

Whole body-brain functional imaging and interrogation platform: technology development, analysis methods and applications

Virginia Marie Sophie Ruetten

**Gatsby Computational Neuroscience Unit
University College London**

**Janelia Research Campus
Howard Hughes Medical Institute**



This dissertation was supervised by
Prof. M. Sahani and Dr. M. B. Ahrens

This dissertation is submitted for the degree of
Doctor of Philosophy

Declaration

I, Virginia Marie Sophie Ruetten, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Acknowledgments

I attribute and have to thank for the best parts of this thesis to the environment I had the opportunity to work in, and to the people I had the chance to interact with. The flaws and gaps in this work lie on my shoulders, which I hope I will be able to address with time, and the reader forgive for now. Regardless of outcome, I am grateful for the years that I was given to create and build up this project; through it I found meaning, a sense of "Ikigai" and much joy.

Beyond immediate contributions, countless discussions with mentors and friends both at the Gatsby, Janelia and outside shaped this thesis and experience. A list of names would feel too reductionist for those included and meaningless by its length if comprehensive. Thus, I refrain and thank you, the reader, if you are part of my world.

Last but not least, I thank my supervisors, Maneesh Sahani and Misha Ahrens. Giving someone space in a lab is one thing; giving them space to grow is a whole other. By far my greatest critics, they have also been my strongest supporters. Most labs run a single ship, have a single mission, hold one overarching theme; they, in contrast, are willing to let out buoys with sails set for uncharted lands and, as long as the mission is sufficiently motivated and planned, equip, help, and cheer them on. Mistakes, dead ends, U-turns, of which I have to admit quite a few, were not met with frowns or an avalanche of stop signs; instead, I was asked how I intended to address the issues, what barriers needed to be lifted to move forward, and considered advice offered. Allowing someone to succeed is cheap and easy; allowing them to fail in the hope of achieving something greater, that comes at a cost. They gave me enough time to turn mistakes into lessons rather than having to bury them in regrets. When doubts or disagreements arose, questions abounded, and discussions became relentless, but never would the buoy wheel be threatened to be taken back ashore. Such things have an impact; they humble you and fill you with the desire to do better. Beyond their mentorship and friendship, I am also grateful for having had them as examples: at sea, it is easy to get distracted by ripples, to pay too much attention to the speed of neighboring boats, or mistake an island for a continent. Both seem to be oblivious to these sirens. "Why do we care?" was one of the most recurring questions posed. In a time when much pressure can be felt for the need to publish, the need to cash in quick results, cut corners, through them both, through getting to work at the Gatsby and at Janelia, I got to live my Ph.D. in another world, a world in which the only apparent metrics of achievement are the depth of one's understanding and one's distance to some core truth of nature. Both goals are still far, perhaps receding horizons, perhaps even slightly misaligned with society's current short-term reward system, but they are ones that, in the end, let you look back at the last 6 years of your life with great fondness and gratitude, rather than doubt if you should have ever started the trip. I can only hope that they too enjoyed the journey.

Other direct contributors to work presented in the thesis are outlined below.

Whole body imaging methods. This section of work received help from Caiying Guo and Sujatha Narayan who assisted with cloning. The Janelia vivarium team oversaw the husbandry of the animals as well as the injections of constructs. Optical advice was received from Philipp Keller and Michael DeSantis. The registration algorithm was designed and coded by Wei Zheng (Guoqiang Yu lab, Virginia Tech). Mika Rubinov, Talley Lambert and Greg Fleishmang helped by adapting their code bases (Voluseg, ND2 and BigStream respectively), to make them compatible with whole-body imaging pipeline.

Expansion microscopy work. This work was done under mentorship and in collaboration with Paul Tillberg, Amy Hu, and Mark Eddison. The 3D model was produced with the help of Alyson Petruncio and Aparna Dev who manually annotated ExM datasets and Igor Siwanowicz who from these datasets built the 3D mesh models of the fish.

SCA/GPFADS. The SCA and GPFADS work was done under mentorship and in collaboration with Guillaume Hennequin and Alberto Bernacchia. The hippocampal data used to demonstrate the advantages of SCA was provided by Brian Lustig and Albert Lee.

Funding. I thank the Gatsby Charitable Foundation and Janelia for funding me through my Ph.D. studies and for my Janelia Graduate Fellowship which offered me extra flexibility.

Contents

1	Introduction	7
1.1	Distributed bodily computations	7
1.2	A simple vertebrate platform to probe whole-body physiology	8
1.3	Seeding new grounds	8
2	Methodology: Whole Body Functional Imaging	10
2.1	Introduction	10
2.2	Transgenic Lines	11
2.3	Microscopy choice for whole body imaging	12
2.4	Registration	13
2.4.1	Motion correction	14
2.4.2	Image processing scheme	15
2.4.3	Validation	16
2.5	Post-registration data processing pipeline	19
2.5.1	Segmentation	19
2.5.2	Clustering	19
2.5.3	De-mixing	19
2.6	Additional Method Details	21
2.6.1	Experimental Model and Subject Details	21
3	Anatomical Foundations: Establishing an anatomical and molecular reference model for Whole Body Imaging via Expansion Microscopy	22
3.1	Introduction	22
3.2	Background to Expansion Microscopy	23
3.3	Engineering a Whole Body ExM protocol	24
3.4	Whole Fish ExM IHC and FISH Protocols	25
3.5	Results and Applications	25
3.6	Next steps	33
4	Analysis Tools: Developing computational methods for analyzing Whole Body Dynamics	34
4.1	Introduction	34
4.2	Coherence based clustering	35
4.2.1	Coherence based K-means clustering	35
4.2.1.1	Introduction	35
4.2.1.2	Derivation	36
4.2.2	Coherence based Spectral Clustering	37
4.2.2.1	Introduction	37
4.2.2.2	Derivation	39
4.3	Sequential Component Analysis	42
4.3.1	Sequentiality and space-time decomposition	43
4.3.2	Results	44
4.4	Non-reversible Kernels and GPFADS	46
4.4.1	Introduction	46
4.4.2	Background: Learning latent dynamical models	46
4.4.3	Background: Gaussian Process Factor Analysis (GPFA)	47
4.4.4	Nonreversible Gaussian processes	48

4.4.5	Quantifying non-reversibility and decomposing multi-output GP covariances	48
4.4.6	Planar non-reversible processes	50
4.4.7	Fourier domain interpretation	51
4.4.8	Higher-dimensional non-reversible priors	51
4.4.9	Experiments	52
4.4.9.1	Non-reversible GP priors better capture dynamics	53
4.4.9.2	GPFADS: recovering embedded latent dynamical systems	53
4.4.9.3	Uncovering rotational dynamics in M1	54
4.4.10	Discussion	55
5	Application: Studying distributed physiology using Whole Body Imaging	57
5.1	Introduction	57
5.2	Label-free tissue type identification assisted by 3D anatomical model	59
5.3	Spectra of dynamics	60
5.4	Motor Coupled Visceral Activity	66
5.5	Ependymal Coupled Dynamics	68
6	Conclusion	73
7	References	75
8	Appendix	83
9	Research paper declaration	115

Abstract

Organisms evolved to overcome physiological challenges and exploit their environments. Many solutions are not localized to a specific cell type or a specific organ; instead, they are implemented in distributed networks of cells across different organ systems. Traditionally, physiological subsystems of organisms have been studied in isolation, making it challenging to connect different subfields of biology. Here we introduce a system for imaging cellular dynamics of virtually all cells in a vertebrate, read out via calcium fluctuations. Larval zebrafish expressing calcium indicators under a ubiquitous promoter were imaged in spontaneous conditions and under physiological stressors (e.g.: cold). We imaged calcium dynamics across brain and body, and developed a registration pipeline for computationally stabilizing the tissue across the entire animal. We observed a plethora of spatio-temporal footprints of multicellular processes, including pulsatile traveling waves along the kidney and ultraslow fluctuations in ependymal cell activity locked to whole-brain dynamics and motor activity. We adapted clustering and dimensionality reduction methods to better capture the types of multiscale dynamics prevalent in the data, proposing coherence and non-time reversibility as two informative metrics, as compared to correlation- or variance-based approaches. To assist in data interpretation and to be able to follow up such observations at cellular and sub-cellular scale, we developed two whole fish ExM protocols compatible with animals as old as 21dpf for immunostaining and RNA probing. Together, these tools provide a powerful platform to study whole-body physiology and physiology-driven decision making.

Impact Statement

This work sets out an alternative way to study vertebrate physiology which we believe will accelerate fundamental research. We show how tools extensively used in neuroscience including Ca^{2+} imaging, expansion microscopy and modelling can be applied at scale and take into consideration cellular activity from all across the body without excess loss of interpretability or need for labor-intensive post-hoc labelling. This approach, generalizable to other sensors and other transparent and optically accessible organisms, permits for a less biased screening and interrogation of cellular players involved in biological processes. Most cell types have not been functionally imaged to this date, fewer still functionally imaged *in-vivo* in an intact unanesthetized organism, or within the context of the ongoing activity of other cells types. We see this platform as a seed generator of new fields, particularly working at the interface of organ systems. This platform also lends itself to chemical and drug screening.

