) subject to causal connections lost in the measurement noise explained above even for very low filtering cutoffs.

### 5.3.3 Discussion

In this section we have investigated Granger Causality in a network of interconnected cells, validating it's performance in a multivariate setting. We have shown that the ability of

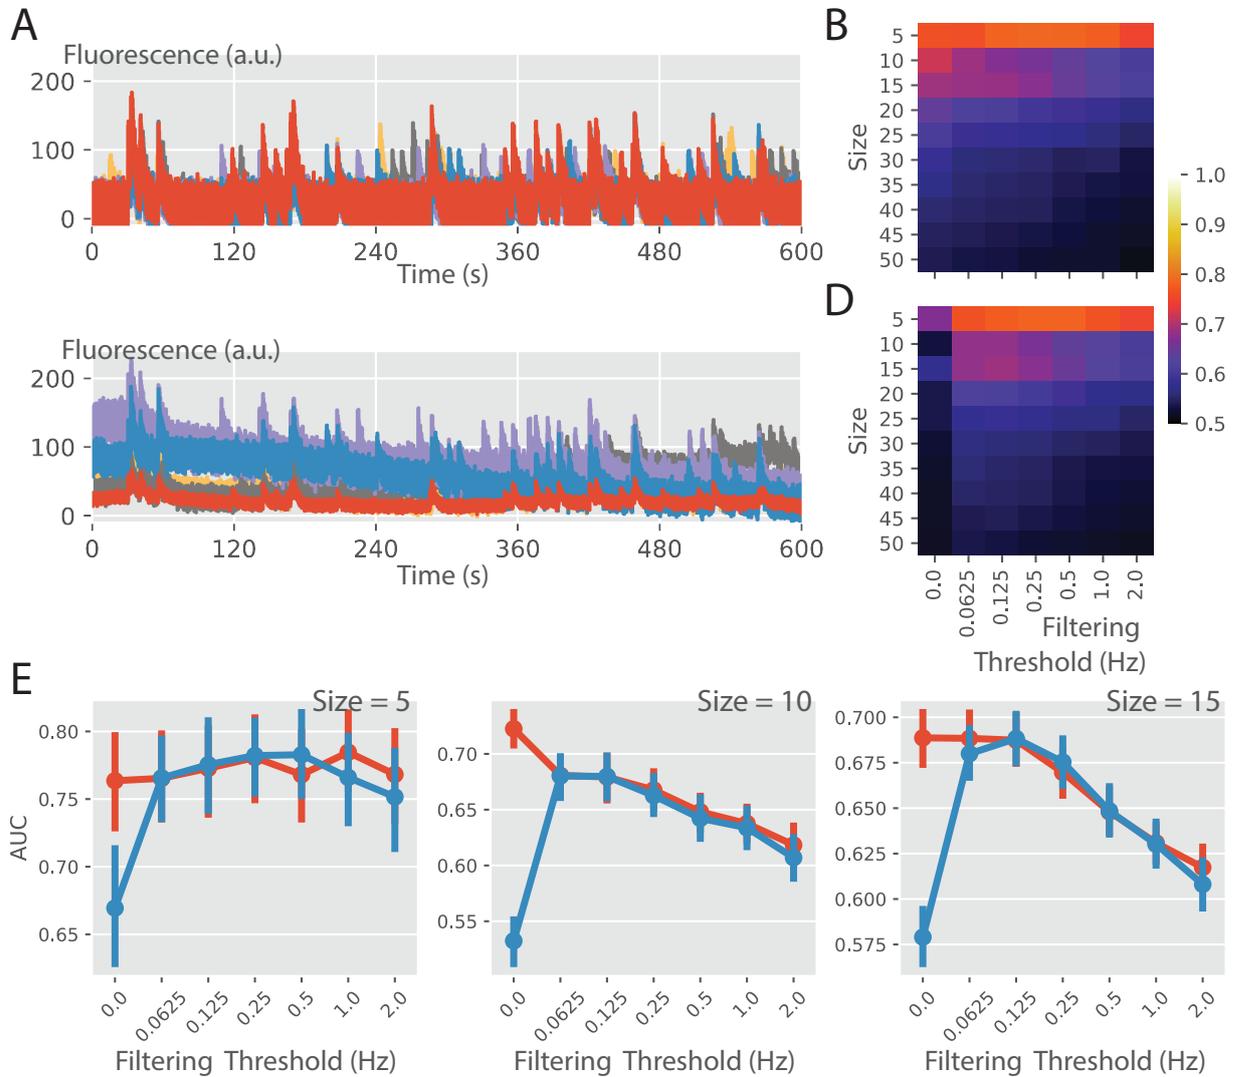


Figure 5.6: **High-pass filtering improves signal stationarity:**

**A.** A noisy calcium trace with no added signal drift. **B.** AUC as a function of network size and filtering threshold with GC estimated on calcium traces without any signal drift. X-axis is shared between sub-figure B and D. 0 signifies no filtering. **C.** Fluorescence trace with added signal drift. **D** AUC diagram analogous to B with signal drift present. **E** AUC as a function of filtering threshold for networks with 5, 10 and 15 neurons based on calcium traces with no drift (red) and added drift (blue).

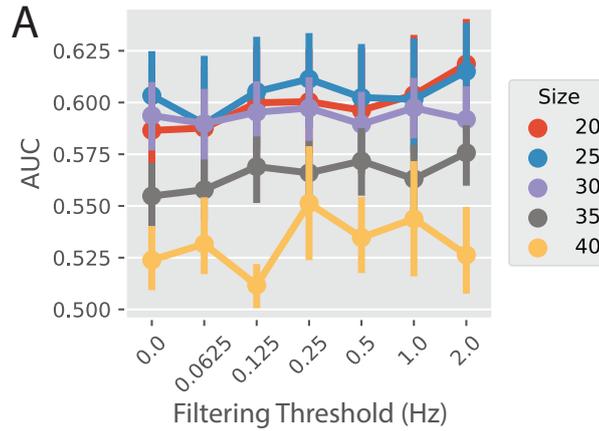


Figure 5.7: **GC is invariant to filtering in the absence of noise:**

**A.** Performance for different sizes of networks (color) does not change with filtering at different thresholds.

GC to discern synaptic connections decreases rapidly with the size of the network. This is somewhat limiting as the number of neurons in nearly all live animals is several orders of magnitude higher than that explored here. We further address this in the following section, providing possible solutions to the issue.

We further identified an adverse effect of measurement noise on the GC inference which may be a potential challenge for *in vivo* acquisition: it is generally possible to reduce recording noise by either a) increasing illumination power of the excitation laser or b) by increasing the exposure time and slowing down the rate of acquisition. The former method may, however, increase photo-bleaching of the calcium reporter, producing a non-stationary recording which is then more challenging to analyse. Similarly, the second method is limited as in the prior section we outlined the importance of high acquisition rate for correct identification of the time ordering of the causal processes. Under slower acquisitions (although with improved SNR), it could then be impossible to infer directed functional connectivity.

Lastly, while we show that high-pass filtering can have some negative effects, it provides an important means of improving the stationarity of the calcium signal, especially long timescale trends present in *in vivo* data. It is important to note that our polynomial fit used to capture trends in the acquired data only captures very slow timescales and some faster trends may be present in the *in vivo* data, which have not been included. Using a faster filtering threshold (e.g. 0.25Hz) than that found to be sufficient here (0.125Hz) may therefore produce more conservative results.

## 5.4 Scalable GC analysis methods

The previous two models dealt primarily with bivariate and multivariate settings with a low number of processes where full conditioning is possible. However, in calcium recordings, it is common to identify hundreds or even thousands of neurons (Chen et al., 2018b) which are generally organized into groups of neurons with similar activity or function (Kunst et al., 2019). Conditioning on all sources of such large datasets is combinatorially and computationally not viable, necessitating the applications of GC with partial conditioning. In this section we explore such methods applied to simulated structured networks where groups of neurons are functionally similar, providing a solution which can be used for large *in vivo* acquisitions.

For all simulations in this section we included measurement noise s.t.  $SNR = 10.0$  (see Section 5.3.1) and we pre-processed all traces with a high-pass filter with cutoff of 0.125 Hz.

### 5.4.1 Methods

#### Simulation of the spiking neuron model

We use the neuronal network model formulated in Section 5.3.1, however, here we group neurons in the network into sets of 10 'functionally similar' neurons which have high probability of synaptic connection (30%) within the group. Each pair of groups then has a 30% chance of a directed connection where neurons from one group have a 10% chance of forming a directed synaptic connection with neurons in the other group. In the absence of group-to-group connection, neurons only have 1% chance to form a synaptic connection (Fig. 5.8A). This model provides a setting in which sub-sampling and partial conditioning can be easily examined.

#### Group-wise conditional GC

Provided we are only interested in the interactions of groups considered as ensembles of neurons and interactions of individual neurons are not of interest, we can compute Granger Causality between each group pair, such that we condition only on the neurons from the groups other than the two groups considered (Fig. 5.8B). This prevents us from identifying individual neuron-to-neuron connections, particularly because GC is never computed for connections between individual neurons as it is only calculated between each two groups considered in the form of two multi-dimensional time series rather than two sets of one-

dimensional time series where individual neuron-to-neuron GC values would be computed (Geweke, 1982, 1984). As such, only one GC value (for each direction) between the two groups in the pair considered is computed.

This method comes with two major limitations. First, the total neuron count is still likely to be higher than that which is computationally viable. A second limitation lies in that the results of analyses conducted on models with the same number of groups however of varying group sizes are not comparable, requiring the selection of a smallest common group size over all the recordings we wish to compare. This may not be an issue in this simulated model as all of our group sizes are consistently of 10 neurons each, however *in vivo*, the number of neurons can not only vary significantly between formulated groups but also between recordings. As such, we might miss causal connections due to selecting a subset of neurons smaller than that offered by the recordings.

### **Sub-sampled group-wise GC with partial conditioning**

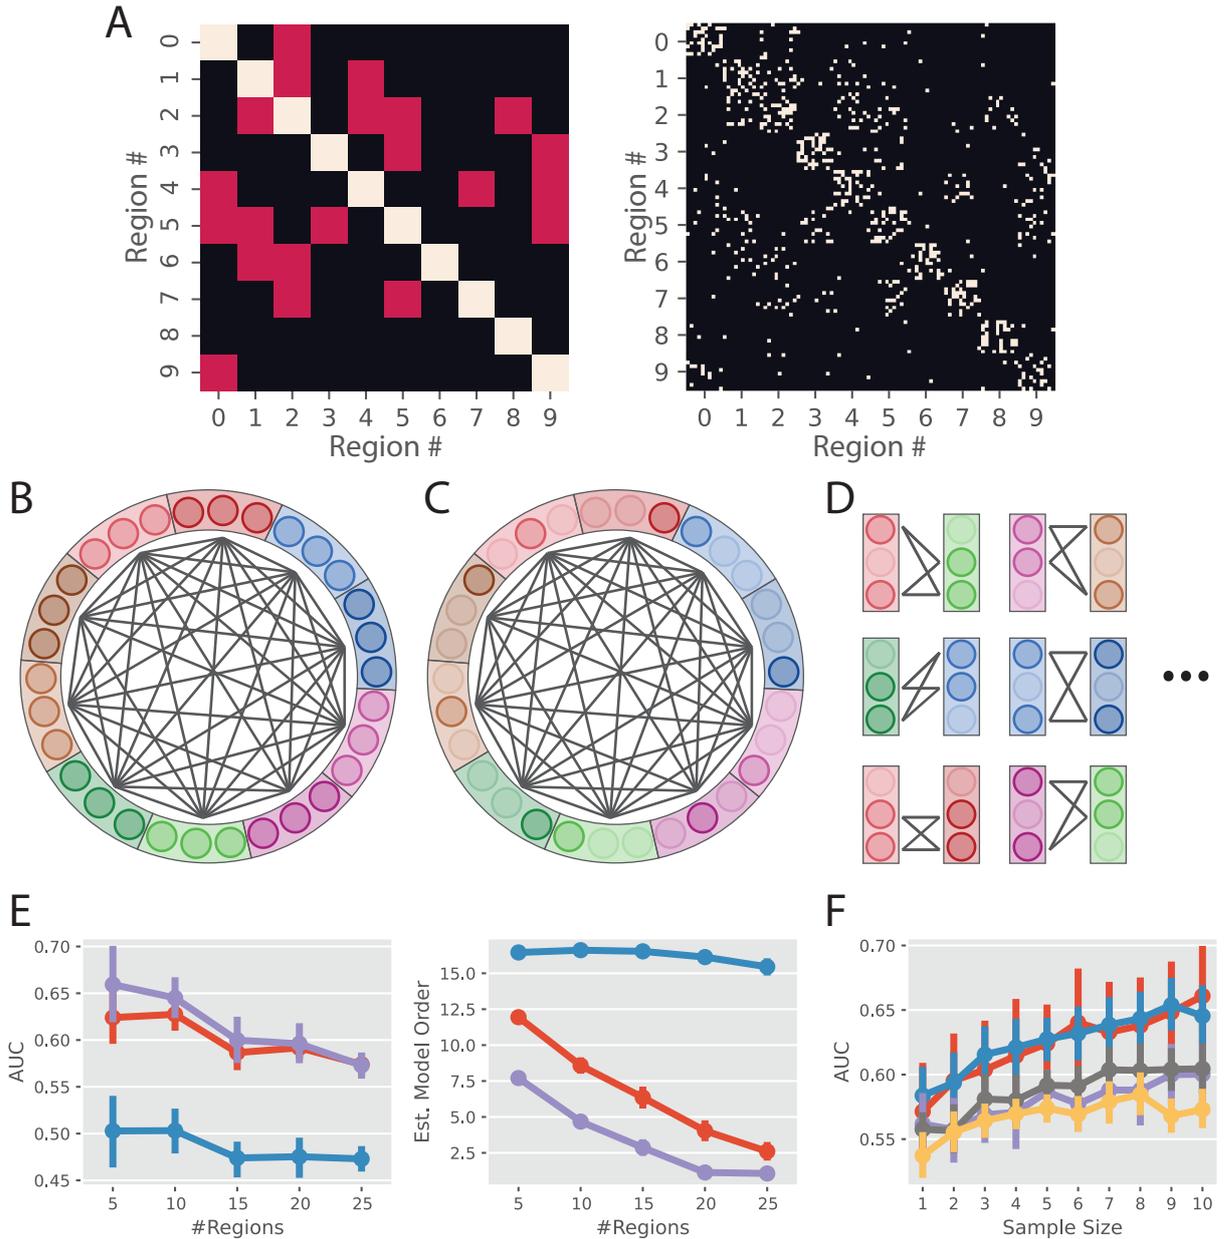
To mitigate the limitations of the group-wise method formulated above, we can select a common number of neurons we then sub-sample from each group which is equal or smaller than the size of the smallest group across all of our recordings. We can further decrease the sub-sample size to accommodate for combinatorial and computational constraints (Fig. 5.8C). We then apply the group-wise method to this sub-sample and repeat this process of sub-sampling and GC estimation several times, averaging the estimated causal value between each pair of groups across all trials.

This method ensures that we both maximize the potential of our acquired data and our results can be compared and aggregate findings be obtained across multiple recordings.

### **Sub-sampled unconditional group-wise GC**

The above two methods both condition on neurons from groups other than the two considered. As previously stated, this is advantageous as indirect causal connections can be correctly identified, however in some cases this may not be necessary, particularly in scenarios where relations between groups are sparse and easy to identify. In such cases, computing Granger Causality between each pair of groups, not conditioning out neurons of other groups may be sufficient and be combinatorially and computationally advantageous.

For this reason, we consider an unconditional approach analogous to the sub-sampled conditional method described in the previous section: For a given pair of groups, we select an equal sub-sample of neurons from each and then compute GC between them. We repeat



**Figure 5.8: Conditioning methods in large networks:** **A.** Example topology of a simulated network with 10 groups. The region (group) connectivity matrix is shown on the left where 30% of connections between regions exist with directed synaptic connections (red) and lack thereof (black). The connectivity matrix of the individual neurons is shown on the right, based on the group connectivity. Directional synaptic connections exist with 30% probability for neurons within the same region (diagonal), 10% for directed connections between regions and 1% everywhere else. **B.** Group-wise approach to conditioning out indirect connections. Groups considered are considered as a whole, conditioning on neurons of all remaining groups. In this illustration all groups are of the same size, however this need not be the case in *in vivo* data. **C.** Sub-sampled conditional approach. Similar to method in B. but only operating on a sub-sample of neurons from each group, in this example with sample size of 1. This process is repeated for several trials with different sub-samples of neurons. **D** Unconditional approach where each pair of groups is considered independently. For each pair of groups, the selected sub-samples are used to compute GC (in this case sample size is 2), not taking any other group into account, possibly leading to misidentification of connections. This process is repeated for several trials with different sub-samples of neurons.

Figure 5.8: **E.** Performance in methods shown in B. C. and D. with the group-wise approach shown in purple, conditional approach in red and unconditional approach in blue, respectively. Left plot display the performance (AUC) for networks with different numbers of regions. Similarly, the figure on the right shows the estimated model order for different network sizes of each method. **F.** Performance of GC when predicting structural connectivity (AUC) for the sub-sampled conditional approach shown in B. for different sub-sample sizes (x-axis) and sizes of the network: 5 groups (red), 10 groups (blue), 15 groups (purple), 20 groups (gray) and 25 groups (yellow).

this process for a set number of trials for each pair of groups, then averaging each causal estimate across all trials (Fig. 5.8D).

### Performance evaluation of the proposed methods

We investigated all three conditioning methods over a networks with 5, 10, 15, 20 and 25 groups, noting their performance in terms of AUC when comparing the structural and inferred directed functional connectivity (Fig. 5.8E). For both the sub-sampled conditional and unconditional method, we selected samples of 5 neurons from each group, repeated across 10 trials.

#### 5.4.2 Results

##### Unconditional approach does not track structural connectivity

When comparing the performance of the unconditional and group-wise approach (blue and purple in Fig. 5.8E, respectively) across all trials, we identified that the unconditional approach suffers from indirect connections to and from other groups of neurons which it is unable to properly condition out. This resulted in  $AUC \approx 0.5$ , equivalent to pure chance. In contrast, the AUC when conditional using the group-wise method was close to 0.66 for smaller networks (5 groups) and 0.57 for larger networks (25 groups).

##### Model order decreases with network size for the conditional but not for the unconditional approach

The computed model order (Fig. 5.8E, right) was found to not change significantly with network size for the sub-sampled unconditional approach. As the sample size selected from each group has been kept constant and the model order is computed on a per-pair basis, the number of neurons considered in the estimation process remains constant. This would remain true even if the size of group was varied as long as the sub-sample size remains

constant (not shown), maintaining constant model order.

For both the group-wise and sub-sampled conditional approach the model order decreased with network size as the number of sources grows linearly with the number of groups. The model order was then significantly smaller for the full group-wise approach when compared to the sub-sampled alternative due to the larger amount of neurons considered in the full approach.

### **Sub-sampled conditional approach can provide performance similar to group-wise conditioning method**

We found the sub-sampled conditional method to perform similarly but generally worse when compared to the full group-wise approach for all sizes of the simulated networks when the default case of sample size of 5 neurons was considered (Fig. 5.8E). We then considered different sub-sampling sizes to show its effect on the inference, observing similar trends across all network sizes (Fig. 5.8F): the performance generally improved with sample size (note that sample size of 10 is equivalent to full group-wise approach).

### **5.4.3 Discussion**

We presented a limited neural model with groups of densely connected neurons which allowed us to consider three different approaches of causal analysis in structured networks. We identified a conditional method with sub-sampling as particularly appealing, as it provides means for maximizing the use of data in our dataset, allows flexibility in terms computational complexity (controlled by the sample size and trial count) and performs well even with small sample sizes. The unconditional approach maintained model complexity (model order) across a variable number of regions which would make it a good candidate for application to sets of recordings where number of imaged regions varies greatly. However, this method did not perform on the selected model, highlighting the need for proper conditioning of indirect causes.

The investigation of the network parameters, such as connection probability between the neurons and groups and their impact on the applicability of the individual approaches was very limited as we kept most parameters constant, including the number of neurons in each group. A more thorough analysis could offer insight into settings in which the unconditional approach is applicable, offering better understanding of its utility for *in vivo* recording analysis.

Likewise, the full group-wise approach, in the way it was applied here, serves as a best-

case scenario rather than how it could be applied in practice. As mentioned, applicability of this method in a manner that is comparable across recordings would require selecting a smaller (consistent) number of neurons from all groups and recordings. For the model network here, suppose we can under this constraint only pick 5 neurons of the 10 neurons in each group. Instead, we could use the conditional sub-sampling method with, for example, sample size of 3, sampling from up to 10 neurons when available. This comparison is more akin to the application to real data recordings and it is possible that multiple trials of the conditional method would outperform the one trial of the 'full' group-wise method.

Finally, we did not investigate the impact of the number of trials in the sub-sampling approaches on the quality of the result. It is likely that with increasing number of trials, the performance of these methods would improve with diminishing returns. A formal investigation of this relationship could guide the selection of sufficiently large trial count while minimizing the computational requirements of these methods.

## 5.5 Conclusion

In this chapter we demonstrated the applicability of Granger Causality to calcium imaging data, identifying and addressing a number of general and specific limitations and constraints. First, we confirmed that GC can be reliably inferred from calcium data using a simple two cell model in Section 5.2, next we have addressed scalability and data processing using networks of interconnected neurons in Section 5.3 and finally we present methods for the application of GC at scale to structured neural networks in Section 5.4.

### 5.5.1 Recording requirements and considerations

In our simulations, sampling (acquisition) rate had a profound impact on the accurate estimation of the directed functional connectivity. With optimal acquisition rates of 50 to 250 frames per second, the highest possible acquisition rate of modern light-sheet microscopes setups (generally 100Hz) may be a favourable choice. This contrasts many behavioural studies (Portugues et al., 2014; Haesemeyer et al., 2018; Favre-Bulle et al., 2018) which utilize very slow acquisition rates ( $<10\text{Hz}$ ) as undirected connectivity measures such as correlation or mutual information do not suffer from the same constraints. In such conditions, Granger Causality is unlikely to be applicable, however it is possible that long-timescale interactions in live animal brains may allow for meaningful directed connectivity results as these are not present in our simple simulations.

With the acquisition rate of 100Hz, we have shown that the inference results for short recording durations ( $<5$  minutes) are very poor and recordings of at least 10 minutes are advisable. This amounts to 60000 samples per recording and longest possible recordings are desirable especially with non-ideal acquisition parameters such as high amounts of noise present in the acquired traces or in the presence of non-stationary changes throughout the acquisition.

We identified non-stationarity in the signal as a factor with significant negative impact on the GC estimation. Several sources for non-stationarity exist. Firstly, changes in recordings that cannot be corrected during pre-processing, such as slow z-drift during plane acquisition. Secondly, long time-scale brain processes such as integration of stimuli or adaptation. Generally these can be accounted for by additional signal filtering (described below). Finally, photo-bleaching which, apart from decaying baseline fluorescence, causes a decrease in amplitudes of the calcium reporter in response to spiking activity. This is unlikely to be correctable by filtering alone and caution must therefore be exercised when selecting illumination power and duration of the recording.

Finally, with lower signal-to-noise ratios, the estimation of directed functional connectivity also suffers as many of the small causal events become masked by the noise. This can be mitigated by either increased illumination power during acquisition which, in general, yields better SNR or by increasing the duration of the recording as the directed functional connectivity can be better estimated with increasing amounts of data. However, again, both increased illumination and recording duration will result in a more pronounced photo-bleaching of the calcium reporter, resulting in a direct trade-off between the SNR, recording duration and data stationarity.

### 5.5.2 Processing requirements and considerations

We identified zero-phase high pass filtering as a potential method for improving the stationarity of the calcium data by removing long term drifts and trends from the data. In the presence of noise, filtering with higher frequency thresholds had a progressively worsening impact on the inferred directed functional connectivity, making the selection of an appropriate filtering threshold important. In our simulation, even filters with threshold of 0.125 Hz were sufficient in removing the non-stationarity in the data, however, it is possible that in *in vivo* acquisitions, faster trends may be present, justifying use of a 0.250 Hz or even 0.5 Hz filtering threshold as it may be more conservative.

In multivariate models, we have demonstrated the importance of proper conditioning as indirect causal connections may skew the results of the directed functional connectivity inference. In particular, we showed that unconditional pair-wise computation of a Granger Causality in a network with multiple functionally similar groups yields results no better than chance and a conditional approach is therefore required. We show that this can be achieved by repeated sub-sampling of neurons from each groups offering a method that is not only computationally viable but also comparable across multiple recordings.

### 5.5.3 Relation to existing applications of GC *in vivo*

We identified a number of studies which attempted to apply Granger Causality to *in vivo* datasets. Here we briefly discuss the relationship between our findings and results presented in the related published literature.

[Fallani et al. \(2014\)](#) examined sections of the tail in zebrafish embryos, computing Granger causality between 11 neurons in the spinal cord captured using a spinning disk confocal microscope at an acquisition rate of 4 Hz for a 250s duration. The study reports an ipsilateral connectivity directed along the tail (and away from the brain) that is

characteristic of the tail beat motion larval zebrafish exhibit. Tail beat requires delayed contraction of the muscles along the tail which is different for the left and the right side of the tail. Our second simulation setting of several interconnected neurons is very similar to the setting presented. Since we identified the need for fast acquisition ( $> 50\text{Hz}$ ) for accurate inference of directed functional connectivity, it is possible that the comparably slow acquisition rate of  $4\text{ Hz}$  resulted in the identification of spurious connections in the neurons of the tail. It is also important to note that the neurons along the spine receive common input from the brain. As the authors do not record the brain activity, it would be impossible to condition it out, leaving the results susceptible to the identification of indirect connections.

[Vanwallegem et al. \(2017\)](#) investigated the responses of zebrafish larvae to auditory stimuli using a light-sheet microscope, acquiring 4 volumes of 25 planes each spanning depth of  $250\text{ }\mu\text{m}$  per second ( $4\text{Hz}$ ). GC was applied in this study over a number of brain regions where signals from each of the regions were averaged. The authors performed a number stationarity and co-integration test prior to GC to establish that the data is sufficiently stationary, however give no information on how or whether the sampling time of the planes in the volumes acquired were corrected. When the light-sheet is scanned through the brain, the activity of neurons captured at higher depths is captured later than at lower z-depths when scanning down and vice-versa when scanning up. This alters the true time ordering of the events captured, which, if not corrected, will result in spurious depth-dependant causal relations between brain regions at different depths.

[Oldfield et al. \(2020\)](#) investigated prey capture in naive and experienced animals and attempt to quantify the differences in directed functional connectivity by applying Granger Causality to both settings. The authors identify the acquisition rate of  $3.6\text{Hz}$  used is much slower than the timescale at which the action potentials resulting in the recorded activity occur, further arguing that bursts of action potentials may occur at slower timescales, closer to that of the calcium dynamics (and the acquisition rate used). We have identified that despite the slow dynamics of calcium reporters, directed functional connectivity can still be inferred, even in a two neuron model where no complex dynamics are present where we have also shown that accurate inference of GC requires fast acquisition rates (Sec. 5.2) this closely follows previously published results ([Seth et al., 2013](#)).

#### 5.5.4 Limitations and future directions

In our simulations, we mainly focused on the limiting factors Granger Causality rather than the underlying network used for simulations. While this allowed us to identify many recording and processing guidelines listed above, we have not investigated the effects of properties of the simulated network of integrate-and-fire neurons. Investigations into how the density of synaptic connectivity affects the quality of the inferred directed functional connectivity could provide a better understanding of how applicable GC is in different connection densities. The probability of synaptic connections we used in most of our simulations (30%) was likely higher than that present in most neural systems and thus provides quite conservative results. It is likely that in a network with sparser connections, GC would perform better given the smaller number of indirect connections for each given pair of neurons and/or groups of neurons. More insight would also be given into the applicability of the unconditional approach proposed in Section 5.4 which may perform better when sparser connections between groups of neurons are considered.

In the final step of our investigation, we considered sub-sampling methods as means to produce comparable and computationally plausible results, showing that smaller sample sizes yielded results of worse quality. The networks considered were still rather small (10 neurons per functional group), limiting the extent to which the sub-sampling methods could be investigated. Application of these methods to networks with a larger number of neurons e.g. 200 per functional group would provide a better understanding of the limitations of these methods as in such models, the sample size relative to the total number of neurons in each group could be much smaller, akin to *in vivo* recordings.

We mentioned photo-bleaching as one of the potential factors that needed to be considered, however the impact of which is not investigated as we do not include it in our simulations. Since photo-bleaching generally follows an exponential decay (Patterson and Piston, 2000), it may be possible to account for it by division of the signal trace by the fitted exponential. Further investigation is therefore required and could yield significant improvements in terms of attainable signal-to-noise ratios or recording durations in recordings suitable for causal analysis.

## Chapter 6

# Application of Granger Causality to *in vivo* Calcium Fluorescence Data

### 6.1 Introduction

In the previous chapter exploring the applicability of Granger Causality to simulated calcium fluorescence data, we established that the slow calcium dynamics do not prevent us from estimating GC accurately, as long as sufficiently high sampling rate is used and an appropriate amount of data is collected. Having shown a scalable sub-sampling method of Granger Causality inference on large datasets in simulation, it is necessary to demonstrate its applicability to *in vivo* data as well.

While this study is not the first to attempt to do this in larval zebrafish ([Fallani et al., 2014](#); [Vanwallegem et al., 2017](#); [Oldfield et al., 2020](#)), past applications were limited to very slow recording speeds of less than 5 Hz. It remains unclear whether Granger Causality can be inferred at such slow recording speeds as it would need to be mediated through some emerging long time-scale neural processes. This is because GC cannot be inferred from the relatively fast underlying synaptic connectivity at such low sampling rates ([Seth et al., 2013](#)) as we have also shown in our investigation in the previous chapter.

In this chapter, we acquire a dataset of single-plane light-sheet recordings of spontaneous brain activity in larval zebrafish. Using existing zebrafish atlases, we obtain a common reference and superimpose previously identified functionally similar brain regions onto our recordings. We then assign neurons identified in each recording to these brain regions based on their spatial position, allowing to then perform Conditional Granger Causality

with sub-sampling to obtain a directional functional connectivity map of the brain. We validate our results against previously identified brain region connectivity and show evidence suggesting that Granger Causality can be inferred from calcium fluorescence data using our methodology.

## 6.2 Methods

### 6.2.1 Acquisition of light-sheet data

All larval zebrafish samples were maintained and prepared for light-sheet acquisition according to protocols described in Section 3.9 which also describes the optical design of the light-sheet microscope. All brain activity data was collected from 6 and 7 dpf larvae using the MicroManager platform (Edelstein et al., 2010) and the control of the light-sheet microscope was facilitated by the  $\mu$ SPIM plug-in (Chapter 3).

For each larval sample, multiple plane recordings were acquired at different depths, 10-20  $\mu\text{m}$  apart as limited by the photo-bleaching of the calcium reporter. We used the maximum acquisition speed, 98 Hz, allowed by the microscope hardware, acquiring a total of 60000 images per recording over a 10.2 minute period following the results of our theoretical study in Chapter 5. We allowed for a 5 minute rest period between acquisitions and provided a 60s acclimatization period prior to each recording during which the illumination laser was scanned above the tissue of the fish in order to minimize the transient brain activity response to the changing light intensity at the start of the recording. Since our setup forms the light sheet by scanning laser in the x-y plane, different parts of the illuminated region are acquired at different points in time during the acquisition of each frame. To achieve best possible time ordering of acquired data and to ensure the data is acquired as uniformly as possible the laser was scanned 4 times for each frame (392 times per second). Over-illuminated regions at both ends of the light sheet caused by the reversal of the mirror galvanometer were removed using the laser masking functionality of  $\mu$ SPIM.

The volume for registration was acquired after all plane recordings for a given sample were finished, covering the same x-y region and depth of 170  $\mu\text{m}$  from the top-most part of the brain. Each sample was then culled and safely disposed in accordance with UK Animal Act 1986.

For the purposes of this study, brain activity was recorded strictly in the absence of stimulus to assure stationarity of the data.