1 Introduction

1.1 Distributed bodily computations

Organisms evolved to overcome physiological challenges and exploit their environments. Many solutions are not localized to a specific cell type or a specific organ; instead, they are implemented in distributed networks where multiple cell types located across different organ systems jointly solve problems faced by organisms. Understanding organismal function must therefore consider at some level the organism as a whole. Indeed, a state of health is a state of dynamic stability,

instantiated through balanced interactions across the networks of molecules, cells, and tissues that make up an organism. However, with critical advances in molecular biology in the 20th century, the study of physiology, the study of functional circuits spanning the body, was overtaken by a more reductionist approach focusing on individual tissue and cell-types (Billman, 2020, Buchman, 2003, Dantzer, 2018). Indeed, entire fields became defined by their tissue type of interest: neuroscience, immunology, nephrology, cardiology, etc. However, in the last decades, a growing body of research is beginning to reemphasize the extent of cross-talk between organ systems: for instance, scientists have uncovered the mechanisms by which the brain can suppress the immune system, enhance antibody production, modulate liver regeneration, or how in turn inflammatory signals secreted by the immune system can alter neural activity to induce depression-like symptoms (Dantzer, 2018, Rosas-Ballina et al., 2011, Izumi et al., 2018, Zhang et al., 2020). Technological advances in the ability to image, monitor, control, and perturb tissue across many scales, as well as the increased computational power to handle and model big data, offer the opportunity to integrate what has been learned from individual fields back into a less cell-type centric framework to dissect the feedback mechanisms which restore balance and health at the scale of the organism.

1.2 A simple vertebrate platform to probe whole-body physiology

To understand a circuit, biologists as engineers, ideally would have: a) intricate knowledge of basic building blocks of the system, b) an ability to identify states and monitor changes in these, c) models and theory through which to extract principles and form hypothesis, d) means to perturb variables to test said models and theory. When combined, this allows for iteration between hypothesis generation and hypothesis testing, for causality to be established, effects of new perturbations predicted, and ultimately control over the system gained. Such affordances are often not features of the system of real interest (e.g.: human body), and thus researchers resort to studying simplified systems, models, in which they believe that the principles of interest are at work. An ideal model system should be as simple as it can be while retaining the key aspects of the original system that one wants to understand and endowed with as many of the above features. The larval zebrafish, first established as a model organism in genetics and development biology, is one example of such a simplified vertebrate system. Small, robust, it comes equipped with the full array of vital organs found in the human body and, crucially, is relatively transparent making it optically accessible (Choi et al., 2021). In the last decade, developments in microscopy have made it possible to image whole-brain activity using light-sheet imaging (Ahrens et al., 2013). As an already established model organism, there exists a large corpus of literature about the larval zebrafish, as well as a toolbox for genetic manipulation (Li et al., 2021b). Together these features make it a practical and versatile platform to dissect the input-output function of vertebrate physiology, offering the possibility to interrogate the organism as whole, while retaining cellular level access. Here, we make a first attempt at scaling existing methods used primarily to dissect neuronal computations, extending them to the rest of the body in the hope of providing a new inroad into the study of physiology and distributed bodily computations.

1.3 Seeding new grounds

The most significant contributions of this Ph.D. cannot be said to have been planned or to have arisen from great foresight but instead arose out of a curiosity to see more, coupled with the good fortune of being in an environment that had the resources and mindset to foster and enable this. With my initial tentative steps at studying body-brain dynamics and the vagus nerve, the

same challenges presented themselves over and over again: an inability to access the inner workings of the body, to see what was going on inside. Despite the incredible tools available to access and monitor neuronal activity, there was comparatively little available to interrogate the rest of the body. Indeed, while I could, for instance, precisely manipulate neural activity of genetically defined vagal subpopulations using optogenetics, record during such experiments the ongoing activity across the brain at respectable frame rates, have a suite of well-established methods to process and analyze the resulting neural data, lookup the projection and even RNA transcript expression patterns of these neurons in online neural atlases; in contrast to all of those tremendous resources, when wanting to find out what consequences such neuronal activity might have on any organ other than the brain or on general host physiology, one was faced with a relative dearth of tools and hard-to-find references for even the most basic of questions that one would hope to bypass. Thus, much of the Ph.D. came to be about filling this gap.

2 Methodology: Whole Body Functional Imaging

2.1 Introduction

Here we present how we adapted and combined existing methods to allow for whole-body functional imaging. An idea, even if simple in concept, can turn out to be unfeasible in practice. While doubts were present as to the possibility or practicality of whole-body functional imaging, many of the required building blocks were already in place, established by previous researchers, waiting to be combined. Indeed, the initial doubts as to the practicality of whole-body functional imaging turned out to be unfounded, others overcome; equally so, its limitations and strengths also became better delineated.

Imaging entire developing embryos at single cell resolution has been accomplished now in multiple species revealing to a greater extent the level of cellular coordination at such ages and allowing researchers to study at scale developmental trajectories of cells and their interactions (Keller et al., 2008, McDole et al., 2018, Wu et al., 2011). In parallel to this, imaging Ca^{2+} dynamics across thousands of cells is now main-stream in neuroscience (Ahrens et al., 2013, Sofroniew et al., 2016) and a range of highly optimized Ca^{2+} sensors have been developed enabling such experiments (Miyawaki et al., 1997, Chen et al., 2013, Dana et al., 2016, Dana et al., 2019). We took example from both of these fields, as well as others, to establish a means of imaging cellular activity across the body of an organism. The first requirement was to generate a stable, healthy transgenic line with pan-cellular GECI (genetically-expressed calcium indicators) expression (Section 2.2). Pan-cellular promoters had already been identified, and pan-cellular FP (fluorescent proteins) lines established (Mosimann et al., 2011) supporting the idea that this was possible. The next doubt was as to whether a) the sensors' dynamic range would capture non-neuronal Ca^{2+} dynamics and b) as to whether the density and non-specificity of the GECI expression would make the data unparsable, or difficultly interpretable. There was precedence for the first and fortuitously the second empirically turned out not to be an issue. Indeed, a growing body of labs had started expressing GCaMP, a GECI, in non-neuronal tissue (Mu et al., 2019, Lorincz et al., 2018, Pozo-Morales et al., 2023) so it was known that pan-cellular GCaMP expression could capture meaningful Ca^{2+} fluctuations in at least a subset of non-neuronal tissues. The first WBI experiments corroborated these findings: we found that most cells in the body exhibited some fluorescence changes, therefore that Ca^{2+} fluctuations were within the sensors range (Section 5.2); furthermore, a surprising amount of different tissue types could be easily differentiated by tissue texture alone (Section 5.3).

However, the first experiments run using light-sheet based microscopes also made it clear that optics optimized for whole-brain imaging were not well suited for WBI (whole-body imaging) due to the difference in geometry and scattering properties of the tissue. Thankfully, a broad range of microscopes exist from which to choose from, each with differing optical capabilities and many of them commercially available. We tested a few optical systems (Section 2.3) and found that spinning disk confocal or two-photon microscopes were the best suited for WBI, and we further adapted a two-photon microscope to enable acquisition of orthogonal sample views to gain more coverage. Upon data acquisition, we were then faced with the challenge of dealing with data from highly mobile and contractile tissues (e.g.: gut, gallbladder) which induce large motion artifacts. To facilitate registration, we sought a way to acquire a reference channel that would only be subject to motion and not Ca^{2+} fluctuations. The more wide-spread strategy of labeling nuclei, which allows users to extract a set discrete points, turned out not to work well so we opted for a different approach using an existing transgenic line labeling the membranes

of all cells. We then adapted a registration algorithm to estimate the motion field of the data (Section 2.4). Having overcome the worst of motion artifacts, we subsequently segmented the data into cell-like units using a pre-existing scheme (Section 2.5.1). The next challenge was to efficiently label the data, i.e.: enable one to easily extract dynamics from specific tissue types. Somewhat surprisingly, manual annotation based on morphology alone turned out to be much less informative and accurate than one could expect; instead, we found that clustering the data in activity space using spectral methods (Section 2.5.2) and backprojecting cluster labels into anatomical space resulted in clusters that relatively cleanly mapped back to specific tissue types, and that through this more distinct labels could be assigned, giving an easy way to annotate the data, adding much interpretability. Finally, we iterated over a final de-noising step to address the fact that high-amplitude signal from muscle activity can bleed into neighboring cells due to the axial extent of the microscope’s point spread function (Section 2.5.3). Together, these pieces address the main initial doubts and hurdles posed when approaching whole-body imaging, leading to a readily adoptable platform.

2.2 Transgenic Lines

Animals expressing genetically encoded fluorescent sensors enable researchers to image fluctuations in the concentration of intra- or extracellular molecules (Chen et al., 2013). Many lines have been made targeting neurons, yet few exist with respect to imaging such sensors in the rest of the body. Despite the simplicity of the idea of making a pan-cellular GECI line it took some attempts to succeed. Ubiquitous promoters have already been identified, and two papers have even driven GECI expression (Li et al., 2021a, Chen et al., 2017), however, expression was either limited to the few-cell stage or was extremely sparse. We tried three different pan-cellular promoters (b-actin, ubi, olactb), with them either failing or giving sparse expression, as in part previously reported (Yoshinari et al., 2012). To drive higher expression, we chose the most promising of the three, the ubiquitin promoter, and added the amplificatory element tTA/TRE (which is less prone to silencing than the usual gal4/uas system, as it is less dense in GC bases). The GCaMP variant GCaMP7f was used (Dana et al., 2019). The constructs were injected into fertilized eggs and embryos grown up. Once maturity reached, the animals were bred and screened for germ-line expression. From these adults, the ones whose progeny had the most homogenous pan-cellular expression profiles were kept. For 1-photon optogenetic experiments, exciting with a shorter wavelength than that used for imaging is preferred in order to minimize unwanted excitation of the opsin; thus, we thus also generated a stable pan-cellular red GECI line Tg(ubi:tTA;TRE:jREGCO1b) compatible with the use of channelrhodopsin and its variants (Fig. 17) (Dana et al., 2016).




Figure 1: Whole body imaging and label-free tissue type identification.

(A1) Pan-cellular Ca^{2+} indicator transgenic line allows for *in-vivo* functional whole body imaging. A1 show a sagittal section of 7 days post-fertilization Tg(ubi:tTA;TRE:GCaMP7f) larval zebrafish. Data is acquired with spinning-disk confocal. Some of the major organs identifiable without further labelling are annotated.