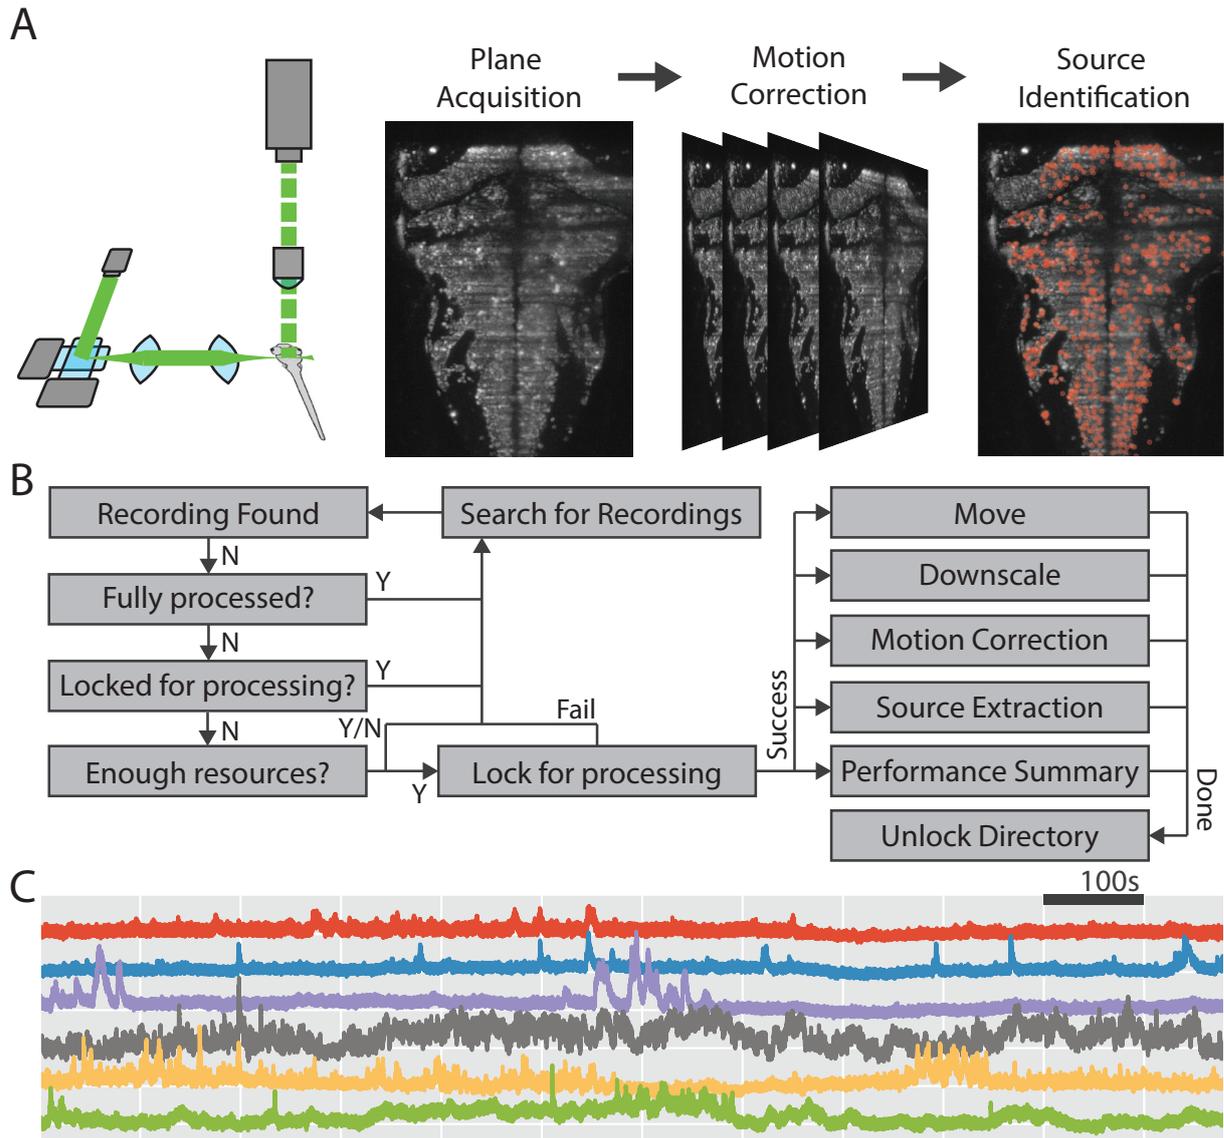


Figure 6.1: Acquisition and processing pipeline of light-sheet imaging data:

**A.** An overview of data acquisition and processing. Data acquired using light-sheet microscope controlled by  $\mu$ SPIM software has been first registered to avoid movement artifacts. Individual sources (neurons) were then extracted using the CaImAn toolbox (Giovannucci et al., 2018). **B.** Flow diagram of the scheduling system used for automated data processing at scale. A scheduler process monitors preset directories for unprocessed recordings. When a recording that is not already being processed is found and enough processing resources are available, the recording directory is locked such that no other processing task can be launched against it. The scheduler then launches the next step of the processing pipeline as a new task, allowing multiple processing tasks for different recordings to run in parallel. When the current processing step is completed, the recording directory is unlocked again, allowing for the next step of processing to be scheduled. **C.** Example fluorescence traces (a.u.) of neurons identified during processing of the fluorescence data.

### 6.2.2 Motion correction and cell identification

Prior to data processing, we binned the original recordings with 2x2 binning such that each pixel covered  $\sim 0.64 \mu\text{m}$ . We found this to have no adverse effect on the processing results and yielded a significant improvement in terms of computational speed and decreased storage requirements.

To correct x-y motion in the recordings, we utilized the elastic motion correction algorithm supplied with the CaImAn toolbox [Giovannucci et al. \(2018\)](#) with patch size of 48x48 px and 24x24 px overlap. We allowed for rigid shift of maximum 10 pixels and elastic shift of 4 pixels in each dimension. As we did not acquire data from multiple depths, z-registration of the imaging data was not possible. Instead, the data was screened prior to processing for drift along the z-axis and recordings with poor quality were removed from processing.

Identification of individual sources (neurons) was facilitated by the constrained non-negative matrix factorization (CNMF) algorithm ([Pnevmatikakis et al., 2016](#)) implementation of CaImAn toolbox ([Giovannucci et al., 2018](#)). We disabled merging of components and eliminated all sources with  $\text{SNR} < 2.5$  as computed by the CaImAn toolbox. The traces of the identified neurons were used without any de-noising or other pre-processing (Fig 6.1C).

### 6.2.3 Automated scheduling of processing tasks

Due to the large nature of datasets produced by light-sheet acquisitions, averaging  $\sim 100$  GB per 10 minute recording, the transfer and processing were slower than viable for manual scheduling of the processing tasks with transfers taking up to 20 minutes and motion correction and source identification several hours each. With several tens of recordings produced on any given day, processing of the recordings could take several days on a single machine. To accelerate and automate our data processing pipeline, we developed a simple scheduling system (Fig. 6.1B) based on the Docker container architecture ([Merkel, 2014](#)).

We chose Docker as a basis for our processing due to its wide adoption and availability on modern UNIX-based operating systems. Further, it allows us to package full execution environments for our processing tasks independent of packages and software installed on the host operating system, ensuring consistency of execution between different machines and operating systems even as software is updated over time. A centralized Docker repository then allows for consistent updating of processing tasks and their execution environments across multiple machines without the need for manual intervention on

each machine.

We split the steps in our data ingestion and processing pipeline into individual Docker containers such that each step is completely independent of any other. General steps in the processing of light-sheet single plane recordings would be:

- *Move*: Data is moved from general storage (such as local storage or acquisition storage) into shared network-attached storage.
- *Downscale*: All tiff stacks of the given recording are downscaled and incomplete meta-data produced by ImageJ is optionally removed using the *mogrify* tool supplied with the ImageMagick toolbox.
- *Plane Motion Correction*: The single plane recording is motion corrected in x-y using the CaImAn toolbox as described in the previous sub-section.
- *Plane Source Extraction*: Neurons are identified from the single plane activity recording using the CaImAn toolbox as described in the previous sub-section.
- *Performance Summary*: A brief summary of the processing steps execution is generated, including estimated motion over time and number of sources identified. A small, motion-corrected, image stack is generated by time sub-sampling of the registered recording for the purposes of manual assessment of any motion artifacts present.

To manage the scheduling of the of the processing tasks, a separate scheduler container was used, monitoring both the general storage and shared network-attached storage for unprocessed recordings. If one is found, the scheduler then uses the host operating system Docker service to launch a new container for the next step in the processing pipeline of the recording. The information about the state of processing of each recording was kept in a JSON manifest file in its respective directory which was updated at the end of each step in the processing pipeline.

In order to allow distributed processing, we implemented a locking mechanism based on the creation of a folder in the recording directory which is an atomic operation on synchronous UNIX filesystems. The locking folder would then again be removed at the end of each processing step, allowing for the next step to be scheduled. Upon failure during the execution (exhaustion of resources, loss of power or restart of the operating system), the recording directory would remain locked until it is examined and corrected by the user and the locking folder is manually removed.

Step	Parameters	Values
General	interpolation	BSpline
	use-histogram-matching	true
	winsorize-image-intensities	[0.005,0.995]
1. Rigid	metric	MI[s,d,1,32,Regular,1.0]
	convergence	[500x250x100,1e-6,10]
	shrink-factors	4x2x1
	smoothing-sigmas	2x1x0vox
2. Affine	metric	MI[s,d,1,32,Regular,1.0]
	convergence	[500x250x100,1e-6,10]
	shrink-factors	4x2x1
	smoothing-sigmas	2x1x0vox
3. SyN	metric	MI[s,d,1,32,Regular,1.0]
	convergence	[200x100x50,1e-6,10]
	shrink-factors	4x2x1
	smoothing-sigmas	2x1x0vox

Table 6.1: **ANTs registration parameters:**

Detailed description of each parameter can be found on the ANTs documentation website: <http://stnava.github.io/ANTsDoc/>.  $s$  and  $d$  in the MI metric denote the source and destination file, in this case the recording volume and the atlas volume, respectively.

Since each processing task has different requirements in terms of the computational resources, multiple tasks can be executed in parallel on any given machine. By defining a capacity of each machine based on network link speed, storage speed, processing power and memory available, we can ensure an optimal utilization of all available resources by scheduling tasks with different demands.

We processed our recordings using our scheduling system on two machines, although more machines could be used to accelerate processing:

- A Dell machine, Dual Intel(R) Xeon(R) Gold 6132 CPU (28 cores total @ 2.60GHz), 512GB DDR4 RAM, 10GBit/s Network Link
- A custom machine, Dual Intel(R) Xeon(R) CPU E5-2690 (16 cores total @ 2.90GHz), 256GB DDR3 RAM, 1GBit/s Network Link

#### 6.2.4 Atlas registration

The registration volume acquired from each fish sample was registered against the Z-Brain Atlas (Randlett et al., 2015) using a non-rigid volumetric registration provided by the ANTs toolbox (Avants et al., 2009). Parameters used in the registration call can be found in Table 6.1. Some of our recordings contained predominantly the hindbrain as we avoided illuminating the eyes of the larva directly. Such recordings often failed to register against the atlas. Registering against a sub-volume of the atlas containing only the hindbrain and part of the tectum resolved this issue.

While the Z-Brain atlas provides a number pre-defined anatomical regions, functionally similar regions defined in the MPIN atlas (Kunst et al., 2019) are better suited for the analysis of functional connectivity in this study. To superimpose MPIN region masks onto the volumes recorded, we registered the MPIN atlas volume against the Z-Brain reference using the same ANTs registration process as for our recorded volumes. With the transformations between MPIN, Z-Brain and our recorded recordings known, the MPIN masks were transformed using the *antsApplyTransforms* tool supplied with the ANTs toolbox. While the general depth of each plane recording was known in relation to the registration volume for a given fish sample, the acquisition of the plane and volume recordings would often take place up to 2 hours apart. During this time, the larva could sink, significantly altering the relative z-position between the recordings. To avoid establishing false affiliation between the neurons in the plane recording and MPIN brain regions, we re-established the relative depth between the plane and volume recordings as follows: First the background of both recording was removed by removing 60th percentile luminance value from all pixels. Next, we registered each plane of the volume recording against the single-plane recording using fast phase cross-correlation algorithm provided by the scikit-learn toolbox (Pedregosa et al., 2011) to account for any x-y drift during the period elapsed between recordings. Finally the similarity between the single-plane recording and each motion-corrected plane of the volume recording was computed as L1 loss and the loss values over all planes in the volume were smoothed using a hanning window spanning 10  $\mu\text{m}$ . This yielded a clear single best optimal plane in the case of all recordings which was manually validated.

To overlay the MPIN region masks onto each single-plane recording, the optimal plane from each mask volume was selected and translated according to the x-y shift computed above. The region affinity of each neuron was based on its centre of mass. As the MPIN atlas does not contain overlapping regions, the affinity of each neuron was unique.

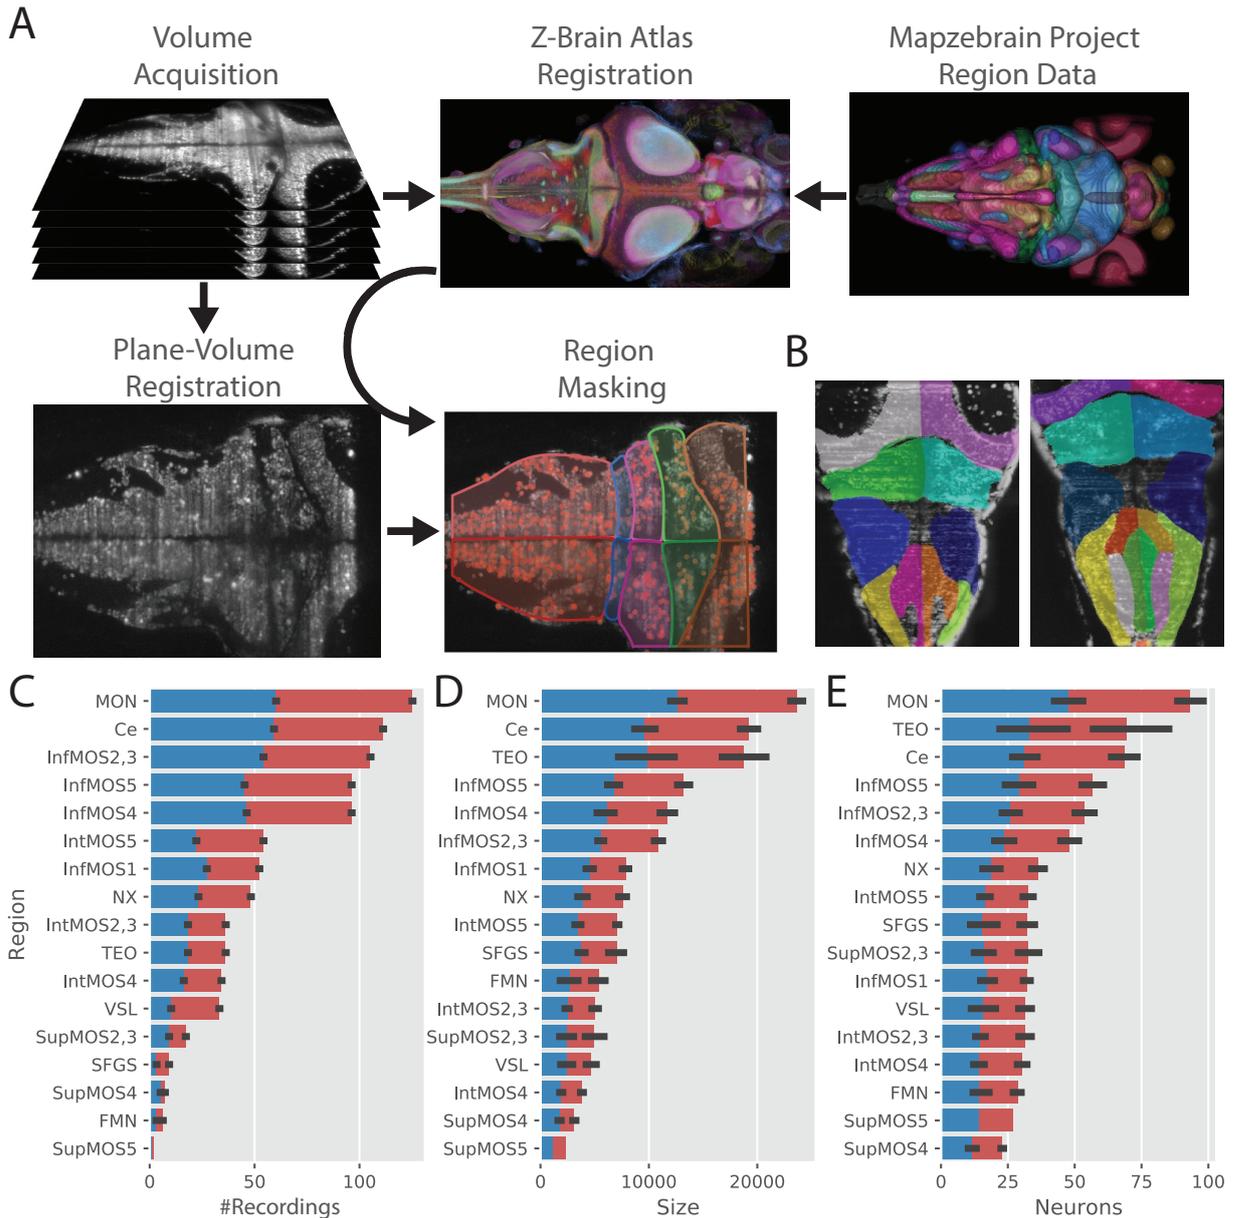


Figure 6.2: Atlas registration and neuron grouping:

**A.** After acquisition of functional plane recordings, a volume of the brain was captured. This was then registered against the Z-Brain Atlas (Randlett et al., 2015) using Elastix registration toolbox (Avants et al., 2009). Registering functional region data from MPIN atlas (Kunst et al., 2019) against the Z-Brain atlas using the same method then provides a common reference for mapping the functional regions onto volumes of individual recordings. By then registering the functional plane recording onto the volume recording, we overlay the MPIN region masks onto the plane recording which can then be used to determine region affinity of individual neurons. **B.** Region masks mapped onto two individual light-sheet recordings. **C.** Number of recordings in the dataset used for Granger Causality inference in which a given functional region was present in the left hemisphere (blue) and the right hemisphere (red). **D.** Mean size of each region in pixels in the left (blue) and right (red) hemisphere for recordings in which it was present. **E.** Mean number of neurons for each region in the left (blue) and right (red) hemisphere for recordings in which the region was present.

### 6.2.5 Brain region selection

We acquired 105 single-plane recordings from 30 fish which we manually screened and selected good recordings based on a number of criteria: We discarded recordings which were of generally poor quality as a result of excessive scattering of the light sheet through the agarose, poor camera focus or poor calcium reporter expression. We eliminated all recordings with motion artifacts, excessive z-drift and x-y drift which could not be corrected by the motion correction protocol. Samples with low brain activity resulting from poor recovery following the sample preparation or poor adaptation to the restraint conditions were also removed. 82 recordings passed the screening with satisfactory quality.

Following the registration protocol, we ensured that all recordings were accurately registered against the Z-Brain atlas. Poorly registered recordings were either retried or discarded. We obtained a sufficient atlas registration for 69 recordings, failures were mainly attributed to missing or poor quality registration volumes and uneven z-drift between acquisition of the single-plane recording and registration volume which resulted in non-trivial plane-volume registration.

We determined brain region affinity for each neuron in all recordings, treating parts of each region contained in the left and right hemisphere separately and retained all region-recording pairs containing at least 10 neurons. We then examined the number of recordings which contain a given region (Fig. 6.2C), the size of each region in pixels (Fig. 6.2D) and the number of neurons affiliated with each region (Fig. 6.2E).

To only select regions which are well represented across our recordings, we required each region to be present in both hemispheres across at least 15 recordings.

### 6.2.6 Functional measure inference

For each recording, Granger Causality was established between available regions using sub-sampled conditional approach established in Chapter 5.4 with sample size of 5 and 50 trials per recording (Fig. 6.3A,C). To eliminate the non-stationary period caused by the start of the recording, we discarded the first 20 seconds of each acquisition. Additionally, we define two measures of the asymmetry in the estimated directed functional connectivity,

$$GC_{diff} = GC_{a \rightarrow b} - GC_{b \rightarrow a} \quad (6.1)$$

Abbreviation		Region
Left	Right	
teo	TEO	Tectum
ce	Ce	Cerebellum
mon	MON	Medial Octavolateralis Nucleus
intmos2,3	IntMOS2,3	Intermediate Dorsal Medulla Oblongata stripe 2,3
intmos4	IntMOS4	Intermediate Dorsal Medulla Oblongata stripe 4
intmos5	IntMOS5	Intermediate Dorsal Medulla Oblongata stripe 5
supmos2,3	SupMOS2,3	Superior Dorsal Medulla Oblongata stripe 2,3
supmos4	SupMOS4	Superior Dorsal Medulla Oblongata stripe 4
supmos5	SupMOS5	Superior Dorsal Medulla Oblongata stripe 5
infmos1	InfMOS1	Inferior Dorsal Medulla Oblongata stripe 1
infmos2,3	InfMOS2,3	Inferior Dorsal Medulla Oblongata stripe 2,3
infmos4	InfMOS4	Inferior Dorsal Medulla Oblongata stripe 4
infmos5	InfMOS5	Inferior Dorsal Medulla Oblongata stripe 5
nx	NX	Vagus Motor Nucleus
fmn	FMN	Facial Motor Nucleus
sfgs	SFGS	Stratum Marginale
vsl	VSL	Vagal Sensory Lobe

Table 6.2: **MPIN atlas brain region abbreviations:**

Abbreviations for regions present in the data recorded. Abbreviations for the left hemisphere are in lower case and capitalized for the right hemisphere, following the convention in the MPIN cellular-resolution atlas publication (Kunst et al., 2019). Regions shaded in gray were selected as they contained at least 10 neurons in each hemisphere in 15 recordings.

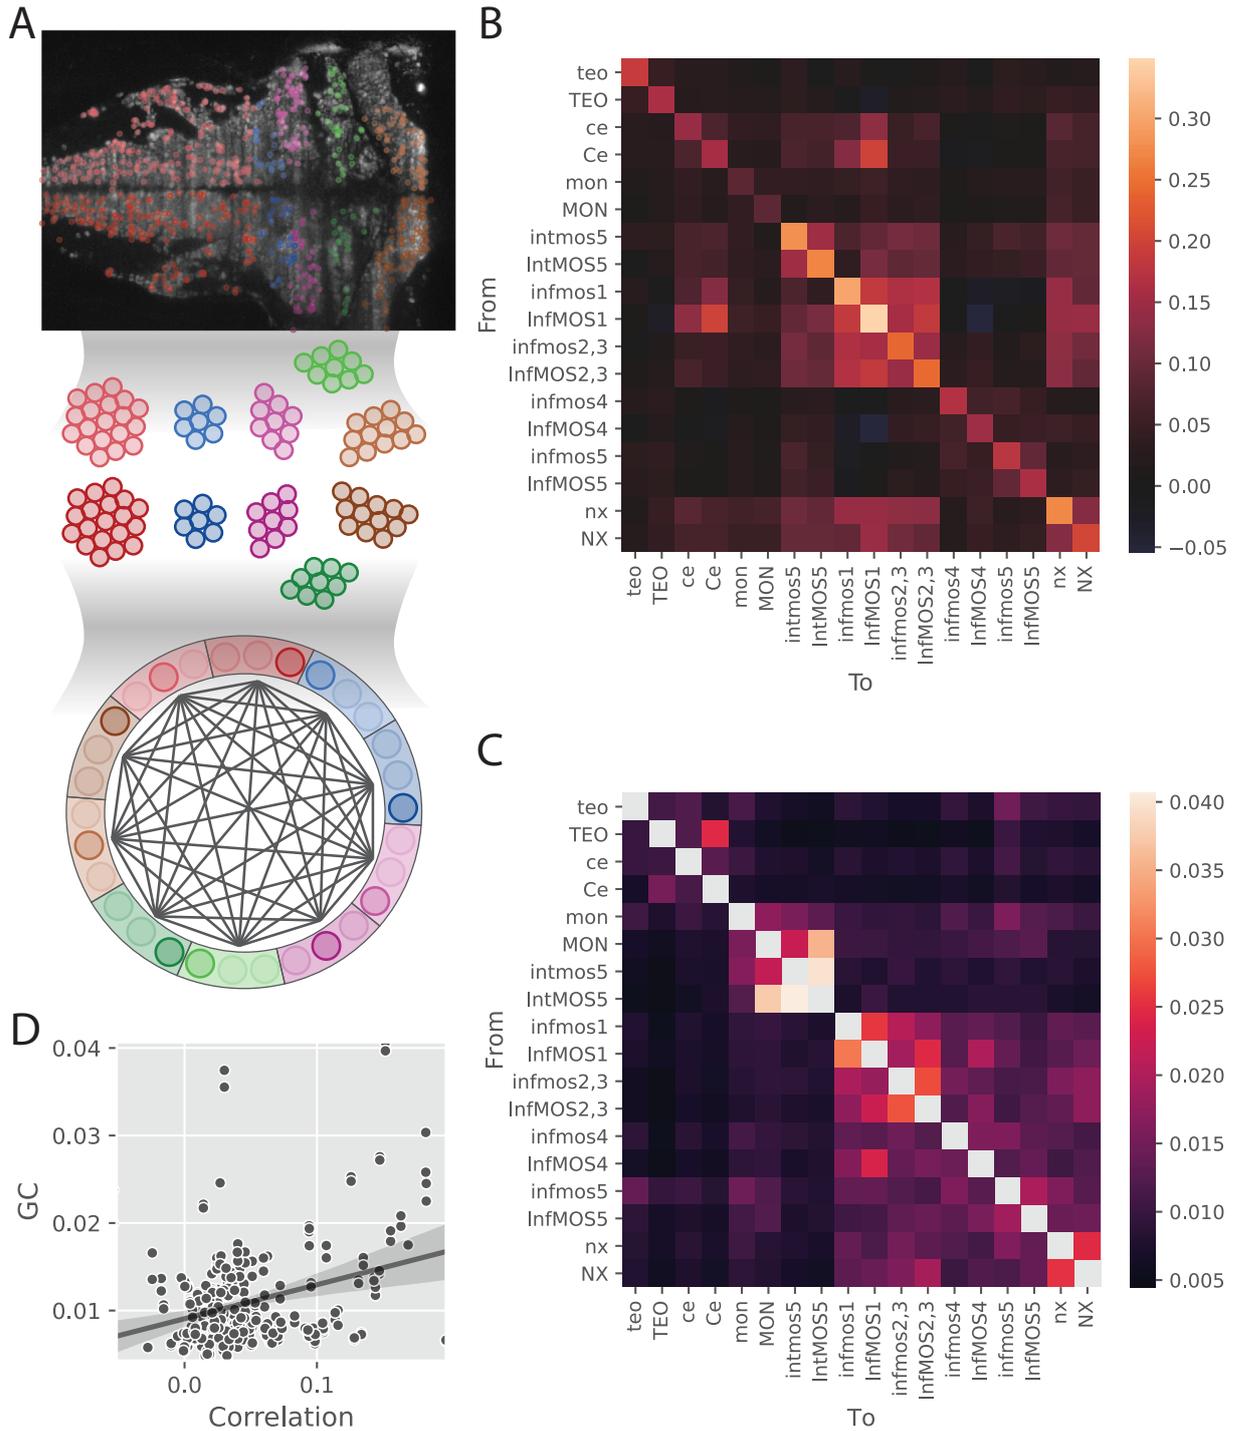


Figure 6.3: **Functional connectivity estimated *in vivo*:**

**A.** Neurons in each recording are grouped based on their region affinity. Granger Causality is then calculated using the conditional method with sub-sampling described in Chapter 5.4. **B.** Correlation matrix of the examined brain regions (names of regions corresponding to abbreviations shown can be found in Table 6.2) averaged over all recordings ( $n = 69$ ). **C.** Granger Causality for the selected regions averaged over all recordings ( $n = 69$ ). **D.** Granger Causality values plotted against Correlation. Each data point represents a single directed connection.

which denotes the difference between forward ( $GC_{a \rightarrow b}$ ) and backward  $GC_{b \rightarrow a}$  functional connectivity and its normalized version:

$$nGC_{diff} = \frac{2 \times (GC_{a \rightarrow b} - GC_{b \rightarrow a})}{GC_{a \rightarrow b} + GC_{b \rightarrow a}} \quad (6.2)$$

For comparison, we also computed correlation between the neurons of brain regions using the de-trended calcium trace of each (Fig. 6.3B). De-trending here was done using the same procedure as defined in Section 5.3.1 where high-order polynomial was fitted to the trace and then subsequently subtracted to obtain the de-trended version.

## 6.3 Results

### 6.3.1 Inferred directed functional connectivity is bilaterally similar

Prior studies have identified bilateral symmetry in the structural organization of the zebrafish brain (Kunst et al., 2019; Naumann et al., 2016; Hildebrand et al., 2017). While the organization is symmetrical, this does not guarantee bilaterally symmetrical neural responses in the presence of asymmetrical stimuli, such as sideways drift which elicits turning behaviour (Naumann et al., 2016).

To examine bilateral symmetry in the directed flow of information across hemispheres, we examined the  $nGC_{diff}$  values for ipsilateral (Fig. 6.4A) and contra-lateral (Fig. 6.4C) connections between regions in both hemispheres. If the brain was bilaterally functionally symmetrical, we would expect similar overall directionality in all connection pairs across the two hemispheres. We therefore plotted  $nGC_{diff}$  values for ipsilateral connections originating in the left hemisphere against ipsilateral connections originating in the right hemisphere (Fig. 6.4B), showing positive trend indicating bilateral symmetry ( $corr = 0.64$ ). Similar relationship ( $corr = 0.52$ ) was established between contra-lateral connection originating in the left and right hemisphere (Fig. 6.4D).

It is important to note, however, that in a true symmetry, we would expect a linear relationship with unit coefficient and zero intercept for both ipsi- and contra-lateral connections,

$$nGC_{L \rightarrow L} = 1.0 \times nGC_{R \rightarrow R} + 0.0$$

$$nGC_{L \rightarrow R} = 1.0 \times nGC_{R \rightarrow L} + 0.0$$

which was not reproduced by the linear fits (shown in Fig. 6.4B,D):

$$nGC_{L \rightarrow L} = 0.5259 \times nGC_{R \rightarrow R} + 0.0021$$

$$nGC_{L \rightarrow R} = 0.6467 \times nGC_{R \rightarrow L} + 0.0130$$

This could be explainable by the asymmetrical visual stimulation caused by the illumination laser of the light-sheet microscope which which was incident on the fish only from one (the left) side.

### 6.3.2 Granger Causality provides information different from Correlation

When plotting Granger Causality against Correlation values for the same connections (Fig. 6.3D), we can establish some linear relationship between the Correlation and Granger Causality values. It is important to note, however, that this does not mean that the information provided by both measures is the same or similar. GC provides directed information about the functional connectivity whereas Correlation does not.

The linear relationship between GC and Correlation observed here is not ubiquitous for all types of data. Its presence in calcium data likely stems from the slow dynamics of the calcium reporter: every time activity in a region elicits activity in a functionally-connected downstream region, the calcium fluorescence increases for both regions for several seconds, following the slow fluorescence decay. This then appears as though the regions are active at the same time, increasing correlation between the regions as well. Correlation cannot, however, discern the time ordering of the events.

### 6.3.3 Cerebellum shows functional efferents consistent with previously found structural connectivity

To form an overview of the asymmetric transfer of information in the larval brain (Fig. 6.5A) we identified statistically significant differences in Granger causalities  $GC_{diff}$  using Wilcoxon signed-rank test (Fig. 6.5C), corrected for multiple hypotheses, selecting connections with  $p < 0.05$ . The statistically significantly asymmetries  $nGC_{diff}$  (Fig. 6.5B, highlighted in red) were plotted on a connectivity diagram (Fig. 6.5D).

Previous cell tracing studies have found efferent connections from cerebellum to many regions of the Medulla Oblongata both ipsi-laterally (stripes 2, 3 and 4) and contra-laterally (stripes 1-5) (Kunst et al., 2019). Here we observed similar efferent connections from the left cerebellum, although not from the right cerebellum. The cause of this bilateral asymmetry is not clear, however it is plausible that the asymmetric light-sheet illumination (from the left side) is the cause.

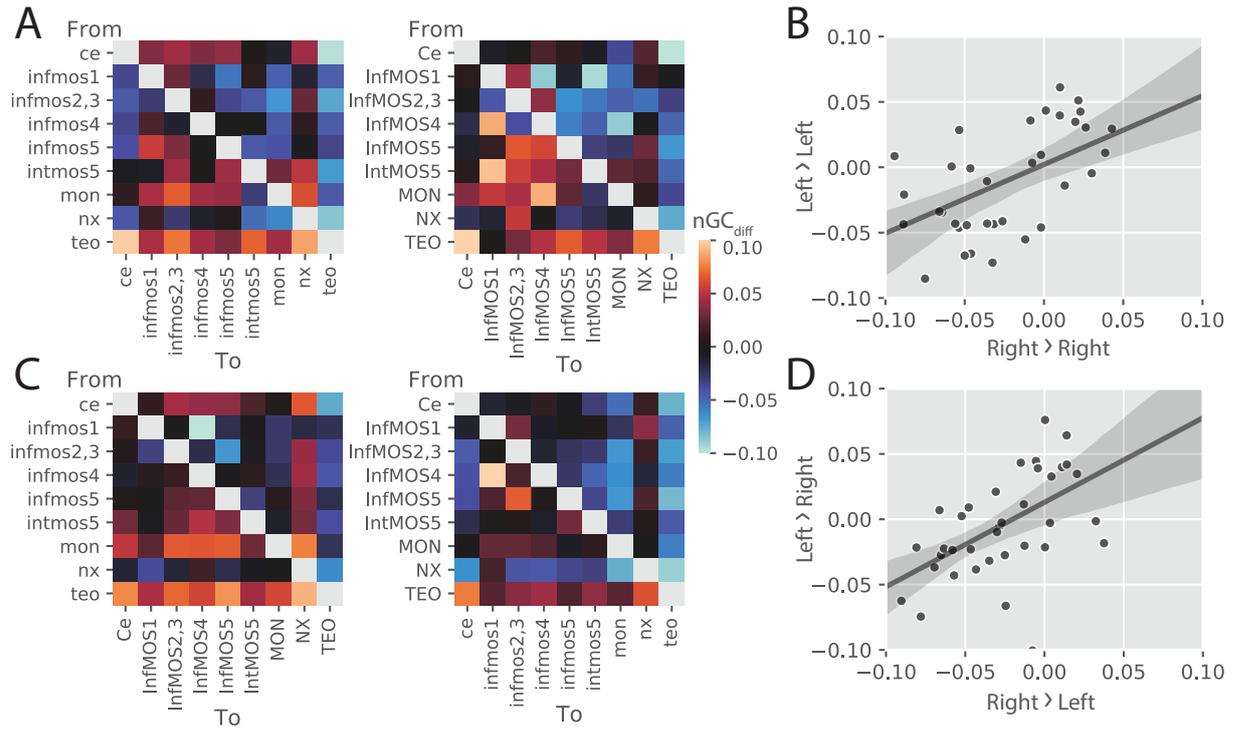


Figure 6.4: **Cross-lateral comparison of Granger Causality results:**

**A.** Normalized difference between Granger Causality values ( $nGC_{diff} = 2 \times (GC_{a \rightarrow b} - GC_{b \rightarrow a}) / (GC_{a \rightarrow b} + GC_{b \rightarrow a})$ ) for ipsilateral connections originating in the left hemisphere (left) and in the right hemisphere (right). **B.** Correlation plot for  $nGC_{diff}$  values for ipsilateral connections of the two hemispheres. **C.** Normalized difference between Granger Causality values for contra-lateral connections originating in the left hemisphere (left) and in the right hemisphere (right). **D.** Correlation plot for  $nGC_{diff}$  values for contra-lateral connections of the two hemispheres.