(A2) 3D model of a larval zebrafish generated from expansion microscopy datasets ([Section 5.2](#)). The organs annotated in A1 are shown in the 3D model.

(B) Larger magnification view of a sagittal section of the sample (same sample as shown in A1). The region encompasses the proximal tubules of the kidney, the liver and the gallbladder (area within a black square in A2 labelled B). This section is medial to that shown in A1. B2 shows the same region within the 3D model.

(C) Larger magnification view of a sagittal section of the sample (same sample as shown in A1). The region encompasses the cloaca (area within a black square in A2 labelled C). Distinct bands are visible corresponding to different tissues types. From dorsal to ventral: muscle, spinal cord, notochord, aorta, connective tissue, hindgut, skin epithelium of ventral fin. While one can infer the major cell-types in most of the bands (e.g.: myocyte, neuron, gut epithelial cell, etc), the band labeled "connective tissue" illustrates how this is not always the case.

2.3 Microscopy choice for whole body imaging

Various microscopes were tested and compared to assess which would provide the best overall performance in terms of sample coverage and signal to noise ratio (SNR). Due to the complex geometry of the sample, light-sheet-based microscopes (the field's standard), which require two perpendicular access directions (for illumination and collection), were found to suffer heavily

from scattering in the interior of the sample. We found that the better sectioning afforded by confocal microscopes outweighed the disadvantages of slower speeds and higher light doses. Visceral Ca^{2+} dynamics were found to be generally slower than that of neurons, mitigating the impact of low sampling rates. Therefore, we used primarily a spinning disk confocal (Nikon W1) to record activity throughout the body at about 0.3Hz, using a 20x 0.75 NA objective with resolution of $0.36\mu\text{m} \times 0.36\mu\text{m}$ (x-y). The samples could be continuously recorded for periods of up to 4-5 hours. Despite the high dose of light, the fish remained healthy and the bleaching was minimal (with GCaMP, although not jRGECO1b). The field of view typically imaged was $\sim 800\mu\text{m} \times 800\mu\text{m}$, which covers 1/4 of the fish. This small fraction covers most organ systems, as these are heavily concentrated in the anterior part of the fish (much of the length of the fish is its tail). For most recordings, a depth of up to about $200\mu\text{m}$ was imaged, at intervals of $9\mu\text{m}$. While this only covers half of the fish in terms of depth, organs are either present bilaterally (e.g. liver, kidney) or on the right side of the fish (e.g. endocrine pancreas, gallbladder); thus, by placing the sample with its right side parallel to the objective most organs could be surveyed simultaneously. Using an inverted spinning disk confocal has the down side of requiring the animal to be laid on its side which is not optimal for imaging neural activity as much of the midbrain is occluded by the heavily scattering ear of the fish. To access both the brain and the viscera, one ideally needs to image two orthogonal planes. For this, with the assistance of Thorlabs, a dual plane, dual wavelength two-photon microscope was assembled (with minor modifications from the original design [Sofroniew et al., 2016](#)). We could then mount the sample upright next to a 45-degree mirror and scan both brain and mirror, resulting in the acquisition of two orthogonal views. While such configuration is likely the way to go, data presented in this thesis comes from spinning-disk confocal microscopes; later work will make use of the ortho-view mesoscope.

2.4 Registration

Compared to the brain, the viscera exhibits larger, more rapid, and nonrigid motion (e.g.: intestinal/gallbladder contraction). As a result, changes in fluorescence can arise from changes in activity or motion. This tissue warping could not be corrected using standard registration algorithms developed for neural tissue registration. To address this, we first acquired a second channel that was only affected by motion and secondly developed a tailored algorithm to take into account the non-homogeneous nature of tissue motion. More precisely, to be able to more readily uncouple changes from motion from fluctuations in activity, we performed dual color imaging of both GECI and a (nonsensor) fluorescent protein from which to infer ground truth motion. Typically nuclear-localized fluorescent proteins are used for this; however, we found membrane-targeted FP (fluorescent protein) to be more information rich; therefore, pan-cellular GECI lines were crossed and co-imaged with a pan-cellular membrane-targeted static FP line: Tg(bactin2:mCherry-CAAX) fish ([Fig. 2A](#)). Next, we developed and validated an algorithm that iteratively estimates the motion flow field from the mCherry channel, which can then be applied to the GECI channel ([Fig. 2B](#)). This was done in collaboration with the Guoqiang Yu Lab (Virginia Tech), specifically with Wei Zheng. The algorithm was designed and coded by Wei Zheng. We analyzed and interpreted the results together and iterated on the basis of the findings. Given that this is not primarily my work, only a brief overview of the pipeline is presented.

The pipeline can be divided into two sections. One section deals with the actual issue of motion estimation, the other is the scaffolding around it to ensure that one converges to a good solution. This latter part includes the use of pyramid downsampling to avoid local minima, removing detractor elements from the data, ignoring areas with too little information, and using a floating

reference to account for slow changes in intensity. We first outline how motion is estimated and then present the full algorithm.

2.4.1 Motion correction

We call "template" the image which is considered static, and "moving image", the data that we aim to transform.

Iterative patch-wise optical flow

We want to minimize the intensity difference between the template and the moving image while keeping the motion field to vary smoothly over space. Given an intensity function $I(x, t)$ where x designates the spatial coordinate and t the time. We can formulate an optimization problem with the following loss function:

$$\mathcal{L}(\Delta x) = \sum_x^X (I(x, t + \Delta t) - I(x + \Delta x, t))^2 \quad (1)$$

where X is the total pixel number.

This original objective function is non-convex and cannot be solved directly. Here we use the optical flow algorithm to transform the original problem into a set of locally convex optimization problems, through the use of the Horn-Schunck function (Horn and Schunck, 1981). By iteratively solving the locally convex optimization problems and updating the moving image and mask, we can approach a local optima of the original non-convex problem. Optical flow assumes for some intensity function I which is a function of x and t :

$$\frac{dI(x, t)}{dt} = \frac{\partial I(x, t)}{\partial x} \frac{\partial x}{\partial t} + \frac{\partial I(x, t)}{\partial t} \frac{\partial t}{\partial t} \quad (2)$$

$$= \frac{\partial I(x, t)}{\partial x} \frac{\partial x}{\partial t} + \frac{\partial I(x, t)}{\partial t} \quad (3)$$

$$= 0 \quad (4)$$

Thus the loss can be defined as:

$$\mathcal{L}(\Delta x) = \left[I(x, t + \Delta t) - I(x, t) - \frac{\partial I(x, t)}{\partial x} \Delta x \right]^2 \quad (5)$$

We can iteratively estimate Δx by setting $I(x, t + \Delta t)$ as our reference image and $I(x, t)$ as our estimate of the moving frame. For brevity, we only consider the x coordinate, but in practice this is done in 3D.

To be able to get a better estimate of the spatial intensity gradient, we use a locally averaged estimate: the image is divided into patches of $(r \times r)$, and the average spatial gradient is computed. A single Δx is estimated and optimized per patch. Once convergence reached, we linearly interpolate motion between patch centers. We also want the estimated motion to be somewhat smooth over space - we therefore augment the loss function with a penalty term weighted by β which penalizes the difference between each motion estimate of a patches motion and the average of the immediately adjacent patches. The augmented loss now becomes:

$$\mathcal{L}(\Delta X) = \sum_n^N \sum_x^Z \left[I_n(x, t + \Delta t) - I_n(x, t) - \frac{I_n(x, t)}{\partial x} \Delta x^n \right]^2 - \beta(\Delta x^n - \overline{\Delta x^n})^2 \quad (6)$$

where $\Delta X \in R^N$, N is the total number of patches, Z is the number of pixels per patch, Δx^n is the estimate of the spatial derivative for patch n and $\overline{\Delta x^n}$ is the average of Δx^n for the neighbouring patches of n .

This is a quadratic function and thus has a closed form solution. Once ΔX computed, we linearly interpolate between patch centers and apply the motion to the moving frame, and iterate until the loss ceases to decrease beyond a user defined threshold. To take into account that it is the accumulated motion which ought to be spatially smooth, the terms weighted by β are the accumulated patch and patch neighbourhood motion. Furthermore, to be able to parallelize calculations, the gradients of the previous timesteps are used for the neighborhood motion estimation. The loss function is thus defined as:

$$\mathcal{L}(\Delta X) = \sum_n^N \sum_x^Z \left[I_n(x, t + \Delta t) - I_n(x, t) - \frac{\partial I_n(x, t)}{\partial x} \Delta x_a^n \right]^2 - \beta(\Delta x_a^n - \overline{\Delta x_a^n})^2 \quad (7)$$

where Δx_a^n is the accumulated motion of patch n and $\overline{\Delta x_a^n}$ the accumulated average of Δx^n for the neighbouring patches of n .

2.4.2 Image processing scheme

To make the algorithm more robust to noise and results more consistent careful preprocessing is needed. Here we outline the main steps of the scheme:

Method parameters

- β : smoothness penalty
- r : motion correction patch size (in units of number of pixels)
- N_L : number of pyramid layers used for motion correction
- N_{template} : number of frames to include in floating template
- N_{moving} : number of frames to consider when initializing motion

Algorithm 1 Image Processing for Motion Analysis

Require: sequence of time frames, model parameters ($\beta, r, N_L, N_{\text{template}}$)

for pyramid layer l from 1 to N_L **do**

Pyramid down-sampling. Down-sample raw data by 2^{N_L-l} in x and y

if $l > 1$ **then**

 Scale the motion flow field estimated from the previous iteration layer by half in x and y, apply it to the downsampled data

end if

for each time frame **do**

Moving Foreground Estimation. Define as background a) very bright connected components as background (often moving immune cells), b) regions without reliable texture, with statistically small spatial intensity gradient comparable to that of background data, c) sufficiently dark regions. Regions not classified as background are considered raw foreground.