### 6.3.4 Medial Octavolateralis Nucleus exhibits similar functional connectivity to Cerebellum

The Medial Octavolateralis Nucleus (MON) receives a strong afferent input from the lateral line (Bell et al., 2008) and has been shown to receive afferent connections from the Cerebellum (Ce) ipsi-laterally (Kunst et al., 2019) as well as provide efferent output to cerebellar granule cells (Dohaku et al., 2019). Several studies have identified MON to have a Cerebellum-like structure, suggesting that MON may function as a part of the Cerebellum (Mikami et al., 2004; Bell et al., 2008; Bae et al., 2009). The results of our directed functional connectivity support this view as we observed ipsilateral, bilaterally symmetrical, efferent flow of information to many parts of the Inferior Dorsal Medulla Oblongata (Fig. 6.5D). Whether this is facilitated by direct structural connections which have not yet been identified or is mediated by the Cerebellum is unclear, however.

### 6.3.5 Efferent functional connections from the Tectum is greater than afferent ones

The Tectum plays a major role in virtually all visually driven behaviours as it receives numerous inputs from the retina, integrating the visual input which is often sent to downstream motor areas in the mid- and hindbrain (Gahtan et al., 2005; Severi et al., 2014; Dunn et al., 2016a; Naumann et al., 2016). While we have not observed many significant functional connections to and from the tectum (likely due to its low appearance in our recordings,  $n = 16$ ), we observed a relatively smaller incoming functional connectivity in both the left ( $GC_{in} = 0.88 \times GC_{out}$ ) and the right ( $GC_{in} = 0.87 \times GC_{out}$ ) tectum (Fig. 6.5E). These results reflect the feed-forward nature of the tectal region, although the fish was not explicitly stimulated in any of our recordings.

We also observed some significant outward functional connectivity to the Vagal Motor Nucleus and Inferior Dorsal Medulla Oblongata stripes 2 and 3 from the left tectal region. While no known direct pathways exist between these regions, it is possible that the observed functional connectivity is mediated by the nucleus of the medial longitudinal fasciculus (nMLF) which has not been included in our recordings but mediates many of visually-evoked motor behaviours (Severi et al., 2014; Naumann et al., 2016).

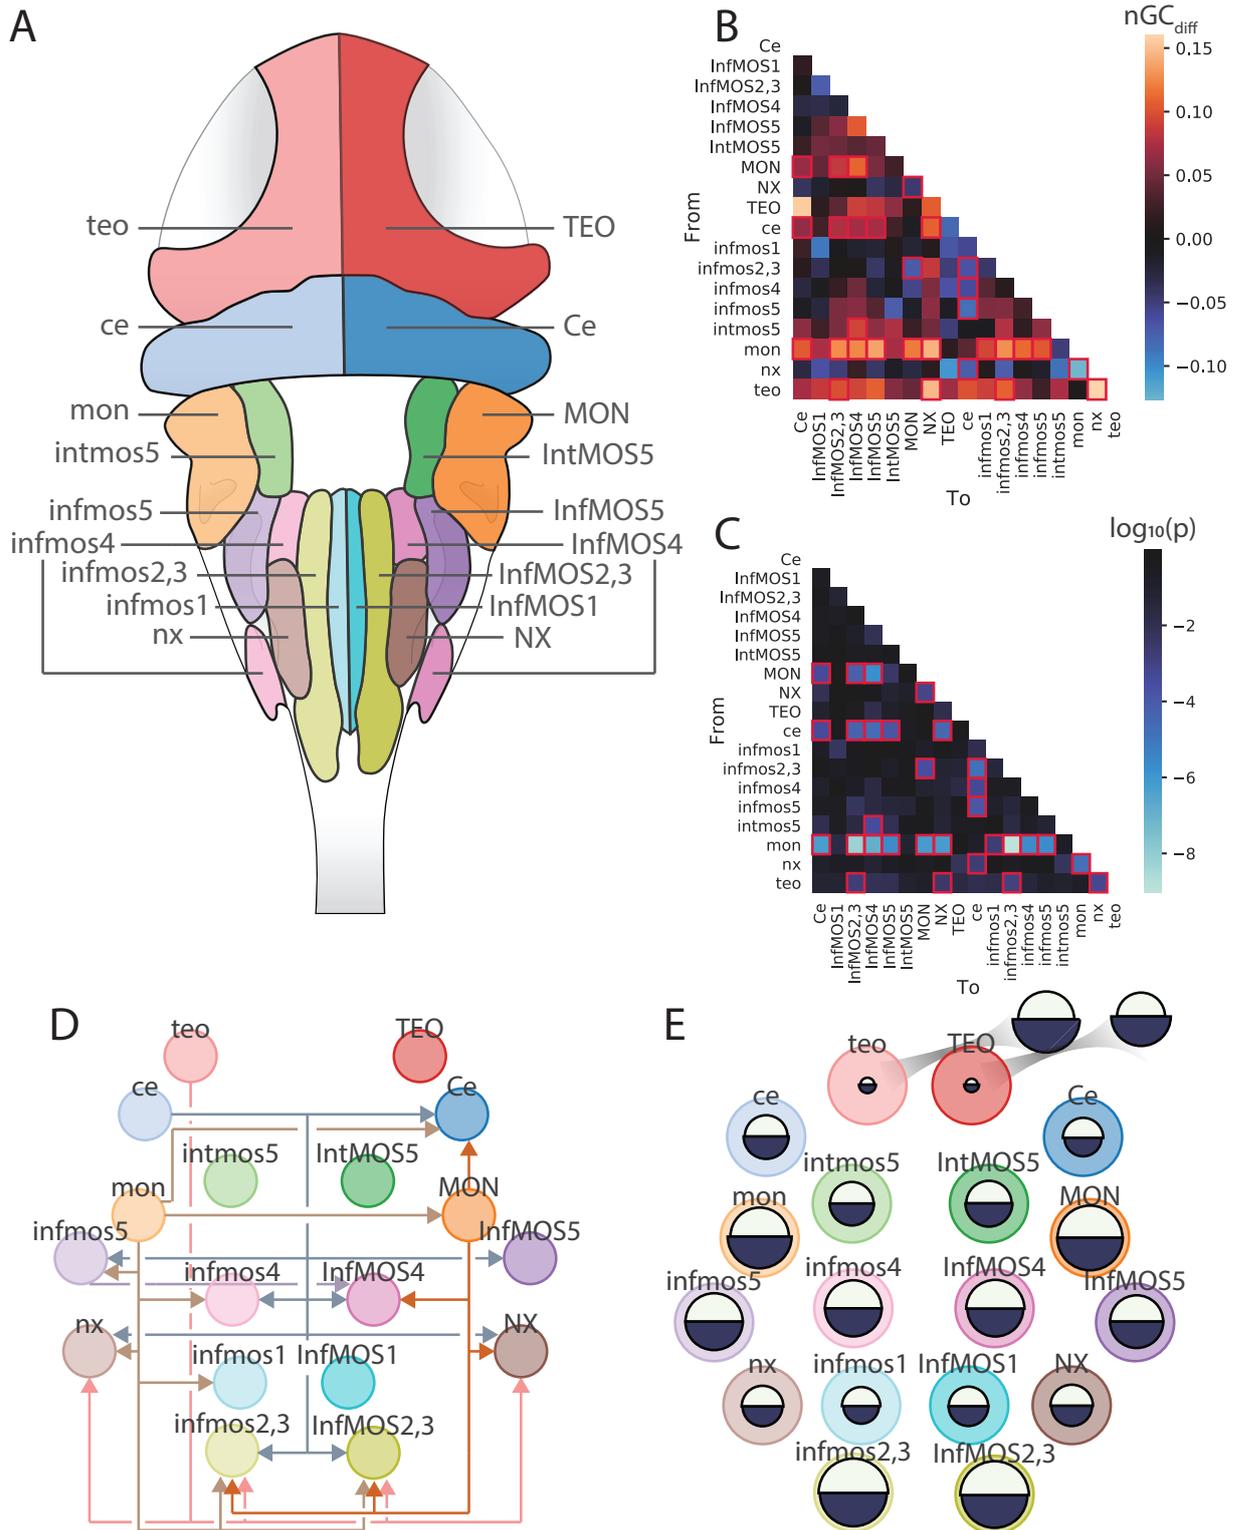


Figure 6.5: **Directed functional connectivity map of larval zebrafish brain:**

**A.** Diagram showing the brain regions analysed using Granger Causality. Explanation of brain region abbreviations can be found in Table 6.2. **B.** Normalized difference between Granger Causality values ( $nGC_{diff}$ ) between all regions considered. Connections highlighted in red are statistically significant ( $p < 0.05$ , corrected for multiple hypotheses). **C.** Statistical significance of the asymmetry in Granger Causality ( $GC_{a \rightarrow b} - GC_{b \rightarrow a}$ ) computed using the Wilcoxon signed-rank test. Values highlighted in red are statistically significant. **D.** Asymmetries shown in B. that are significant according to Wilcoxon test presented in C. were plotted on a connectivity diagram. **E.** Relative plot of aggregated incoming (white) and outgoing (dark blue) directed functional connectivity (GC) to and from each brain region.

## 6.4 Discussion

In this study, we applied Granger Causality to *in vivo* calcium imaging data collected using a light-sheet microscope setup and we show supporting evidence for the applicability of GC inference as a means to establish directed functional connectivity between brain regions.

We have shown that Granger Causality is bilaterally similar for both ipsi- and contra-lateral connections. The measure was not, however, truly symmetrical as the linear relationship between the values observed in the two hemispheres had coefficient of 0.5259 and 0.6467 for ipsi- and contra-lateral connections, respectively whereas a coefficient of 1.0 would be expected in a symmetrical setting. It is unclear as to why this is the case, however it is likely that the visible light-sheet illumination produced light pollution which affected the one eye to a greater extent than the other, effectively serving as an asymmetrical stimulus, whereby affecting the symmetry in the directed functional connectivity. A setup which has two illumination arms, each illuminating the brain from one side of the fish sample could mitigate this issue as the brain would be illuminated more uniformly, the perceived stimulus would be symmetrical and the illumination power of each arm could be halved, minimising the effect the visible light has on the sample.

By investigating the directed connections, we found connections between regions which are asymmetrical in a statistically significant manner, suggesting that GC provides more information than an undirected functional measure such as correlation. We have identified neurologically plausible connections, namely originating in the tectum and we have identified that the tectum, which is often regarded as a feed-forward processing structure, has larger outgoing than incoming G-causal magnitude in both hemispheres, consistent with this assumption. For the Medial Octavolateralis Nucleus (MON) we found similar functional connectivity to that of Cerebellum, reflecting its hypothesised functional similarity. In aggregate, these result point toward the ability of Granger Causality to uncover real directed functional connectivity in the larval brain.

We reported some connectivity that could not, however, be explained by previous research. In particular, we observed a directed functional connectivity from the left tectum to parts of the Medulla Oblongata and Vagus Motor Nucleus. This could have been perhaps a functional connection mediated by another region which has not been included in our recordings, such as nMLF and the direction of the connectivity generally follows the common flow from sensory to motor regions. The directed contra-lateral connection between Inferior Dorsal Medulla Oblongata stripe 5 (left) and Inferior Dorsal Medulla Oblongata stripe 4 (right) is not known, however. It is possible that this could be a spurious

connection or a novel functional connection that has not been previously identified.

It is also important to note that a large portion of the connections shared common directionality between ipsi- and contra-lateral connections (Fig. 6.2D) i.e. a connection from region A (left) would tend to have the same directionality to region B (left) as to the region B (right). This goes against the structure of many established behavioural pathways where brain regions often interact either ipsi- or contra-laterally (Severi et al., 2014; Naumann et al., 2016)

While our results have been mostly suggestive of correct directed functional connectivity inference, we offer no good guidance on the selection of the hyper-parameters for the GC inference. For our analysis, sample size 5 and 50 trials seemed to yield a good performance, albeit it is unclear whether these parameters are optimal. As guidance, we estimate the probability of a neuron being missed in a region of size,  $N$ , sample size,  $S$ , and number of trials,  $T$ , as

$$P = \left(1 - \frac{S}{N}\right)^T \quad (6.3)$$

With the largest region in our dataset of  $\sim 50$  neurons (Fig. 6.2E), the probability of a neuron being missed in such region would be  $(1 - 5/50)^{50} \approx 0.52\%$ . The advantage of additional trials quickly diminishes and it is possible that a lower number of trials would be sufficient in our case. To calculate the number of trials needed for a given probability  $P$ , we can solve the equation above for  $T$ :

$$T = \frac{\log(P)}{\log(1 - \frac{S}{N})} \quad (6.4)$$

and likewise, the equation can be solved for the sample size,  $S$ , as

$$S = -N \times (\sqrt[T]{P} - 1) \quad (6.5)$$

however it is important to note that  $S$  is bound by the size of the smallest number of neurons in any given brain region in the dataset and will likely be better determined by the amount data collected, with more data allowing for larger sample sizes that produce more complex models. A good protocol to establish the hyper-parameters could therefore be to first establish the sample size  $S$  given the constraints of the data and then calculate the number of trials  $T$  based on a desired value of  $P$ , for example 5%. A theoretical investigation between the interaction of the hyper-parameters  $S$  and  $T$  and the performance of the GC inference could provide more insight into the optimal parameter selection.

We limited this study to the minimal stimulus input given by the visible light of the light-sheet-forming laser. The method of estimating directed functional connectivity

would, however, prove to yield more insightful results when applied to stimulated fish samples as this would allow comparisons between directed functional connectivities under different conditions, possibly providing more insight into how the stimulus is interpreted and how response behaviours are formed. We identify this as a major future direction for both further validation of directed functional connectivity measures on calcium fluorescence data *in vivo* and as a way to guide future research of the behaviours in question.

## Chapter 7

# General Discussion

In this work, we established a robust platform for light-sheet imaging and closed loop virtual reality. We then used our light-sheet microscope control solution to establish the applicability and limitation of using Granger Causality (GC) as a measure of directed functional connectivity *in vivo*, supported by our prior investigation of GC inference on calcium fluorescence data in the simulated neuronal networks. While our results advance the understanding of directed functionality measures in calcium data, our results lead to many open questions which were not within the scope of our work.

In this chapter, we summarize our findings and provide a more general interpretation in form of guidance for application of Granger Causality to *in vivo* calcium fluorescence recordings. We conclude by outlining potential future directions of research with the aim to identify the role of sensory-motor circuits in the control of behaviour.

### 7.1 New tools to study embodied behaviours using functional light-sheet microscopy

While the applicability of light-sheet microscopy has been demonstrated for functional imaging of distributed neural activity ([Ahrens et al., 2013b](#); [Vladimirov et al., 2014, 2018](#); [Chen et al., 2018a](#)), significant technical challenges have limited its wide-spread use. We have identified a gap in the available software for light-sheet microscope control and acquisition as currently no open source solutions which provide functionality focused on functional calcium imaging currently exist. This increases the barriers of entry for both new and established laboratories looking to adopt functional light-sheet imaging for their research.

To this end, we developed  $\mu$ SPIM in Chapter 3 (and [Saska et al. \(2021\)](#)), a flexible

open-source light-sheet control extension for the MicroManager acquisition platform. The aim for  $\mu$ SPIM has been two-fold: firstly, to provide complete control of all parts of the light-sheet microscope, providing the user with functionality focused on functional imaging and secondly to do so in a hardware-agnostic manner, giving the user an opportunity to adapt the microscope setup to their particular needs.

We achieve the first aim by providing an intuitive graphical interface which seamlessly integrates with the existing MicroManager software and provides control over all functionality of our software.  $\mu$ SPIM provides both basic functionality for imaging of brain activity including single plane imaging as well as custom user-defined volume imaging and advanced functions such as laser masking to attain more uniform illumination over the light-sheet, an optional light adaptation period which reduces transient brain activity at the start of the recording caused by sudden change in illumination and sequential acquisition allowing the acquisition of multiple recordings in succession.

To give the user control over the hardware used while maintaining the quality of acquisition,  $\mu$ SPIM includes a number of calibration procedures which assure correct synchronization of all control components of the microscope. The flexibility of the acquisition hardware is provided by the MicroManager platform which integrates with many camera and shuttering systems out of the box.

To provide means to study embodied behaviour in head-restrained preparations, we develop a closed-loop virtual reality in Chapter 4. We base our solution on a physically plausible model of motion parametrized using parameters obtained from the motion of freely-swimming zebrafish larvae, implementing the VR system in a single combined stimulation-acquisition arm, offering ease of integration into existing solutions. Although we do not use the VR stimulator for our study of Granger Causality *in vivo* in this work, it could serve as a part of future studies of directed functional connectivity in the presence of stimuli.

## 7.2 Validation of Granger Causality in calcium fluorescence data

The second main aim of this work is to establish the applicability of Granger Causality to calcium fluorescence data. GC provides an appealing alternative to undirected functional connectivity measures such as correlation as it offers directed causal information which is not provided by most analysis methods applied to calcium data to-date. In the past,

GC has been successfully applied in EEG, MEG and Electroencephalography data (Hesse et al., 2003; Barnett et al., 2020; Potes et al., 2014), often using data acquired at 50-500Hz sampling rate which is attainable by modern light-sheet calcium imaging systems (up to 100 Hz). Nevertheless, attempts to apply GC to calcium fluorescence data have been scarce with no studies systematically validating its applicability to large datasets and identifying the limitations imposed by the calcium reporter dynamics.

### 7.2.1 Validation in simulated neuronal network models

We first establish three simulated neuronal network models where the ground truth structural connectivity is known.

By examining a simple bivariate case of two neurons with reciprocal connections, we were able to establish several prerequisites in order for the inference of GC to be possible. Firstly, we found the optimal range for sampling rates to fall between 50Hz and 250Hz with significant decrease in performance outside either end of this range (akin to theoretical results published by Barnett and Seth (2017)). Secondly, recording durations had large impact on the statistical power of the inference with longer recording durations offering better results. Generally, we found recordings of 10 minutes to be sufficient. Next, we identified that in order for the structural connectivity to give rise to the functional transfer of information, the target neuron may need to receive multiple inputs. This has an important implication on validating the results of directed functional connectivity based on the known structural connectivity as some structural connections may only be functionally active if the target neuron is in a specific (excited) state or when multiple incident connections are active simultaneously. Finally, we found synaptic delay to only be a limiting factor for slower acquisition rates while acquisition rate of 100 Hz remained unaffected.

Our second simulated model focused on multi-variate scenario where multiple neurons are connected by a random synaptic topology. In this model we examined the effects of measurement noise and non-stationary trends in the data, showing that gradual decrease in the ability of the inferred functional connectivity to track the underlying structural connectivity with increasing noise. Further, we have shown that adding polynomial trends extracted from our *in vivo* recordings had significant adverse impact on the performance of the GC inference, however we demonstrated that this can be corrected by a low-order zero phase high-pass filter applied prior to the Granger Causality inference. Following analytical studies (Barnett and Seth, 2011), our results confirmed that filtering with faster cut-off frequencies results in progressively worse inference performance in the presence of

noise, making the selection of slowest possible cut-off frequency which corrects the non-stationarity in the data very important.

In the last simulated model, we examined structured networks where neurons belong to functionally similar groups. We then proposed and evaluated three approaches applicable to large datasets. We have shown that without conditioning on other variables when computing GC between two groups, it becomes impossible to infer structural connectivity from the estimated directional functional connectivity. With the limitations in comparability and result aggregation across datasets (recordings) of the full approach, we concluded that averaged multiple trials of the conditional approach on sub-sampled data is most suitable for application to large datasets.

### 7.2.2 Validation in *in vivo* light-sheet imaging

Having established the basic principles of the approach of estimating Granger Causality on calcium fluorescence, we collected a dataset of 69 light-sheet single plane recordings from unstimulated 6-7 dpf larval zebrafish. Using existing tools, we extracted activity traces of individual neurons and by registering our recordings against the Z-Brain atlas (Randlett et al., 2015), we were able to establish affinity of the neurons in each recording to functional regions from the MPIN atlas (Kunst et al., 2019). Selecting only regions well represented across our recording dataset, we then computed conditional Granger Causality using repeated sub-sampling.

While we have observed bilateral similarity between the average directionality of the connections between pairs of regions, we were not able to establish true symmetry, likely due to the unilateral illumination of the fish samples by visible light. By identifying connections that were statistically larger in one direction than the other, we were able to show some similarity between our results of directed functional connectivity and known structural connectivity, namely in the efferent projections from the cerebellum. Further, we have demonstrated overall information leaving the tectum to be larger than the incoming connections from other brain regions, following the notion that the tectum mainly provides feed-forward processing of incoming sensory inputs. It is important to note, however that the amount of information leaving and entering any other brain region was very similar, that either a) the directional information provided by Granger Causality under the constraints of *in vivo* calcium imaging is very weak or b) not many of the examined regions show significant directional flow of information in the presence of no stimulus (limited to that of the visible light-sheet).

## 7.3 Applicability and limitations of GC in *in vivo* calcium recordings

This section puts our results into the broader perspective of calcium imaging, highlighting their importance in terms of ways in which GC is applied and its possible limitations. We base our suggestions on the results presented in this work and where applicable we support our arguments by past research in other recording modalities and provide examples of past uses of Granger Causality on calcium fluorescence data where applicable.

### 7.3.1 Granger Causality and Two Photon microscopy

In standard two-photon (2P) microscopy, a point is scanned in the x-y plane to form each frame of the recording. This acquisition setup is appealing due to virtually no visible light stimulation of the imaged sample, providing full control over the visual stimulus input. Granger Causality analysis of such recordings may be possible, however extra caution must be exercised when acquisition speeds significantly slower than the minimum transmission (synaptic) delay are used. Scanning of the excitatory point across the x-y plane results in significant difference in time at which the signal is acquired from different parts of the imaged area. This may result in inference of spurious directional connections, particularly in the reverse to the scanning direction of the two-photon microscope. As this issue is likely to affect most, if not all, attainable imaging speeds of standard two-photon microscope systems, use of two-photon light-sheet microscopy ([Wolf et al., 2015](#)) is advisable for studies which require greater control over the visual stimulus as the applicability of Granger Causality is better understood through the results of our study. The linear acquisition of data in the frame is not covered in our validation on simulated neural networks and further investigation would be needed prior to application to 2P data acquired at high acquisition speeds using modern resonance scanner setups ([Piyawattanametha et al., 2006](#)).

### 7.3.2 Granger Causality on volumetric light-sheet recordings

A similar issue arises in light-sheet volume recordings where a sheet of light is scanned, sequentially collecting multiple planes to form a volume. This affects the study by [Vanwalleghe et al. \(2017\)](#) where the authors present causal analysis results based on volumetric recordings. Without providing an adequate solution to this issue, it is likely that the presented results were significantly affected, where simple correlations in the brain activity may appear to be causal due to the untrue time ordering of the analysed data.

Re-sampling the individual plane traces in the volume to a common time offset (e.g. that of the center plane of the volume) could reduce this effect but it is unlikely to fully mitigate it. As this topic has not been extensively discussed in past literature, a study into potential solutions to this problem may be necessary before the application of GC to sequentially acquired data (such as volumetric light-sheet or two-photon recordings) can be justified.

### 7.3.3 Acquisition setting and parameters for single plane light-sheet recordings

Single plane light-sheet acquisition, which scans a light beam to form a plane, does not suffer from the issue of different parts of the frame being acquired at different points in time as much. This is because modern mirror galvanometers allow for much faster movement of the light beam than once per frame (even at 100 frames per second). In our study we scanned the laser back and forth 4 times per frame (or 392x per second). With our theoretical results suggesting that 500 Hz is significantly faster than the causal timescale (inferred from worse GC inference at this sampling rate compared to 250 Hz recordings, also see [Barnett and Seth \(2017\)](#)), it is likely that scanning the light beam at this speed is sufficient to avoid this issue.

The simulated neural network results presented in this study show that inference of GC from calcium fluorescence data is possible, despite the slow dynamics of the calcium reporter. This does not, however, imply that slow sampling rates are sufficient to infer causal connectivity and we show this by exploring the directed functional connectivity inferred at different sampling rates with optimal sampling frequencies being above 50 Hz and inference of causality being nearly impossible for sampling frequencies below 20 Hz. Similar results were found in fMRI data for which GC analysis was demonstrated to be invariant to the slow hemodynamic convolution but not low sampling rate ([Seth et al., 2013](#)) and the application of GC to fMRI data which is commonly collected at a very low sampling frequency has been critiqued for this reason ([Seth et al., 2015](#)). Despite this, virtually all applications of GC to *in vivo* calcium data use sub-5 Hz acquisition rate ([Fallani et al., 2014](#); [Vanwallegem et al., 2017](#); [Oldfield et al., 2020](#)) which is hard to justify in the context of the millisecond timescale of the underlying neuronal responses. While it is unclear whether some slow brain region interactions occur at such timescales, the application is likely to be limited since many behaviours are processed in less than 400 ms ([Burgess and Granato, 2007](#); [Severi et al., 2014](#)). For this reason, we suggest using at

least 100 Hz acquisition speeds following the results of our *in silico* studies.

We have shown that 10-minute recording duration is sufficient for establishing the relative strength of synaptic connections in the simulated bivariate model with little improvement for longer recordings. In multi-variate models, the optimal recording duration is likely to be much higher as proper conditioning on all variables is needed to avoid indirect connectivity issues. Generally, longer recordings will benefit the inference of GC in terms of statistical power and it is therefore advisable to select the longest possible recording duration allowed by the experimental setup.

In our experimental validation of GC, we observed bilateral similarity but not true symmetry. We hypothesized this may be because of the unilateral visible light illumination. Illumination of the sample from both sides would not only make the potential stimulus reaching the eyes of the fish symmetrical but would also allow for decrease of the power of each of the two lasers, possibly making the effect of the stimulus on the functional connectivity less pronounced. A two photon light-sheet microscopy (Wolf et al., 2015) could serve as a good alternative as it does not use visible light unlike regular light-sheet microscopes.

### 7.3.4 Signal de-trending and pre-processing

In Chapter 5 we showed that non-stationary long timescale drifts in the data have a detrimental effect on the inference of Granger Causality and we introduced zero-phase high pass filtering as a means to negate these effects. We suggest use of a cut-off of at least 0.125 Hz for the high-pass filter which is applied prior to the GC inference. Similarly, we show that filter cut-offs of 0.5 Hz and higher are not likely to be desirable as they significantly affect the subsequent inference due to the presence of noise in the signal.

### 7.3.5 Photo-bleaching of calcium reporters

Due to photo-bleaching, both the baseline and the calcium response amplitude can decrease over time, resulting in non-stationary data. While high-pass filtering can likely correct for the decaying signal baseline, it is unlikely to fully compensate for the decrease in amplitude of the calcium response, requiring additional compensatory pre-processing steps (see Future directions below) to fully correct the affected calcium activity traces. Without correction, it is up to the experimenter to assess the stationarity of the signal and adjust the recording duration or laser power down accordingly. Since we have not included photo-bleaching in our simulated models, it is impossible to provide quantitative guidance here

without further research.

### 7.3.6 Analysis of connectivity without conditioning

In Section 5.4 on structured neural network models, we explored several methods of estimating Granger Causality between individual brain regions, finding that conditioning on all other sources is necessary to accurately estimate the directed functional connectivity between two regions since our unconditional approach was not able to provide any information about the structural connectivity of the underlying networks. This has potential implications for studies utilizing the bivariate Granger Causality (Fallani et al., 2014; Oldfield et al., 2020) as the results of this method are susceptible to estimating indirect connections. Similarly, single recordings are unable to capture the full brain, making completely conditional inference impossible. It is therefore necessary to keep in mind that it is possible that some estimated functional connections will be facilitated by unobserved regions.

## 7.4 Future directions

Here we outline several lines of research which could extend the methodology established or use it to further improve understanding of functional connectivity in larval zebrafish.

### 7.4.1 The effects of the signal-to-noise ratio to data stationarity trade-off

A brief outline of the issue of photo-bleaching was given in the context of the negative impact non-stationarity has on the quality of the Granger Causality inference in Section 7.3.5. While the effects of photo-bleaching are known (Patterson and Piston, 2000), we did not include them in our simulated calcium data model. Extending the model to include photo-bleaching could provide further opportunities to understand the trade-off between the resulting degree of non-stationarity in the data, the SNR of the signal and the duration of the recording, providing information on the optimal choice for illumination power and recording length for *in vivo* acquisition.

### 7.4.2 Methods to compensate for calcium reporter photo-bleaching

Since illumination power and recording duration both increase the apparent effects of photo-bleaching, a method for negating its effect could not only improve the stationarity

of the data but also allow for longer recordings or increase in the laser power (yielding higher SNR) and thus increasing the statistical power of the interference. Since the general dynamics of the calcium reporter over time with different power of the laser have been well described (Patterson and Piston, 2000), estimating the decay constant for a given recording setting would then be sufficient to construct an exponential decay of the fluorescence. Division of the individual neuron traces by the exponential trend may then be sufficient to correct for the photo-bleaching effects without affecting the time precedence of the signal.

### 7.4.3 Comparison of GC with other directed functional connectivity metrics

While we have shown the applicability of GC to calcium fluorescence data in this study, we have not established that Granger Causality yields the best results in terms of the inferred directed functional connectivity. Other measures such as Transfer Entropy (Schreiber, 2000) or Cross-correlation (Silfverhuth et al., 2012) or their extensions (Orlandi et al., 2014) could perform better and it is therefore necessary to investigate them comparatively against Granger Causality.

### 7.4.4 Directed functional connectivity in behaving animals

In our *in vivo* study we have focused on Granger Causality in samples with no explicit external stimulus. While this ensures stationarity in the calcium signal for the validation of our GC inference method, recordings where a stimulus is presented are necessary to examine specific behaviours. Granger Causality can then be used to compare sensory-motor circuit interactions in terms of the directed functional connectivity across different behaviours and settings. For example, comparing GC on recordings which have been acquired in the presence of either forward or backward stimulus grating and comparing them to each other or to the reference connectivity with no explicit stimulus could help identify the functional changes that happen in order to suppress swimming or to swim more frequently in fish in the two different conditions.

## 7.5 Concluding remarks

With the increasing research interest in the brain-wide dynamics underlying animal behaviour, we believe that our open-source tools will aid wide adoption of light-sheet microscopy for whole brain imaging, advancing our understanding of more complex behaviours. As the complexity of the examined neural interactions grows, the importance of efficient

analysis methods will only become more apparent. We present Granger Causality as a means to establish directed functional connectivity in calcium fluorescence data, serving as a useful tool for guiding and confirming hypotheses alongside existing tools such as targeted neuronal ablation or opto-genetic stimulation. While our results on *in vivo* data are limited, the applicability and performance of Granger Causality will only increase as the acquisition hardware and the calcium reporter properties improve over time. Through the proposed future directions, we provide guidance both on further improving the methodology established in this work as well as setting a clear path for future works aimed at identifying the role of sensory-motor circuits in the control of behaviour through the unique opportunities offered by the directed functional connectivity measures.