Floating Template Calculation.

if there are more than N_{template} motion corrected frames **then**

 Update the floating template by the median of the last N_{template} motion corrected frames. Update the initialization of the motion field by the motion field closest to the center of the last N_{template} motion fields.

else

 Use the first frame as the floating template.

end if

Motion Field Initialization:

if there are more than N_{motion} motion corrected frames **then**

 Update the initialization of the motion field by the motion field closest to the center of the last N_{motion} motion fields.

else

 Set the motion field initialization to zero.

end if

Template Foreground Estimation:

Motion Mask Calculation:

 Find the mask of the area over which to estimate motion: the intersection of valid foreground of the template and raw foreground of the moving image.

Estimate motion ([Section 2.4.1](#))

end for

end for

Post-processing We apply the estimated motion field to the functional channel interpolating linearly.


2.4.3 Validation

To validate the algorithm, we simultaneously imaged the reference channel (pan-cellular membrane marker: Tg(bactin2:mCherry-CAAX)) and a line with sparse expression throughout the body (especially the gut, which is the hardest area to register: Tg(phox2bb:eGFP)) under normal conditions for an hour. We estimated the motion field using only the red (reference) channel and applied it to the held out green (sparse) channel and then computed the mean-squared error (MSE) between the first frame of the green channel and subsequent frames and


compared that to the MSE between the first frame and the motion corrected data (Fig. 2C). Despite strong gut motion, the registered data accumulates drift at a much slower rate (Fig. 2D). While not perfect (registration during strong gut movement will likely remain a challenge for some time) the results were deemed sufficiently good to proceed. All registered datasets are manually inspected to ascertain quality before proceeding to the next steps.

Assumptions and Advantages This method assumes that all the non-global fast intensity changes in the static membrane channel are only caused by motion. The method takes into account the nonisotropic nature of the measurements (lower resolution in z). The method can be parallelized patchwise so that calculation for a typical frame takes about 4 seconds to complete with v100 GPU. Pyramid downsampling and the smoothness penalty make the approach more likely to converge close to the global optima rather than being trapped in bad local optima, which leads to artifacts. The floating template allows us to gracefully deal with slow intensity changes that occur during long acquisitions due to photobleaching. The masking steps allows us to ignore unwanted intensity changes such as wandering immune cells. The code was written and run in MATLAB.

A Dual Color Acquisition



B Registration



C Validation




Figure 2: Registration of Whole Body Imaging datasets

(A) Dual color acquisition with reference channels allows for robust data registration. Image of a double transgenic animal expressing the functional sensor GCaMP (left) and the red static fluorescent protein mCherry targeted to cellular membranes (right).

(B) Example of typical tissue warping within the viscera and correction that can be achieved through registration. The motion field is estimated iterative. The left column shows an example of the flow field estimation at t=0 and after an hour over the region of swimbladder which deflates progressively over time. The middle column shows an example of typical drift which happens. The image shows the overlay of the GCaMP channel at the beginning of the experiment (magenta) and after one hour of recording (green). While the dorsal part of the sample remains in a similar location, the anterior portion of the gut and gall bladder are heavily shifted (shown in the inset of higher magnification below). After registering the data (right column), less positional drift is visible.

(C) Validation experiment. To quantitatively assess the registration algorithm, we performed an experiment in which a sparse line (Tg(phox2bb:eGFP)) was imaged along with the registration channel (pan-cellular membrane-targeted mCherry). The motion flow field is estimated using the reference channel and applied to held out sparse data in which motion can be easily quantified. We overlay the data from the sparse channel acquired at $t=0$ (green) and at $t=1h$ (magenta). After one hour there is significant drift, particularly visible in the gut region (higher-magnification views of the brain and gut are shown to the right of the image). Post registration (right column) there is less misalignment. We quantify this by computing the MSE between the first and subsequent frames captured (left column) of the sparse channel. The large changes in the blue curve (raw data) correspond to timepoints of movement which accumulate over time - these are not visible in the orange (registered) trace; nonetheless, there remains still some drift. Each registered dataset is manually inspected to ensure that any downstream analysis isn't unduly misled by uncorrected motion artifacts.

2.5 Post-registration data processing pipeline

2.5.1 Segmentation

Next, we functionally segment the data using Voluseg, a pipeline that implements locally constrained NMF to cluster neighboring correlated pixels into functional segments. This pipeline was published in [Mu et al., 2019](#). In short, an intensity-based brain mask is created and divides the volume into spatially contiguous three-dimensional blocks, which overlap slightly to capture cells on the borders, and runs cell detection in parallel on these, making use of the Spark cluster computing framework. The algorithm fits the following model:

$$V_{(n \times t)} \approx W_{(n \times c)} H_{(c \times t)} + X_{(n \times 1)} I_{(1 \times t)} \quad (8)$$

where V is the full spatiotemporal fluorescence matrix for each block, W and H are, respectively, the spatial location and time series of segmented cells, and X and I are rank-1 spatiotemporal models of the background signal.

For a subset of the data sets, volumes were hand-annotated to provide ground-truth labels for tissue types using the software Amira.

2.5.2 Clustering

To identify clusters, we run Spectral Clustering as described in [Section 4.2.2](#) using coherence as a similarity metric ($\tau = 0.3$ and $N_{\text{clusters}} = 300$). To find spatially continuous clusters, we run density-based clustering (DBSCAN) on each cluster with parameters: $\epsilon = 10\mu\text{m}$ and $\text{min_samples} = 5$ and keep the largest connected cluster discarding other cells.

2.5.3 De-mixing

Each camera pixel captures photons according to the point spread function (p.s.f.) of the microscope. As a result of the extent of the p.s.f. (especially in the axial direction), signal from out-of-focus planes can come to contaminate pixels, thereby leading to mixing of signals from various tissue types. This is typically not an issue if the numerical aperture of the optics

(which determines the optical sectioning) is adequately chosen for the application; however, such a problem can arise if there are large differences in fluorescence intensity between adjacent planes, as the number of photons from the out-of-focus plane can then come to dominate the measurement. In the context of whole body imaging, such a case occurs upon large muscle contractions: muscle cells are densely packed and, at first approximation, are synchronously recruited, resulting in large Ca^{2+} bursts. As a result of this, camera pixels that ought to collect light only from non-muscular cells can, if they are in close proximity to muscle, collect signal from both muscle and non-muscle cell, which then confounds downstream analysis. To demix such muscle contaminations, we first cluster the data (Section 2.5.2) and identify the clusters that correspond to muscle tissue to use as regressors. These are the most easily identifiable thanks to their stereotyped high variance dynamics and anatomical footprint. Once muscle clusters are identified, a clean up step is run pixel-wise to make sure that these clusters themselves are not contaminated. Muscle activity resembles neuronal activity in that it is well modeled as a series of delta functions convolved by a decaying exponential kernel resulting in positively skewed data (Stringer and Pachitariu, 2019). Most other cells types don't show so exclusively such stereotyped transients, such that a pixel which is not positively skew is unlikely to come from muscle tissue. We therefore exclude pixels whose skewness is below a user-set threshold. Next, we assume that at the time points of peak muscle contraction, the contribution to the change in fluorescence of an individual unit $\Delta y_t = y_t - y_{t-1}$ is a linear combination of the change in muscle fluorescence $\Delta x_t = x_t - x_{t-1}$ and some cell intrinsic changes, ϵ . In other words, for $t \in T_c$, where T_c is the set of time points at which muscle spikes occurred:

$$\Delta y_t = \beta \Delta x_t + \epsilon \quad (9)$$

Contamination is purely in the form of a linear and positive contribution (i.e., changes in muscle fluorescence can only increase the change in pixel fluorescence); thus, we constrain β to be positive. Given that for any plane, there are many muscle traces (more than corrupted time points), we are at the risk of overfitting, and thus we run LASSO regression, setting $\alpha = 0.5$. Using $\hat{y}_t$ as the uncontaminated activity, we then have the following:

$$\Delta \hat{y}_t = \Delta y_t - \beta \Delta x_t \quad (10)$$

$$= (y_t - \beta x_t) - (y_{t-1} - \beta x_{t-1}) \quad (11)$$

As the model only learns to fit the differences, one needs to adjust the offset c :

$$\hat{y}_t = y_t - \beta(x_t) + c \quad (12)$$

We set c in order to minimize the difference between the raw timeseries (y and the decontaminated timeseries $\hat{y}$) at the time points of no contraction:

$$c = \frac{1}{N} \sum_{t \notin T_c} (y_t - \hat{y}_t) \quad (13)$$

where N is the number of time points not in the contaminated set. To avoid any potential overshoots (i.e.: removal of more signal than that caused by muscle contamination) we remove

any time point where the corrected timeseries has large negative derivatives, if, at the same time-point a large positive deflection was predicted: sudden decreases in Ca^{2+} activity are extremely rare; thus, it is safe to assume that if one occurs at the point of a large negative derivative, then it is likely because the model overestimates the amount of muscular contamination and in those instances we assume that we cannot be confident in the prediction results and instead ignore those data-points and linearly interpolate across those time points.

2.6 Additional Method Details

2.6.1 Experimental Model and Subject Details

Zebrafish Larvae were reared at 28.5 °C in 14-10h light-dark cycles (Westerfield, 2007). Zebrafish from 5 to 14 d.p.f. were fed rotifers and used for experiments. All experiments complied with protocols approved by the Institutional Care and Use Committee at Janelia Research Campus (JRC). Zebrafish sex cannot be determined until ~4 weeks post-fertilization (Anderson et al., 2012), so the sex of the experimental animals were unknown.