# Bibliography

- Ahrens, M. B., Huang, K.-H., Narayan, S., Mensh, B. D., and Engert, F. (2013a). Two-photon calcium imaging during fictive navigation in virtual environments. *Frontiers in neural circuits*, 7:104. [8](#), [10](#)
- Ahrens, M. B., Li, J. M., Orger, M. B., Robson, D. N., Schier, A. F., Engert, F., and Portugues, R. (2012). Brain-wide neuronal dynamics during motor adaptation in zebrafish. *Nature*, 485(7399):471–477. [1](#), [3](#), [6](#), [7](#), [8](#), [9](#), [14](#), [27](#), [34](#), [35](#), [50](#), [54](#)
- Ahrens, M. B., Orger, M. B., Robson, D. N., Li, J. M., and Keller, P. J. (2013b). Whole-brain functional imaging at cellular resolution using light-sheet microscopy. *Nature methods*, 10(5):413. [1](#), [14](#), [106](#)
- Akaike, H. (1973). Maximum likelihood identification of gaussian autoregressive moving average models. *Biometrika*, 60(2):255–265. [56](#)
- Antinucci, P. and Hindges, R. (2016). A crystal-clear zebrafish for in vivo imaging. *Scientific reports*, 6(1):1–10. [5](#)
- Ardiel, E. L., Kumar, A., Marbach, J., Christensen, R., Gupta, R., Duncan, W., Daniels, J. S., Stuurman, N., Colón-Ramos, D., and Shroff, H. (2017). Visualizing calcium flux in freely moving nematode embryos. *Biophysical journal*, 112(9):1975–1983. [14](#)
- Avants, B. B., Tustison, N., and Song, G. (2009). Advanced normalization tools (ants). *Insight j*, 2(365):1–35. [93](#), [94](#)
- Bae, Y.-K., Kani, S., Shimizu, T., Tanabe, K., Nojima, H., Kimura, Y., Higashijima, S.-i., and Hibi, M. (2009). Anatomy of zebrafish cerebellum and screen for mutations affecting its development. *Developmental biology*, 330(2):406–426. [101](#)
- Barker, A. J. and Baier, H. (2015). Sensorimotor decision making in the zebrafish tectum. *Current Biology*, 25(21):2804–2814. [1](#)

- Barnett, L., Barrett, A. B., and Seth, A. K. (2009). Granger causality and transfer entropy are equivalent for gaussian variables. *Physical review letters*, 103(23):238701. [11](#), [56](#)
- Barnett, L. and Bossomaier, T. (2012). Transfer entropy as a log-likelihood ratio. *Physical review letters*, 109(13):138105. [56](#)
- Barnett, L., Muthukumaraswamy, S. D., Carhart-Harris, R. L., and Seth, A. K. (2020). Decreased directed functional connectivity in the psychedelic state. *NeuroImage*, 209:116462. [72](#), [108](#)
- Barnett, L. and Seth, A. K. (2011). Behaviour of granger causality under filtering: theoretical invariance and practical application. *Journal of neuroscience methods*, 201(2):404–419. [12](#), [70](#), [74](#), [108](#)
- Barnett, L. and Seth, A. K. (2014). The mvgc multivariate granger causality toolbox: a new approach to granger-causal inference. *Journal of neuroscience methods*, 223:50–68. [56](#), [57](#), [59](#)
- Barnett, L. and Seth, A. K. (2017). Detectability of granger causality for subsampled continuous-time neurophysiological processes. *Journal of neuroscience methods*, 275:93–121. [12](#), [63](#), [108](#), [111](#)
- Bell, C. C., Han, V., and Sawtell, N. B. (2008). Cerebellum-like structures and their implications for cerebellar function. *Annu. Rev. Neurosci.*, 31:1–24. [101](#)
- Bianco, I. H., Kampff, A. R., and Engert, F. (2011). Prey capture behavior evoked by simple visual stimuli in larval zebrafish. *Frontiers in systems neuroscience*, 5:101. [1](#), [6](#), [34](#), [35](#)
- Brand, M., Granato, M., and Nüsslein-Volhard, C. (2002). Keeping and raising zebrafish. *Zebrafish: a practical approach*, pages 7–37. [33](#)
- Brewer, G., Boehler, M., Pearson, R., DeMaris, A., Ide, A., and Wheeler, B. (2008). Neuron network activity scales exponentially with synapse density. *Journal of neural engineering*, 6(1):014001. [72](#)
- Burgess, H. A. and Granato, M. (2007). Modulation of locomotor activity in larval zebrafish during light adaptation. *Journal of Experimental Biology*, 210(14):2526–2539. [27](#), [111](#)

- Chang, C. (1962). Isolation of neurotoxins from the venom of *Bungarus multicinctus* and their modes of neuromuscular blocking action. *Arch Pharmacodyn Ther*, 144:241–257. [6](#)
- Chen, T.-W., Wardill, T. J., Sun, Y., Pulver, S. R., Renninger, S. L., Baohan, A., Schreiter, E. R., Kerr, R. A., Orger, M. B., Jayaraman, V., et al. (2013). Ultrasensitive fluorescent proteins for imaging neuronal activity. *Nature*, 499(7458):295–300. [6](#), [12](#), [31](#), [59](#)
- Chen, X., Mu, Y., Hu, Y., Kuan, A. T., Nikitchenko, M., Randlett, O., Chen, A. B., Gavornik, J. P., Sompolinsky, H., Engert, F., et al. (2018a). Brain-wide organization of neuronal activity and convergent sensorimotor transformations in larval zebrafish. *Neuron*, 100(4):876–890. [106](#)
- Chen, X., Mu, Y., Hu, Y., Kuan, A. T., Nikitchenko, M., Randlett, O., Sompolinsky, H., Engert, F., and Ahrens, M. B. (2018b). Brainwide organization of neuronal activity and convergent sensorimotor transformations in larval zebrafish. *bioRxiv*, page 289413. [1](#), [5](#), [8](#), [9](#), [10](#), [77](#)
- Chhetri, R. K., Amat, F., Wan, Y., Höckendorf, B., Lemon, W. C., and Keller, P. J. (2015). Whole-animal functional and developmental imaging with isotropic spatial resolution. *Nature methods*, 12(12):1171–1178. [14](#)
- Cho, J.-H., Swanson, C. J., Chen, J., Li, A., Lippert, L. G., Boye, S. E., Rose, K., Sivaramakrishnan, S., Chuong, C.-M., and Chow, R. H. (2017). The gcamp-r family of genetically encoded ratiometric calcium indicators. *ACS chemical biology*, 12(4):1066–1074. [63](#)
- Cox, K. J. and Fetcho, J. R. (1996). Labeling blastomeres with a calcium indicator: a non-invasive method of visualizing neuronal activity in zebrafish. *Journal of neuroscience methods*, 68(2):185–191. [6](#)
- Cunningham, J. P. and Byron, M. Y. (2014). Dimensionality reduction for large-scale neural recordings. *Nature neuroscience*, 17(11):1500. [10](#)
- Dana, H., Sun, Y., Mohar, B., Hulse, B. K., Kerlin, A. M., Hasseman, J. P., Tsegaye, G., Tsang, A., Wong, A., Patel, R., et al. (2019). High-performance calcium sensors for imaging activity in neuronal populations and microcompartments. *Nature methods*, 16(7):649–657. [7](#)

- Denk, W., Strickler, J. H., and Webb, W. W. (1990). Two-photon laser scanning fluorescence microscopy. *Science*, 248(4951):73–76. [7](#)
- Dohaku, R., Yamaguchi, M., Yamamoto, N., Shimizu, T., Osakada, F., and Hibi, M. (2019). Tracing of afferent connections in the zebrafish cerebellum using recombinant rabies virus. *Frontiers in neural circuits*, 13:30. [101](#)
- Dougllass, A. D., Kraves, S., Deisseroth, K., Schier, A. F., and Engert, F. (2008). Escape behavior elicited by single, channelrhodopsin-2-evoked spikes in zebrafish somatosensory neurons. *Current biology*, 18(15):1133–1137. [6](#)
- Dunn, T. W., Gebhardt, C., Naumann, E. A., Riegler, C., Ahrens, M. B., Engert, F., and Del Bene, F. (2016a). Neural circuits underlying visually evoked escapes in larval zebrafish. *Neuron*, 89(3):613–628. [101](#)
- Dunn, T. W., Mu, Y., Narayan, S., Randlett, O., Naumann, E. A., Yang, C.-T., Schier, A. F., Freeman, J., Engert, F., and Ahrens, M. B. (2016b). Brain-wide mapping of neural activity controlling zebrafish exploratory locomotion. *Elife*, 5:e12741. [31](#)
- Eccles, J. C. (2013). *The physiology of synapses*. Academic Press. [59](#), [64](#)
- Edelstein, A., Amodaj, N., Hoover, K., Vale, R., and Stuurman, N. (2010). Computer control of microscopes using  $\mu$ medelstein2010computeranager. *Current protocols in molecular biology*, 92(1):14–20. [15](#), [16](#), [88](#)
- Fallani, F. D. V., Corazzol, M., Sternberg, J. R., Wyart, C., and Chavez, M. (2014). Hierarchy of neural organization in the embryonic spinal cord: Granger-causality graph analysis of in vivo calcium imaging data. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 23(3):333–341. [3](#), [12](#), [55](#), [84](#), [87](#), [111](#), [113](#)
- Fardet, T., Deepu, R., Mitchell, J., Eppler, J. M., Spreizer, S., Hahne, J., Kitayama, I., Kubaj, P., Jordan, J., Morrison, A., et al. (2020). Nest 2.20. 0. Technical report, Computational and Systems Neuroscience. [58](#)
- Favre-Bulle, I. A., Vanwalleghem, G., Taylor, M. A., Rubinsztein-Dunlop, H., and Scott, E. K. (2018). Cellular-resolution imaging of vestibular processing across the larval zebrafish brain. *Current Biology*, 28(23):3711–3722. [83](#)
- Feierstein, C., Portugues, R., and Orger, M. (2015). Seeing the whole picture: a comprehensive imaging approach to functional mapping of circuits in behaving zebrafish. *Neuroscience*, 296:26–38. [8](#)

- Fetcho, J. R. and O'Malley, D. M. (1995). Visualization of active neural circuitry in the spinal cord of intact zebrafish. *Journal of neurophysiology*, 73(1):399–406. [6](#)
- Florin, E., Gross, J., Pfeifer, J., Fink, G. R., and Timmermann, L. (2010). The effect of filtering on granger causality based multivariate causality measures. *Neuroimage*, 50(2):577–588. [12](#), [70](#)
- Gahtan, E., Tanger, P., and Baier, H. (2005). Visual prey capture in larval zebrafish is controlled by identified reticulospinal neurons downstream of the tectum. *Journal of Neuroscience*, 25(40):9294–9303. [54](#), [101](#)
- Geweke, J. (1982). Measurement of linear dependence and feedback between multiple time series. *Journal of the American statistical association*, 77(378):304–313. [11](#), [56](#), [78](#)
- Geweke, J. F. (1984). Measures of conditional linear dependence and feedback between time series. *Journal of the American Statistical Association*, 79(388):907–915. [11](#), [56](#), [78](#)
- Giovannucci, A., Friedrich, J., Gunn, P., Kalfon, J., Koay, S. A., Taxidis, J., Najafi, F., Gauthier, J. L., Zhou, P., Tank, D. W., et al. (2018). Caiman: An open source tool for scalable calcium imaging data analysis. *bioRxiv*, page 339564. [2](#), [9](#), [89](#), [90](#)
- Granger, C. W. (1969). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: journal of the Econometric Society*, pages 424–438. [3](#), [11](#), [55](#)
- Granger, C. W. (2004). Time series analysis, cointegration, and applications. *American Economic Review*, 94(3):421–425. [56](#)
- Gualda, E. J., Vale, T., Almada, P., Feijó, J. A., Martins, G. G., and Moreno, N. (2013). Openspinmicroscopy: an open-source integrated microscopy platform. *Nature methods*, 10(7):599. [2](#), [8](#), [15](#), [27](#)
- Haesemeyer, M., Robson, D. N., Li, J. M., Schier, A. F., and Engert, F. (2018). A brain-wide circuit model of heat-evoked swimming behavior in larval zebrafish. *Neuron*, 98(4):817–831. [83](#)
- Halloran, M. C., Sato-Maeda, M., Warren, J. T., Su, F., Lele, Z., Krone, P. H., Kuwada, J. Y., and Shoji, W. (2000). Laser-induced gene expression in specific cells of transgenic zebrafish. *Development*, 127(9):1953–1960. [5](#)

- Hannan, E. J. and Quinn, B. G. (1979). The determination of the order of an autoregression. *Journal of the Royal Statistical Society: Series B (Methodological)*, 41(2):190–195. [57](#)
- Harris, C. R., Millman, K. J., van der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N. J., Kern, R., Picus, M., Hoyer, S., van Kerkwijk, M. H., Brett, M., Haldane, A., del R'io, J. F., Wiebe, M., Peterson, P., G'erard-Marchant, P., Sheppard, K., Reddy, T., Weckesser, W., Abbasi, H., Gohlke, C., and Oliphant, T. E. (2020). Array programming with NumPy. *Nature*, 585(7825):357–362. [69](#)
- Helmchen, F. and Denk, W. (2005). Deep tissue two-photon microscopy. *Nature methods*, 2(12):932–940. [7](#)
- Herrera, K. J., Panier, T., Guggiana-Nilo, D., and Engert, F. (2021). Larval zebrafish use olfactory detection of sodium and chloride to avoid salt water. *Current Biology*, 31(4):782–793. [6](#)
- Hesse, J. and Gross, T. (2014). Self-organized criticality as a fundamental property of neural systems. *Frontiers in systems neuroscience*, 8:166. [65](#)
- Hesse, W., Möller, E., Arnold, M., and Schack, B. (2003). The use of time-variant eeg granger causality for inspecting directed interdependencies of neural assemblies. *Journal of neuroscience methods*, 124(1):27–44. [70](#), [108](#)
- Higashijima, S.-i., Masino, M. A., Mandel, G., and Fetcho, J. R. (2003). Imaging neuronal activity during zebrafish behavior with a genetically encoded calcium indicator. *Journal of neurophysiology*, 90(6):3986–3997. [5](#), [6](#)
- Hildebrand, D. G. C., Cicconet, M., Torres, R. M., Choi, W., Quan, T. M., Moon, J., Wetzel, A. W., Champion, A. S., Graham, B. J., Randlett, O., et al. (2017). Whole-brain serial-section electron microscopy in larval zebrafish. *Nature*, 545(7654):345–349. [10](#), [98](#)
- Hu, B., Bolus, D., and Brown, J. Q. (2017). Improved contrast in inverted selective plane illumination microscopy of thick tissues using confocal detection and structured illumination. *Biomedical Optics Express*, 8(12):5546–5559. [29](#)
- Huisken, J. and Stainier, D. Y. (2009). Selective plane illumination microscopy techniques in developmental biology. *Development*, 136(12):1963–1975. [14](#)

- Huisken, J., Swoger, J., Del Bene, F., Wittbrodt, J., and Stelzer, E. H. (2004). Optical sectioning deep inside live embryos by selective plane illumination microscopy. *Science*, 305(5686):1007–1009. [1](#), [7](#)
- Kanal, E., Ozkurt, T., Sciabassi, R. J., and Sun, M. (2009). Detecting granger causality in the corticostriatal learning and rewards network using meg. In *2009 IEEE 35th Annual Northeast Bioengineering Conference*, pages 1–2. IEEE. [70](#)
- Kaur, P., Saunders, T. E., and Tolwinski, N. S. (2017). Coupling optogenetics and light-sheet microscopy, a method to study wnt signaling during embryogenesis. *Scientific reports*, 7(1):16636. [10](#)
- Kawashima, T., Zwart, M. F., Yang, C.-T., Mensh, B. D., and Ahrens, M. B. (2016). The serotonergic system tracks the outcomes of actions to mediate short-term motor learning. *Cell*, 167(4):933–946. [9](#), [59](#)
- Keller, G. B., Bonhoeffer, T., and Hübener, M. (2012). Sensorimotor mismatch signals in primary visual cortex of the behaving mouse. *Neuron*, 74(5):809–815. [50](#)
- Keller, P. J. and Ahrens, M. B. (2015). Visualizing whole-brain activity and development at the single-cell level using light-sheet microscopy. *Neuron*, 85(3):462–83. [14](#)
- Keller, P. J., Schmidt, A. D., Santella, A., Khairy, K., Bao, Z., Wittbrodt, J., and Stelzer, E. H. (2010). Fast, high-contrast imaging of animal development with scanned light sheet–based structured-illumination microscopy. *Nature methods*, 7(8):637. [14](#)
- Kim, D. H., Kim, J., Marques, J. C., Grama, A., Hildebrand, D. G., Gu, W., Li, J. M., and Robson, D. N. (2017). Pan-neuronal calcium imaging with cellular resolution in freely swimming zebrafish. *Nature methods*, 14(11):1107–1114. [8](#), [34](#)
- Kobitski, A. Y., Otte, J. C., Takamiya, M., Schäfer, B., Mertes, J., Stegmaier, J., Rastegar, S., Rindone, F., Hartmann, V., Stotzka, R., et al. (2015). An ensemble-averaged, cell density-based digital model of zebrafish embryo development derived from light-sheet microscopy data with single-cell resolution. *Scientific reports*, 5:8601. [14](#)
- Kubo, F., Hablitzel, B., Dal Maschio, M., Driever, W., Baier, H., and Arrenberg, A. B. (2014). Functional architecture of an optic flow-responsive area that drives horizontal eye movements in zebrafish. *Neuron*, 81(6):1344–1359. [1](#), [6](#)

- Kunst, M., Laurell, E., Mokayes, N., Kramer, A., Kubo, F., Fernandes, A. M., Förster, D., Dal Maschio, M., and Baier, H. (2019). A cellular-resolution atlas of the larval zebrafish brain. *Neuron*, 103(1):21–38. [9](#), [10](#), [77](#), [93](#), [94](#), [96](#), [98](#), [99](#), [101](#), [109](#)
- Lawrence, C. (2011). Advances in zebrafish husbandry and management. *Methods in cell biology*, 104:429–451. [5](#)
- Lehmann, D. (1984). Eeg assessment of brain activity: spatial aspects, segmentation and imaging. *International Journal of Psychophysiology*, 1(3):267–276. [54](#)
- Lemon, W. C., Pulver, S. R., Hockendorf, B., McDole, K., Branson, K., Freeman, J., and Keller, P. J. (2015). Whole-central nervous system functional imaging in larval drosophila. *Nat Commun*, 6:7924. [14](#)
- Levinson, N. (1946). The wiener (root mean square) error criterion in filter design and prediction. *Journal of Mathematics and Physics*, 25(1-4):261–278. [57](#)
- Ling, C. X., Huang, J., Zhang, H., et al. (2003). Auc: a statistically consistent and more discriminating measure than accuracy. In *Ijcai*, volume 3, pages 519–524. [72](#)
- Lister, J. A., Robertson, C. P., Lepage, T., Johnson, S. L., and Raible, D. W. (1999). Nacre encodes a zebrafish microphthalmia-related protein that regulates neural-crest-derived pigment cell fate. *Development*, 126(17):3757–3767. [5](#)
- Lobo, J. M., Jiménez-Valverde, A., and Real, R. (2008). Auc: a misleading measure of the performance of predictive distribution models. *Global ecology and Biogeography*, 17(2):145–151. [72](#)
- Lu, R., Sun, W., Liang, Y., Kerlin, A., Bierfeld, J., Seelig, J. D., Wilson, D. E., Scholl, B., Mohar, B., Tanimoto, M., et al. (2017). Video-rate volumetric functional imaging of the brain at synaptic resolution. *Nature neuroscience*, 20(4):620. [7](#)
- Markov, D. A., Petrucco, L., Kist, A. M., and Portugues, R. (2020). A cerebellar internal model calibrates a feedback controller involved in sensorimotor control. *bioRxiv*. [34](#), [50](#)
- Marques, J. C., Li, M., Schaak, D., Robson, D. N., and Li, J. M. (2020). Internal state dynamics shape brainwide activity and foraging behaviour. *Nature*, 577(7789):239–243. [34](#), [44](#)
- Merkel, D. (2014). Docker: lightweight linux containers for consistent development and deployment. *Linux journal*, 2014(239):2. [90](#)