Transgenics Transgenic zebrafish larvae were in Casper or Nacre background (White et al., 2008). All lines were generated with Tol2 system (Kawakami, 2005). The following lines were generated or used in this work:

Tg(ubi:tTA);Tg(TRE:GCaMP7f) & Tg(ubi:tTA;TRE:jRGECO1b)

The original ubi:tTA (#713) was a gift from Daniel Feliciano and Isabel Espinosa-Medina. The TRE element was cloned from an existing plasmid containing the TRE:Cas9 element (#709) which was modified to replace Cas9 with GCaMP7f. Both plasmids were coinjected at equimolarity ~ 0.027 picomoles. For the ubi:tTA;TRE:jRGECO1b lines, all components were cloned into a single plasmid and injected.

Tg(foxj1a:mRuby) & Tg(foxj1a:GCaMP7f)

The foxj1a:GCaMP6s plasmid was ordered from addgene (catalogue number: 163829) and modified to expression mRuby (codon optimized) or GCaMP7f (codon optimized).

Tg(T2[b-actin2::mCherry-CAAX]; cmlc:GFP)

This line was a gift from Diane Speert.

Codon Optimization Sequences were codon optimized using CodonZ software (Horstick et al., 2015). We denote codon-optimized sequences with (zf).

Sample preparation for imaging For spinning-disk confocal microscopy, samples were embedded in 2% agarose in a glass bottom Petri dish on their side. When imaging in the mesoscope, samples were mounted in a standard Petri dish dorsal side up next to a 45° mirror.

3 Anatomical Foundations: Establishing an anatomical and molecular reference model for Whole Body Imaging via Expansion Microscopy

3.1 Introduction

Interpreting Whole Body Imaging (WBI) data relies heavily on having good anatomical knowledge of larval zebrafish anatomy. Indeed, WBI trades-off cell identifiability for tissue coverage: by labeling all cells one loses ground-truth knowledge of cell-type identity. This trade-off, which could have been very steep, in practice was found to be modest as a surprisingly large amount of tissue types can be distinguished by eye. This however requires a good grasp of the 3D organization of tissues within the animal, and the richer ones prior about tissue and cellular morphology, the more inferences one can make when analyzing WBI data. As such, analyzing functional data became inseparably intertwined with tissue type identification, the latter often becoming a limiting step, and calling for the gain of more granular anatomical knowledge of the fish.

Despite the larval zebrafish being an established model organism, comprehensive anatomical information for tissues other than the brain could be difficult to come by: existing high-resolution zebrafish atlases focus on the brain and central nervous system (e.g. MPI [Kunst et al., 2019](#), Harvard [Randlett et al., 2015](#), NIH [Tabor et al., 2019](#)) and body-wide datasets have limited resolution or have used X-ray (Penn State [Ding et al., 2019](#)) and do not include molecular information. The reason for the lack of published high-quality anatomical data, especially for inner organs, soon became clear as we attempted to collect our own anatomical reference data: despite the claims that larval zebrafish are transparent, the accumulation of changes in refractive index across the sample causes tissue scattering, which only worsens after fixation, making it difficult to collect clean data. Clearing protocols only made partial dents in this problem, and penetration of antibodies across the tissue was also an issue; indeed, many protocols ask users to remove the skin of the sample (e.g.: [Yang et al., 2018](#)) which comes at the risk of causing tissue disruption and warping. As such, this was a methodological problem. This chapter describes how we addressed this gap by leveraging advances in Expansion Microscopy and developing a protocol enabling the acquisition of high-resolution whole-body anatomical and molecular data, crucial to being able to interpret and make the most of WBI data.

Expansion microscopy (ExM) offers a ready solution to accessing deep tissue and probing for proteins and RNA. However, most existing ExM protocols have been optimized for brain tissue, a highly homogeneous tissue. The diversity of tissue types within a whole larva, compared to the brain alone, makes isotropic expansion more difficult to achieve. The challenges arise in part because the cross-link density between the sample and gel depends on the local tissue ultrastructure, and variations in this lead to inhomogeneous expansion and, worse, to tearing. To overcome this, strong proteolysis can be induced to break up protein networks (e.g.: using proteinase K), but at the cost of signal retention. Navigating the landscape of dependencies between recipe parameters is challenging, as the number of parameters is vast and interactions complex, and there is no high-throughput way to measure expansion uniformity and signal retention simultaneously. Furthermore, as the animal grows, tissue heterogeneity only amplifies with, for instance, the rapid expansion of cartilage, which begins around 6 days post-fertilization (d.p.f). Our initial attempts to apply existing protocols, optimized for larval brain tissue, to the entire fish led to large tears and deformations.

Two whole-fish ExM protocols have been published; the first uses the original low monomer ExM recipe and strong digestion with extensive proteinase K incubations, which comes at the cost of protein loss, and the article lacks scrutinizable IHC results (only one supplementary figure is present for a 3 d.p.f. fish, showing a superficial structure) (Sim et al., 2022). This group presents data showing good whole-body morphology; however, the oldest sample shown in detail is 6 d.p.f. (younger than that typically imaged during WBI). The protocol takes about 2 weeks. The second published whole-fish ExM protocol shows post-expansion IHC results (Steib et al., 2022). It overcomes the problem of expanding heterogeneous tissue by increasing the concentration of acrylamide (i.e., making a tougher gel); however, the protocol is only applied to samples of up to 5 d.p.f., which is the age just before when we begin to see large tears with the onset of cartilage growth. The protocol also takes about 2 weeks.

In need of a protocol compatible with both maximizing retention of pre-expansion IHC signals and using older ages of fish, mentored by Paul Tillberg and helped by Amy Hu and Mark Eddison, we developed a new whole-body ExM protocol, achieving both by adopting a new strategy: instead of attempting to make a single tough gel, we decreased the acrylamide concentration and used a multi-embedding approach combined with very high temperature (100 °C) disruption that enables uniform expansion without protease treatment, nor concomitant protein loss. The interlacing of polymers enables a gentler and gradual expansion. The resulting method takes a total of 1 week from fixation to imaging, reaches up to 4x expansion, and also works with older samples (21 d.p.f. zebrafish and mature *Danio rerio*). Finally, to facilitate data analysis, we also optimized sample mounting with the inclusion of an agarose pre-embedding step and enlarging the chamber to ensure consistent sample orientation and avoid edge artifacts. We provide a detailed user-friendly protocol to facilitate method uptake.

3.2 Background to Expansion Microscopy

Expansion microscopy (ExM) is a form of super-resolution microscopy. Unlike standard super-resolution methods, which either focus on altering the hardware or software, ExM focuses on altering the object to be imaged itself, making it physically larger (Chen et al., 2015). ExM considers native tissue as a physical network of biomolecules. Throughout this biological network, an artificial acrylamide-based polymer network can be interlaced. This artificial network can be made to molecularly bind to the biological network via the use of NHS (hydroxysuccinimide ester) chemistry: an NHS moiety is added to the polymer monomer (together called AcX for short). This NHS group can then undertake a nucleophilic attack to amine groups found within the tissue, for instance on lysines, a common amino acid present in essentially all proteins. To be able to move endogenous proteins apart, one needs to first reduce the number of crosslinks that keep the tissue together. Endogenous bonds can be disrupted by detergents, denatured, or cleaved. This reduces the number of forces that hold endogenous proteins together. Finally, by reducing the salt concentration of the medium, the artificial network can be made to swell. The swellable nature of the gel comes from the fact that the acrylate molecules are part of the acrylamide backbone. These charged components naturally want to repel each other; salt masks these charges, thus making the gel more compact at high salt levels. As the natural network of proteins is bound to the acrylamide gel, when the gel swells, the natural network of proteins expands with it. This results in molecules of interest being further apart and thus more readily resolvable. Lipids are typically dissolved away because they are not bound to the gel and the density of proteins left is much reduced, which together results in fewer refractive interfaces, lowering the spatial variation in refractive index, and leading to a cleared, optically accessible sample. The interaction between the properties of biological and artificial

networks and how they are assembled determines the quality of the outcome, including the expansion, homogeneity, and retention of molecules. Most of the work involves navigating this design space with now many variations of original protocol published having been optimized for specific samples and applications such as tissue thickness or expansion factor (Wang et al., 2021, M'Saad et al., 2022).

3.3 Engineering a Whole Body ExM protocol

Challenges to Whole Fish ExM The diversity of tissue types within the organism makes homogenous expansion difficult and there are strong trade-offs. For example, cross-link density varies according to the local tissue structure and tissue-gel anchoring. Such variations can lead to inhomogeneous expansion and at worst tearing. This can be normalized by strong proteolysis but at the cost of signal retention. Navigating the landscape of dependencies between parameters is challenging as the number of parameters is high and the interactions complex. Furthermore, as the animal develops tissue heterogeneity only amplifies with, for instance, the rapid development of cartilage beginning at 6 days post fertilization (d.p.f.).

Gel Engineering We found, as in Sim et al., 2022, that having a harsh digestion with proteinase K (proK) allowed for a homogeneous expansion. However, this turned into a dead end, as it would inevitably come with a severe loss of IHC signal applied before gel embedding and a loss of robustness in expansion uniformity. Therefore, this significantly limited the scope of the method. We then turned to high temperature disruption, by which heat is used to denature proteins and induce proteolysis, as reported in Steib et al., 2022. This type of digestion was better quality, probably because the temperature is more agnostic to protein or the local tissue environment, unlike an enzyme such as proteinase K. Despite this encouraging start, as we began working with older samples, propagating cracks appeared. The inflection point for this was quite striking, occurring and worsening from 6 d.p.f. as cartilaginous structures began to develop. After trying a variety of gel recipes, we found that an intermediate-monomer concentration gel recipe was least prone to crack formation, and that by re-embedding immediately after disruption, without reducing salt concentration or expanding more than 1.5x, subsequent crack formation was strongly reduced and any cracks that did form (typically in cartilage) did not propagate into neighboring tissues. The working model would be that an intermediate monomer concentration may produce a better match to tissue elasticity, which progressively becomes reinforced: polymer chains form within the gel and gradually expand the sample; this modest expansion then enables greater access to the next round of polymers, thereby reinforcing the gel progressively and making the polymer density more homogeneous throughout, rather than forming one very dense gel initially and hoping that the gel strength is sufficient to fully overcome tissue elasticity. For brevity, the protocol is described for two embedding rounds, but more rounds can be added when larger expansion factors are desired. With two rounds of embedding the gel reaches about 2x expansion factor in PBS. For more context, the basics of ExM chemistry is described in [Appendix H](#).

Immuno-histochemistry Antibody staining was initially a problem. Signal amplification through IHC is necessary, as there is endogenous protein loss as well as signal dilution from physical expansion itself. We initially experienced that the skin and outer layer of the fish alone would be stained (the "shell" of the fish), leaving an unstained core. This can occur when there is local depletion of the probes. We found that strongly permeablizing the fish was critical, as

well as increasing antibody (AB) concentration, incubation time, and temperature. Increasing antibody concentration comes at the risk of higher background signal; however, a titration curve of antibody concentration vs. signal showed that the gain in signal far outweighs small increases in background and that at 1:100 we are still not reaching epitope saturation. Since secondary antibodies are often of goat origin, we also found that using goat serum in blocking buffer was preferable. With long incubations in rich serum, bacterial growth is more likely to occur leading to nonspecific background staining; this can be resolved by processing the samples with minimal delay and adding sodium azide from the point of fixation onward.

Orientation and edge effects Since the head of the sample is generally wider than its body, placing samples on their sides often results in slanted samples. Despite this not being a problem for local measurements, this can a) mean that a much larger field of view had to be acquired and b) render alignment of the sample to a reference a highly laborious work requiring hours of manual work. We found that fixed samples can be embedded in low melting temperature agarose (0.7%) where it is easy to optimally orient the sample, stabilizing it in place. Following high-temperature denaturation, agarose pre-embedding did not appear to hinder the ExM expansion process. This process results in more standardized samples. Another related issue was the changes in refractive index that can occur at the edges of the gel as a result of boundary conditions in the gelation process. We could resolve this by increasing the chamber height and ensuring that the first gel was sufficiently large to fill the full extent of the volume that would be the final gel.

3.4 Whole Fish ExM IHC and FISH Protocols

The detailed protocols for whole-fish ExM-IHC and whole-fish ExM-FISH can be found in the appendix ([Appendix H.1](#), [Appendix H.2](#)).

3.5 Results and Applications

We use this protocol throughout [Section 5](#) to help interpret and extend functional findings. Beyond these focused applications, the general protocol opens up multiple avenues. The method, if scaled up, would enable one to build the first whole-body cellular resolution atlas of the larval zebrafish. Compared to electron microscopy (EM), this could be produced at relatively low-cost and at high-throughput (enabling sample averaging and estimation of distributions of statistics) and come with the advantage of including ground truth labels from IHC and FISH. An ExM atlas would thus be complementary to an electron microscopy reference, which still remains the gold standard for subcellular anatomy. Through the use of Whole-Body ExM, such an atlas could include:

The major organ systems and glands. Although these are fundamental backbones of anatomy and have been documented for many years, basic details such as the location of apposition between such major systems are sparsely described: for instance, the contact points between the ventricular system and the vascular system, critical sites of exchange, are difficult to access without ExM or EM. In [Fig. 3](#), we expand a line labeling both the ventricular and vascular system which revealed multiple extensive contact areas (e.g.: [Fig. 3G,H](#)), paralleling those described in the mammalian literature ([Langlet et al., 2013](#)).




Figure 3: Mapping major organ systems and cell types of the body in larval zebrafish using whole body ExM-IHC.

(A) Whole-body ExM of the ventricular system. ExM sample labeling the ependymal cell network throughout brain and spinal cord. In green, total protein stain, NHS-AF488; in magenta, IHC against eGFP expressed in motile ciliated cells by the transgenic line Tg (foxj1a: eGFP) - 7d.p.f. sample.

(B) Inset of A: Midline ventricular ependymal cells that extend dense projections to the base of the brain, structure which has not been previously described in larval zebrafish. cb: cell body, p: projections.

(C) Inset of A: Individual cilia of ependymal cell projecting into the central canal. c: cilia.

(D) Inset of A: Example of ciliated cell outside the ventricular system: multiciliated nephric cell located in the proximal duct of the kidney.

(E) Whole body ExM of both vascular system and the ependymal cell network. In green, IHC against RFP expressed in vascular endothelium by transgenic line Tg(flk1:dsRED); in magenta, IHC against eGFP expressed in motile ciliated cells by transgenic line Tg(foxj1a:eGFP). Dorsal view of 7 d.p.f. sample.

(F) Crop of main cerebral arteries with individual midline ramifications visible. Same sample as in E.

(G) Ependymal cell projections (magenta) that extend to the basilar artery at the base of the brain. Same samples as in E.

(H) Ependymal cells wrapped around a cerebral artery. Same sample as in E.

Visceral neural innervation. The brain and body communicate directly through a dense set of bidirectional neuronal tracts called the autonomic nervous system (ANS), which innervate specific cell types. Despite increasing awareness of the role of the ANS, the visceral cell types which the brain has direct control over, or directly listens to, are still an active area of research and are being addressed piecewise for lack of optical and molecular access. Whole-body ExM allows the labeling of specific branches of the ANS and following their projections to their targets. As an illustration, in Fig. 4, we label the motor and sensory vagus in an ExM sample. For example, Fig. 4G,H show motor vagal fibers (gold) wrapping around postganglionic neurons in the heart (intrinsic cardiac nervous system - first demonstration to our knowledge in fish) and enteric neurons, respectively. In Fig. 4C one can trace out sensory fibers innervating a taste-bud; in contrast to textbook knowledge (Roper, 2013), we find that not only is the base of the tastebud heavily innervated, but projections extend right to the tip of the gustatory cell facing the luminal surface, suggesting that neurons may be directly sensing gastrointestinal luminal content, as well as more traditional indirect means of sensing. Such luminal projections were also observed in the anterior gastro-intestinal tract. This highlights the potential of whole-fish IHC-ExM to uncover yet unknown body-brain communication nodes.




Figure 4: Tracing the connectivity within the autonomic nervous system and neural-visceral innervation using whole-body IHC ExM in larval zebrafish.

(A) Whole body expansion microscopy with total protein stain and dual color immuno-labeling. In orange, total protein stain, NHS-AF488; in green IHC of eGFP expressed in cells of the autonomic nervous system by transgenic line Tg(phox2b:eGFP); in blue, IHC of RFP expressed in cranial motor nerves by transgenic line Tg(isl1CREST1:RFP) - 7 d.p.f. sample.

(B) Reference Atlas. Example of preliminary segmentation and annotation of whole fish, to which further samples would be registered.

(C-E) Innervation patterns of visceral tissues. In green, total protein stain; in magenta, IHC of eGFP (the same as A).

(C) Sensory innervation of the pharyngeal taste bud displaying cilia (c: primary cilia, si: sensory innervation).

(D) Pancreatic endocrine islet showing dense vascularization, containing neural crest derived phox2b+ cells (v: vascularization). (E) Cross-section of gastrointestinal tract lined by mucus-secreting goblet cells, surrounded by enteric neurons with visible bacteria non-specifically stained by anti-GFP antibody in lumen. Sarcomeres in adjacent myotome are also visible. (g: goblet cells, e: enteric neurons, b: bacterium, s: sarcomere)

(F) Major motor vagal branches extending to gills, heart, and anterior intestine. In magenta, IHC of eGFP (same as A), in gold, IHC of RFP (same as A).

(G) Motor vagus innervation of intrinsic cardiac neuron.

(H) Motor vagus innervation of enteric motor neuron.

Global histology. Decades of work and progress have been made through the use of histological stains (e.g.: hematoxylin and eosin stain). Although these stains are not specific, they produce stereotyped staining patterns from their differential tissue binding properties. These patterns are recognizable by a trained eye (or algorithm) and allow for a range of investigations by themselves alone, as well as providing context for more specific stains. Similarly, in ExM, total protein stains provide a similar benefit: staining with different NHS dyes with varying affinities: lipophilic, hydrophilic, maleimides (affinity for sulfur groups), potentially glycan dyes, a rich array of information becomes visible (Fig. 5). This offers a cheap and easy means to navigate through the tissue and assess the histology of the tissue of interest. By its generality, this can be transformative and should not be understated just because of the simplicity of the approach. In Fig. 4A, the general outline of the fish and its main organs are visible through staining with a total protein stain (orange). Finer tissue differentiation can also be made, e.g.: Fig. 4E, shows a cross section of the gut, where muscle tissue sarcomeres are visible (s), or gastrointestinal goblet cells (g) that secrete mucous into the lumen. Subcellular features can also be identified; for example, in Fig. 5B, samples were stained with NHS (NHS-Alex488) and maleimide dye (maleimide-ATTO647N) - one can see the actin network within epithelial cells surrounding the nucleus, typically studied via electron microscopy or using actin-specific labelling (Lam et al., 2015). In Fig. 5C, dense microvilli of the gastrointestinal tract are observed on the apical border of the epithelium forming a canal, while in Fig. 5G, a section of the anterior intestine is shown where dense vesicles are observed that line the border of cells that were later confirmed to be mitochondria, highlighting the aerobic requirements of this tissue.




Figure 5: ExM combined with histological stains.

(A-H) Whole body ExM of wild type 9 d.p.f. larval zebrafish with dual color total protein stain dyes highlighting differential affinity in such dyes. In green, hydrophilic amine binding Alexa 488-NHS; in magenta, lipophilic thiol binding ATTO647N maleimide.

(A) Individual collagen fibers (green lines) passing in front muscle sarcomeres (magenta).

(B) Polygonal skin epithelial cells that line the animal. Nuclei are visible and are surrounded by intracellular actin networks similar to that described in EM studies or studies specifically labelling actin (Lam et al., 2015).