- Metz, C. E. (1978). Basic principles of roc analysis. In *Seminars in nuclear medicine*, volume 8, pages 283–298. Elsevier. [72](#)
- Mikami, Y., Yoshida, T., Matsuda, N., and Mishina, M. (2004). Expression of zebrafish glutamate receptor  $\delta 2$  in neurons with cerebellum-like wiring. *Biochemical and biophysical research communications*, 322(1):168–176. [101](#)
- Miller, K. J., Shenoy, P., Miller, J. W., Rao, R. P., Ojemann, J. G., et al. (2007). Real-time functional brain mapping using electrocorticography. *Neuroimage*, 37(2):504–507. [54](#)
- Morf, M., Vieira, A., Lee, D. T., and Kailath, T. (1978). Recursive multichannel maximum entropy spectral estimation. *IEEE Transactions on Geoscience Electronics*, 16(2):85–94. [57](#)
- Nakai, J., Ohkura, M., and Imoto, K. (2001). A high signal-to-noise ca 2+ probe composed of a single green fluorescent protein. *Nature biotechnology*, 19(2):137–141. [6](#)
- Nalatore, H., Ding, M., and Rangarajan, G. (2007). Mitigating the effects of measurement noise on granger causality. *Physical Review E*, 75(3):031123. [74](#)
- Naumann, E. A., Fitzgerald, J. E., Dunn, T. W., Rihel, J., Sompolinsky, H., and Engert, F. (2016). From whole-brain data to functional circuit models: the zebrafish optomotor response. *Cell*, 167(4):947–960. [8](#), [10](#), [35](#), [98](#), [101](#), [104](#)
- Neil, J. J. (1993). Functional imaging of the central nervous system using magnetic resonance imaging and positron emission tomography. *Current opinion in neurology*, 6(6):927–933. [54](#)
- Oldfield, C. S., Grossrubatscher, I., Chávez, M., Hoagland, A., Huth, A. R., Carroll, E. C., Prendergast, A., Qu, T., Gallant, J. L., Wyart, C., et al. (2020). Experience, circuit dynamics, and forebrain recruitment in larval zebrafish prey capture. *Elife*, 9:e56619. [3](#), [8](#), [10](#), [12](#), [85](#), [87](#), [111](#), [113](#)
- Olive, R., Wolf, S., Dubreuil, A., Bormuth, V., Debrégeas, G., and Candelier, R. (2016). Rheotaxis of larval zebrafish: behavioral study of a multi-sensory process. *Frontiers in systems neuroscience*, 10:14. [6](#)
- Orger, M. B. (2016). The cellular organization of zebrafish visuomotor circuits. *Current Biology*, 26(9):R377–R385. [10](#)

- Orger, M. B. and de Polavieja, G. G. (2017). Zebrafish behavior: opportunities and challenges. *Annual review of neuroscience*, 40:125–147. [1](#), [34](#)
- Orlandi, J. G., Stetter, O., Soriano, J., Geisel, T., and Battaglia, D. (2014). Transfer entropy reconstruction and labeling of neuronal connections from simulated calcium imaging. *PloS one*, 9(6):e98842. [114](#)
- Pachitariu, M., Stringer, C., Dipoppa, M., Schröder, S., Rossi, L. F., Dalgleish, H., Carandini, M., and Harris, K. D. (2017). Suite2p: beyond 10,000 neurons with standard two-photon microscopy. *BioRxiv*. [9](#)
- Paluš, M., Komárek, V., Hrnčíř, Z., and Štěrbová, K. (2001). Synchronization as adjustment of information rates: Detection from bivariate time series. *Physical Review E*, 63(4):046211. [56](#)
- Patterson, G. H. and Piston, D. W. (2000). Photobleaching in two-photon excitation microscopy. *Biophysical journal*, 78(4):2159–2162. [86](#), [113](#), [114](#)
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830. [93](#)
- Peirce, J. W. (2009). Generating stimuli for neuroscience using psychopy. *Frontiers in neuroinformatics*, 2:10. [46](#)
- Pereda, E., Quiroga, R. Q., and Bhattacharya, J. (2005). Nonlinear multivariate analysis of neurophysiological signals. *Progress in neurobiology*, 77(1-2):1–37. [11](#), [57](#)
- Perin, R., Berger, T. K., and Markram, H. (2011). A synaptic organizing principle for cortical neuronal groups. *Proceedings of the National Academy of Sciences*, 108(13):5419–5424. [72](#)
- Pichler, P. and Lagnado, L. (2019). The transfer characteristics of hair cells encoding mechanical stimuli in the lateral line of zebrafish. *Journal of Neuroscience*, 39(1):112–124. [31](#)
- Pitrone, P. G., Schindelin, J., Stuyvenberg, L., Preibisch, S., Weber, M., Eliceiri, K. W., Huisken, J., and Tomancak, P. (2013). Openspim: an open-access light-sheet microscopy platform. *nature methods*, 10(7):598. [2](#), [8](#), [15](#), [27](#)

- Piyawattanametha, W., Barretto, R. P., Ko, T. H., Flusberg, B. A., Cocker, E. D., Ra, H., Lee, D., Solgaard, O., and Schnitzer, M. J. (2006). Fast-scanning two-photon fluorescence imaging based on a microelectromechanical systems two-dimensional scanning mirror. *Optics letters*, 31(13):2018–2020. [7](#), [110](#)
- Pnevmatikakis, E. A., Soudry, D., Gao, Y., Machado, T. A., Merel, J., Pfau, D., Reardon, T., Mu, Y., Lacefield, C., Yang, W., et al. (2016). Simultaneous denoising, deconvolution, and demixing of calcium imaging data. *Neuron*, 89(2):285–299. [9](#), [70](#), [90](#)
- Portugues, R. and Engert, F. (2011). Adaptive locomotor behavior in larval zebrafish. *Frontiers in systems neuroscience*, 5:72. [8](#), [34](#), [37](#), [42](#), [44](#), [49](#)
- Portugues, R., Feierstein, C. E., Engert, F., and Orger, M. B. (2014). Whole-brain activity maps reveal stereotyped, distributed networks for visuomotor behavior. *Neuron*, 81(6):1328–1343. [1](#), [6](#), [8](#), [9](#), [54](#), [83](#)
- Portugues, R., Haesemeyer, M., Blum, M. L., and Engert, F. (2015). Whole-field visual motion drives swimming in larval zebrafish via a stochastic process. *Journal of Experimental Biology*, 218(9):1433–1443. [37](#)
- Portugues, R., Severi, K. E., Wyart, C., and Ahrens, M. B. (2013). Optogenetics in a transparent animal: circuit function in the larval zebrafish. *Current opinion in neurobiology*, 23(1):119–126. [54](#)
- Potes, C., Brunner, P., Gunduz, A., Knight, R. T., and Schalk, G. (2014). Spatial and temporal relationships of electrocorticographic alpha and gamma activity during auditory processing. *Neuroimage*, 97:188–195. [108](#)
- Power, R. M. and Huisken, J. (2017). A guide to light-sheet fluorescence microscopy for multiscale imaging. *Nature methods*, 14(4):360. [14](#)
- Privat, M., Romano, S. A., Pietri, T., Jouary, A., Boulanger-Weill, J., Elbaz, N., Duchemin, A., Soares, D., and Sumbre, G. (2019). Sensorimotor transformations in the zebrafish auditory system. *Current Biology*, 29(23):4010–4023. [6](#)
- Randlett, O., Wee, C. L., Naumann, E. A., Nnaemeka, O., Schoppik, D., Fitzgerald, J. E., Portugues, R., Lacoste, A. M., Riegler, C., Engert, F., et al. (2015). Whole-brain activity mapping onto a zebrafish brain atlas. *Nature methods*, 12(11):1039–1046. [9](#), [93](#), [94](#), [109](#)

- Renninger, S. L. and Orger, M. B. (2013). Two-photon imaging of neural population activity in zebrafish. *Methods*, 62(3):255–267. [7](#)
- Reynaud, E. G., Peychl, J., Huisken, J., and Tomancak, P. (2015). Guide to light-sheet microscopy for adventurous biologists. *Nature methods*, 12(1):30–34. [8](#)
- Rieckher, M., Kyparissidis-Kokkinidis, I., Zacharopoulos, A., Kourmoulakis, G., Tavernarakis, N., Ripoll, J., and Zacharakis, G. (2015). A customized light sheet microscope to measure spatio-temporal protein dynamics in small model organisms. *PLoS one*, 10(5). [14](#)
- Roeser, T. and Baier, H. (2003). Visuomotor behaviors in larval zebrafish after gfp-guided laser ablation of the optic tectum. *Journal of Neuroscience*, 23(9):3726–3734. [10](#), [54](#)
- Ronneberger, O., Liu, K., Rath, M., Rueß, D., Mueller, T., Skibbe, H., Drayer, B., Schmidt, T., Filippi, A., Nitschke, R., et al. (2012). Vibe-z: a framework for 3d virtual colocalization analysis in zebrafish larval brains. *Nature methods*, 9(7):735–742. [9](#)
- Saska, D., Pichler, P., Qian, C., Buckley, C. L., and Lagnado, L. (2021).  $\mu$ spim toolset: A software platform for selective plane illumination microscopy. *Journal of Neuroscience Methods*, 347:108952. [106](#)
- Schindelin, J., Rueden, C. T., Hiner, M. C., and Eliceiri, K. W. (2015). The imagej ecosystem: An open platform for biomedical image analysis. *Molecular reproduction and development*, 82(7-8):518–529. [16](#)
- Schreiber, T. (2000). Measuring information transfer. *Physical review letters*, 85(2):461. [11](#), [56](#), [114](#)
- Schwarz, G. et al. (1978). Estimating the dimension of a model. *Annals of statistics*, 6(2):461–464. [57](#)
- Semmelhack, J. L., Donovan, J. C., Thiele, T. R., Kuehn, E., Laurell, E., and Baier, H. (2014). A dedicated visual pathway for prey detection in larval zebrafish. *Elife*, 3:e04878. [8](#), [34](#), [35](#), [37](#)
- Seth, A. K., Barrett, A. B., and Barnett, L. (2015). Granger causality analysis in neuroscience and neuroimaging. *Journal of Neuroscience*, 35(8):3293–3297. [3](#), [12](#), [55](#), [111](#)
- Seth, A. K., Chorley, P., and Barnett, L. C. (2013). Granger causality analysis of fmri bold signals is invariant to hemodynamic convolution but not downsampling. *Neuroimage*, 65:540–555. [3](#), [12](#), [65](#), [85](#), [87](#), [111](#)

- Severi, K. E., Portugues, R., Marques, J. C., O'Malley, D. M., Orger, M. B., and Engert, F. (2014). Neural control and modulation of swimming speed in the larval zebrafish. *Neuron*, 83(3):692–707. [8](#), [34](#), [35](#), [48](#), [49](#), [51](#), [52](#), [101](#), [104](#), [111](#)
- Silfverhuth, M. J., Hintsala, H., Kortelainen, J., and Seppänen, T. (2012). Experimental comparison of connectivity measures with simulated eeg signals. *Medical & biological engineering & computing*, 50(7):683–688. [114](#)
- Stetter, O., Battaglia, D., Soriano, J., and Geisel, T. (2012). Model-free reconstruction of excitatory neuronal connectivity from calcium imaging signals. *PLoS computational biology*, 8(8):e1002653. [11](#), [55](#), [58](#)
- Stokes, P. A. and Purdon, P. L. (2017). A study of problems encountered in granger causality analysis from a neuroscience perspective. *Proceedings of the national academy of sciences*, 114(34):E7063–E7072. [56](#)
- Stosiek, C., Garaschuk, O., Holthoff, K., and Konnerth, A. (2003). In vivo two-photon calcium imaging of neuronal networks. *Proceedings of the National Academy of Sciences*, 100(12):7319–7324. [7](#)
- Symvoulidis, P., Lauri, A., Stefanoiu, A., Cappetta, M., Schneider, S., Jia, H., Stelzl, A., Koch, M., Perez, C. C., Myklatun, A., et al. (2017). Neubtracker—imaging neurobehavioral dynamics in freely behaving fish. *Nature methods*, 14(11):1079–1082. [8](#)
- Tsodyks, M., Uziel, A., and Markram, H. (2000). Synchrony Generation in Recurrent Networks with Frequency-Dependent Synapses. *The Journal of Neuroscience*, 20(1):RC50. [58](#)
- Vanwallegem, G., Heap, L. A., and Scott, E. K. (2017). A profile of auditory-responsive neurons in the larval zebrafish brain. *Journal of Comparative Neurology*, 525(14):3031–3043. [3](#), [6](#), [12](#), [85](#), [87](#), [110](#), [111](#)
- Vanwallegem, G., Schuster, K., Taylor, M. A., Favre-Bulle, I. A., and Scott, E. K. (2020). Brain-wide mapping of water flow perception in zebrafish. *Journal of Neuroscience*, 40(21):4130–4144. [55](#)
- Vanwallegem, G. C., Ahrens, M. B., and Scott, E. K. (2018). Integrative whole-brain neuroscience in larval zebrafish. *Current opinion in neurobiology*, 50:136–145. [7](#)

- Vladimirov, N., Mu, Y., Kawashima, T., Bennett, D. V., Yang, C.-T., Looger, L. L., Keller, P. J., Freeman, J., and Ahrens, M. B. (2014). Light-sheet functional imaging in fictively behaving zebrafish. *Nature methods*, 11(9):883–884. [35](#), [106](#)
- Vladimirov, N., Wang, C., Höckendorf, B., Pujala, A., Tanimoto, M., Mu, Y., Yang, C.-T., Wittenbach, J. D., Freeman, J., Preibisch, S., et al. (2018). Brain-wide circuit interrogation at the cellular level guided by online analysis of neuronal function. *Nature methods*, 15(12):1117–1125. [106](#)
- Vogelstein, J. T., Watson, B. O., Packer, A. M., Yuste, R., Jedynak, B., and Paninski, L. (2009). Spike inference from calcium imaging using sequential monte carlo methods. *Biophysical journal*, 97(2):636–655. [58](#)
- Wan, Y., McDole, K., and Keller, P. J. (2019). Light-sheet microscopy and its potential for understanding developmental processes. *Annual review of cell and developmental biology*, 35:655–681. [14](#)
- Wang, S., Chen, Y., Ding, M., Feng, J., Stein, J. F., Aziz, T. Z., and Liu, X. (2007). Revealing the dynamic causal interdependence between neural and muscular signals in parkinsonian tremor. *Journal of the Franklin Institute*, 344(3-4):180–195. [70](#)
- Weisenburger, S. and Vaziri, A. (2018). A guide to emerging technologies for large-scale and whole-brain optical imaging of neuronal activity. *Annual review of neuroscience*, 41:431–452. [14](#)
- White, R. M., Sessa, A., Burke, C., Bowman, T., LeBlanc, J., Ceol, C., Bourque, C., Dovey, M., Goessling, W., Burns, C. E., et al. (2008). Transparent adult zebrafish as a tool for in vivo transplantation analysis. *Cell stem cell*, 2(2):183–189. [5](#)
- Whittle, P. (1963). On the fitting of multivariate autoregressions, and the approximate canonical factorization of a spectral density matrix. *Biometrika*, 50(1-2):129–134. [57](#)
- Wiggins, R. A. and Robinson, E. A. (1965). Recursive solution to the multichannel filtering problem. *Journal of Geophysical Research*, 70(8):1885–1891. [57](#)
- Wolf, S., Dubreuil, A. M., Bertoni, T., Böhm, U. L., Bormuth, V., Candelier, R., Karpenko, S., Hildebrand, D. G., Bianco, I. H., Monasson, R., et al. (2017). Sensorimotor computation underlying phototaxis in zebrafish. *Nature Communications*, 8(1):651. [1](#)

Wolf, S., Supatto, W., Debrégeas, G., Mahou, P., Kruglik, S. G., Sintes, J.-M., Beaupaire, E., and Candelier, R. (2015). Whole-brain functional imaging with two-photon light-sheet microscopy. *Nature methods*, 12(5):379–380. [8](#), [110](#), [112](#)