(C) Cross-section of the gut with goblet cells extending into the lumen. Dense microvilli separate walls of the intestinal tract.

(D) Caudal fin with collagen fibers radiating out showing larger scale organization of fibers seen in A.

(E) Vacuolated cells of notochord separated by bridging membranes. Unlike skin epithelial cells whose cellular membranes are thick and continuous, in notochord cells, intercellular attachment points appear discrete and appearing at regular intervals.

(F) Cross-section of the exocrine pancreas. Subcellular organelles are visible in magenta - the identity could not be determined, similarly other neighboring tissue-level structures could not be identified despite extensive literature search. This highlights the fact that some observations still need follow-up studies to be identified as well as the potential for novel observations.

(G) Foregut epithelium with dense staining at apical end of cells identified separately as mitochondria.

(H) Cross section of the liver showing sparse but frequently observed cell types extending thin processes within the tissue (identity yet to be determined).

Hormonal secretion axis and receptors. The ability to probe for RNA expression across the body opens up a deeper and complimentary characterization of cell identity and state as compared to immunohistochemistry (IHC). While IHC allows for tracing morphology and the assessment of the presence of proteins which are often the ultimate effector molecules of interest, good antibodies are incredibly difficult to come by and very expensive. Part of this comes from

the fact that zebrafish-specific antibodies are a niche product, and another equally significant factor is that antibodies are typically tested and validated in the context of pull-down assays or in cell cultures, which do not necessarily predict their affinities in a tissue or ExM setting. Each epitope requires its unique antibody, which have batch-to-batch variability. In contrast to these factors which frustrate IHC studies, FISH probes can be readily designed and ordered, and multiple approaches exist to multiplex and amplify them. In the context of whole-body imaging, ExM-FISH provides an easy way to either identify specific subpopulations via known genes or use it as a discovery tool to find, for instance, which cells have the capacity to sense a specific hormone. In [Fig. 6](#) we applied whole-body ExM-FISH to locate components of the hormonal hypothalamus-pituitary-adrenal (HPA) axis. In [Fig. 6B](#), we stain for the gene responsible for encoding the adrenocorticotrophic hormone (ACTH) that is expressed in the anterior pituitary, the deepest part of the brain that is difficult to otherwise access. In [Fig. 6C](#), we probe for the gene involved in steroid production to identify the steroidogenic adrenal glands. The results came as a surprise, as their distinctive location bordering the central veins differed from the vague description provided in the literature that describes them as "within the neck of the kidney", "is embedded at the anterior portion of the head kidney" ([Yi-Wen, 2007](#)). The spatial footprint of these cells also mapped (by eye) to a cell cluster identified in functional data with distinctive functional biphasic activity, which calls for functional to ExM data registration.




Figure 6: Identifying endocrine cells throughout the body using spatial transcriptomics via whole-body FISH-ExM in larval zebrafish.

(A) Whole-body FISH ExM staining against ribosomal RNA showing whole body access of probes - 7 d.p.f. sample. Protocol described in [Appendix H.1](#).

(B) Locating pituitary ACTH producing cells. In blue, DAPI; In green, NHS-488; in magenta, HCR against POMCA. The location of this gland at the very base of the brain makes it normally particularly difficult to image; individual cells can easily be discerned in 2D plane (B1) or the max-projection of the 3D rendering of gland (B2).

(C) Locating steroidogenic cells. In yellow, NHS-488; in red, HCR against *hsd3b1*; in blue, DAPI. Commonly reported as being within the neck of the kidney, ExM-HCR shows that steroidogenic are located well outside the nephric system, lining the wall of the passing vasculature.

(D) Identification of chondrocytes. In green, NHS-488; in magenta, HCR against col2a1a. Despite the large changes in refractive index of cartilaginous structure, chondrocytes as well as surrounding tissue can be readily imaged.

(E) Results of single cell RNA sequencing of sensory vagus and glossopharyngeal ganglion. Left: UMAP projection of data matrix color coded by Leiden cluster ID. Right: Dot plot of top subpopulation markers.

(F) Example of validation and localization of one subpopulation marker identified from the scRNASeq dataset, *calb2a*. In green, total protein stain; in magenta, *calb2a* HCR. Whole-body staining also revealed an undescribed *calb2a*+ population in the pharyngeal taste buds.

Other developmental time-points and species. The embryo undergoes drastic changes during the first days of its existence. While previously published protocols enabled the study up to about 6 d.p.f. (Sim et al., 2022, Steib et al., 2022), the compatibility of our new protocol with older ages opens up the opportunity to study later stages of development. Neural innervation of visceral tissue, for example, is still developing at 7 d.p.f. (Yang et al., 2018). There is also increasing interest in other species, such as *Danionella cerebrum*, a paedomorphic animal that maintains its small size and relative transparency into adulthood (Britz et al., 2021). We apply whole-body ExM to mature (3 months old) *Danionella cerebrum*. The first attempts to apply the protocol were surprisingly encouraging despite the larger size and more advanced ossification of *Danionella cerebrum* Fig. 7 - more experiments are needed to optimize the mounting, staining, and imaging and assure a robust protocol.




Figure 7: **Whole Body ExM for larger samples.**

(A) Whole-body expansion microscopy of mature female *Danionella cerebrum*. Image showing a 3 weeks old, 2x expanded sample, stained with total protein stain (Alexa 488 NHS) acquired in the SmartSPIM. The larger size of the sample requires a different imaging system to be used as compared to expanded larval zebrafish. (B) Gills. (C) Cardiac ventricles. (D) Ovary. Unlike larval zebrafish whose reproductive organs haven't yet developed, *Danionella* can be functional imaged and expanded at an age at which reproductive functions are mature.

3.6 Next steps

Much work remains. The next immediate key points in terms of protocol developments will include a) being able to register functional data to ExM data, b) running multi-round IHC-ExM by staining post-expansion, c) engineering the protocol to be compatible with both IHC-ExM and FISH-ExM, d) testing the limit of expansion degree to be able to access a larger array of intracellular structures, and e) including multi-round HCR-FISH. In terms of data processing, a workflow to register samples together as well as a database needs to be established.

4 Analysis Tools: Developing computational methods for analyzing Whole Body Dynamics

4.1 Introduction

We are faced with the challenge of analyzing high-dimensional time-series data coming from a dynamical system, in our case the body. Two common goals of researchers analyzing high-dimensional time-series data are: a) trying to summarize high-dimensional activity into lower dimensional latent processes, a.k.a.: dimensionality reduction, and b) grouping time-series in meaningful ways, a.k.a.: clustering. Here, we develop methods for both dimensionality reduction and clustering which address existing gaps which surfaced when analyzing Whole Body Imaging (WBI) data. A first feature of the data is that it captures activity traces coming from a range of cell-types whose numerosity and propensity to exhibit transients vary widely. As a result, such methods which focus on e.g.: accounting for variance in the data, often end up being dominated by a few cell-types, and capture only small fraction of the range of dynamics observed. A second feature of the data is that exhibits a variety of waveforms, timescales and propagating population sequences for which we don't have strong priors, resulting in poor performance of methods which for instance either disregard temporal correlations or have pre-defined range of time-lags. These features are of course not unique to WBI data and are prevalent in biological datasets; however, their particularly prevalence in WBI data compelled us to focus on these issues directly.

In [Section 4.2](#), we note that activity often propagates through a population of cells over time, rather than the whole population being synchronously active. The latter would arise if one driving force would reach all cell at the same time which in most tissue seems not to be the case. As a result of this, coupled cells would exhibit time-lagged activity and would be coherent, rather than then being synchronous and correlated. We thus adapt two well established methods k-means clustering and Spectral clustering to use coherence as a similarity metric, thereby better reflecting the nature of the data.

In [Section 4.3](#) and [Section 4.4](#), we address the second challenge: finding a low-dimensional representation of high-dimensional data. In the ideal case, we would want that the true phase portrait of the system be faithfully reflected in the lower-dimensional representation that we learn. With this in mind, we asked ourselves what statistical property is common to all dynamical systems and that we might want to preserve in our low dimensional representation. One such property stood out: most natural processes are temporally non-reversible. Indeed, time-reverse the majority of film-clips or song and there is little doubt in the audience as to whether it was reversed or not. Indeed, a key property of dynamical systems is that they define a consistent flow field in state space, such that any segment of trajectory produced by the system is unlikely to be produced in the opposite direction — in other words, the trajectories are non-reversible. Noise in contrast is reversible. Despite this fact, few methods use non-time reversibility to inform data compression or as means to distinguish signal from noise. We thus set out to find a way of a) quantifying non-reversibility, b) outline a means to extract non-reversible projections in high-d data via a deterministic scale-able algorithm [Section 4.3](#), c) generalize the concept to the probabilistic setting by expressing non-reversibility as a prior in GP based models [Section 4.4](#). As a first step we illustrate GPFADS on toy data and a well-known and studied publicly available dataset to better highlight the benefits of GPFADS and have some form of "ground truth" of the dynamics in the data. Later work will be to apply GPFADS directly on WBI data.

The work in [Section 4.3](#) on "SCA, Sequential Component Analysis" was presented in a talk at Cosyne 2020. The work outlined in [Section 4.4](#) on "GPFADS, Gaussian Process Factor Analysis with Dynamical Structure" was [published](#) and presented in a talk at NeurIPS 2020.

4.2 Coherence based clustering

4.2.1 Coherence based K-means clustering

4.2.1.1 Introduction

When identifying functional clusters in data, we would want an unsupervised way of clustering pixels whose activity is interdependent. One such proxy of dependence is correlation. However, if one system drives another with a certain spectral content, this is more likely to lead to the driven system showing similar spectral content with some gain and phase modulation, rather than being instantaneously correlated. The measure of coherence captures this idea, yet methods do not often exploit this property to cluster data. Thus, we derive the equivalent of k-means clustering using coherence as a metric of distance. The solution to calculating cluster centroids can be shown to be equivalent to solving an eigenvalue problem. The centroid of the cluster then gives a natural reference phase for calculations and to which to align data to. As one can define the frequency range to consider, this also provides a way of being immune to high-frequency noise. We apply this to real and toy data, where we show that coherence-based k-means gives more meaningful results than standard k-means.

Notation and Definitions We define the following: M : # of latent clusters, K : # of sub-samples used to average over to calculate coherence, T : # of total time points, L : # of time points per sub-sample, $\mathbf{x}_i \in \mathbb{R}^T$: sample i , $\mathbf{x}_{ik} \in \mathbb{R}^L$: sub-sample k of sample i of length L , $\boldsymbol{\mu}_m \in \mathbb{R}^L$: centroid of cluster m , $\mathcal{C}(x, y)(\omega)$: coherence between x and y at frequency ω .

Coherence Coherence is a property of a joint distribution over signals, not of specific signals. Thus, we interpret $\mathcal{C}(x, \boldsymbol{\mu})(\omega)$ as the windowed magnitude-squared coherence estimator for stationary signals (i.e., Welch’s method). That is, we first form K -windowed or tapered sub-sequences from each $\mathbf{x}$ and $\boldsymbol{\mu}$ and find their Fourier transforms $\hat{x}_k(\omega)$ and $\hat{\mu}_k$.

$$\mathcal{C}(\mathbf{x}, \boldsymbol{\mu})(\omega) = \frac{|S_{x, \boldsymbol{\mu}}(\omega)|^2}{S_{xx}(\omega)S_{\boldsymbol{\mu}, \boldsymbol{\mu}}(\omega)} \quad (14)$$

$$= \frac{|\mathbb{E} [\hat{x}(\omega) \cdot \overline{\hat{\mu}(\omega)}]|^2}{\mathbb{E} [\hat{\mu}(\omega) \cdot \overline{\hat{\mu}(\omega)}] \mathbb{E} [\hat{x}(\omega) \cdot \overline{\hat{x}(\omega)}]} \quad (15)$$

$$= \frac{|\frac{1}{K} \sum_{k=1}^K \hat{x}_k(\omega) \cdot \overline{\hat{\mu}_k(\omega)}|^2}{\left(\frac{1}{K} \sum_{k=1}^K \hat{\mu}_k(\omega) \cdot \overline{\hat{\mu}_k(\omega)}\right) \left(\frac{1}{K} \sum_{k=1}^K \hat{x}_k(\omega) \cdot \overline{\hat{x}_k(\omega)}\right)} \quad (16)$$

K-means Algorithm To initialize the group labels, we first run k-means ++ ([Arthur and Vassilvitskii, 2007](#)), an algorithm for choosing the initial values. The cluster centroids are computed as described in the next section. The coherence between each data point and centroids are computed. Cluster allocation reassigns the cluster ID based on coherence. Cluster centroids are recalculated and the process iterates until the convergence of the labels. The number of

clusters can be chosen based on the average silhouette score: $s(i) = \frac{b(i) - a(i)}{\max\{a(i), b(i)\}}$, where $b(i)$ is the distance between a sample i and the closest data point not in its cluster, and $a(i)$ is the mean intracluster distance. The best achievable value would be 1, and the worst: -1.

4.2.1.2 Derivation

We wish to find a centroid μ_m such that the sum of the coherences of points within that cluster $\sum_i \mathcal{C}(x_i, \mu_m)$ is maximal.

This quantity is invariant to an arbitrary real scaling of $\hat{\mu}_{mk}$, so we add the constraint that $S_{\mu\mu}(\omega) = 1$, i.e.: $\mathbb{E}[|\hat{\mu}_m(\omega)|^2] = 1$ that is $\frac{1}{K} \sum_{k=i}^K |\hat{\mu}_{mk}(\omega)|^2 = 1$.

Thus, we aim to maximize:

$$\mathcal{L}(\hat{\mu}_m(\omega)) = \sum_{i=1}^N \mathcal{C}(\hat{x}_i(\omega), \hat{\mu}_m(\omega)) \quad (17)$$

$$= \sum_{i=1}^N \lambda_i \left| \sum_{k=i}^K \overline{\hat{x}_{ik}(\omega)} \cdot \hat{\mu}_{mk}(\omega) \right|^2 \quad (18)$$

$$= \sum_{i=1}^N \lambda_i \left| \sum_{k=i}^K \sum_{k'=i}^K \overline{\hat{x}_{ik}(\omega)} \cdot \hat{\mu}_{mk}(\omega) \cdot \hat{x}_{ik'}(\omega) \cdot \overline{\hat{\mu}_{mk'}(\omega)} \right| \quad (19)$$

$$= \left| \sum_{k=i}^K \sum_{k'=i}^K \hat{\mu}_{mk}(\omega) \left(\sum_{i=1}^N \lambda_i \overline{\hat{x}_{ik}(\omega)} \cdot \hat{x}_{ik'}(\omega) \right) \overline{\hat{\mu}_{mk'}(\omega)} \right| \quad (20)$$

$$= \left| \sum_{k=i}^K \sum_{k'=i}^K \hat{\mu}_{mk}(\omega) A_{k,k'} \overline{\hat{\mu}_{mk'}(\omega)} \right| \quad (21)$$

writing $\lambda_i = \frac{1}{K^2 S_{x_i x_i}(\omega) S_{\mu\mu}(\omega)} = \frac{1}{K^2 S_{x_i x_i}(\omega)}$

$$A_{k,k'}(\omega) = \sum_{i \in \mathcal{S}_m} \lambda_i [\hat{x}_{ik}(\omega) \cdot \overline{\hat{x}_{ik'}(\omega)}] \quad (22)$$

We note that $A_{k,k'}(\omega) = A_{k',k}^*(\omega)$ so the elements form a Hermitian matrix $A(\omega)$.

Vectorizing μ_m , we can rewrite the objective as:

$$\text{argmax } \hat{\mu}_m(\omega)^\top A(\omega) \overline{\hat{\mu}_m(\omega)} \text{ subject to } \hat{\mu}_m(\omega)^\top \overline{\hat{\mu}_m(\omega)} = 1 \quad (23)$$

from which it follows that $\hat{\mu}_m(\omega)$ is the leading eigenvector of $A(\omega)$.

It also follows that the eigenvalue is equal to the sum of the coherences of the centroid with all the included data points. An alternative derivation can be achieved by splitting the problem into real and imaginary parts.

Complexity Computing $A(\omega)$ takes $O(NK^2\Omega)$ sums and as many multiplications. Calculating the highest eigenvector $A(\omega)$ is of $O(\Omega K^3)$ complexity.

Application

We apply coherence-based k-means clustering in [Section 5.5](#) on Ca^{2+} activity from ependymal cells. For purpose of demonstration, we also run coherence-based k-means clustering on bright-field cardiac data. In short, a video of the heart was acquired (15Hz, ventral view), the frames were vectorized, and coherence k-means run. In [Fig. 8A](#), we select the centroid corresponding to the atria of the heart, and show the phase of the coherence (before taking the absolute magnitude of it) between each pixel relative to this centroid, alpha weighted by the pixel's coherence with this centroid. In this image, the shape of the heart is visible as pixels outside of the cardiac chamber are not coherent; furthermore, a traveling wave appears to propagate from the venous sinuses (left) to the bulbus arteriosus (right). In [Fig. 8D](#), the magnitude of the coherence is plotted. We note that here an arbitrary number of clusters was chosen. When applying it for an actual use case, we compute the silhouette score to decide how many clusters to request. In [Fig. 8B](#), we plot the variation in the phase shift relative to the said centroid (inverse of the vector strength R circular statistic). We see that there is more variability in the atria as compared to the ventricle; as the initiation of contraction occurs in the atria, where there are multiple competing initiation nodes, this could be expected. In [Fig. 8C](#), the variance in magnitude is shown; in contrast to phase variability, the magnitude of variance in contraction strength appears to be constant throughout the heart. In [Fig. 8E](#), the clusters themselves are shown.




Figure 8: **Coherence K-means clustering applied to cardiac brightfield data.**

(A) Image shows the phase of the coherence between each pixel and the cluster centroid corresponding to the atria. A traveling wave is observed starting from the entry point of the heart (left) to the bulbus arteriosus (right).

(B) Image shows the variation over time of the phase shift across trials of each pixel relative to the cluster centroid corresponding to the atria.

(C) Image shows the magnitude of the coherence between each pixel and the cluster centroid corresponding to the atria.

(D) Image shows the variance in the magnitude of coherence over time between each pixel and the cluster centroid corresponding to the atria.

(E) The image shows the anatomical footprint of each cluster. Anatomical compartments can be identified: venous sinus (pink), atria (gray), ventricle (green), bulbus arteriosus (right most grey region, red).

4.2.2 Coherence based Spectral Clustering

4.2.2.1 Introduction

Spectral clustering is a well-established method for clustering data ([Von Luxburg, 2007](#), [Niezik and Van Der Schaft, 2011](#)) that requires one to define a similarity function between data points, which is often set to the Euclidean distance. As discussed in the previous section, such a metric

is not necessarily a meaningful distance to use between two time series. Correlation would be a better distance measure, but fails to encode that time-lagged activity traces should ideally be considered as close, which results in the breaking up of such groups of traces into multiple clusters. We thus aim to define a distance based on coherence (small when coherence is high and vice versa) which can be used in spectral clustering.

Background In spectral clustering, one first defines an adjacency graph. A common choice of similarity function is the Gaussian similarity function defined between x_i and x_j as:

$$W_{i,j} = \exp\left(-\frac{|x_i - x_j|_2^2}{\tau}\right) \quad (24)$$

where τ defines the range of distances of interest, or width of the neighborhood.

One then builds the normalized graph Laplacian of the adjacency matrix:

$$L = D - W \quad (25)$$

$$L_{\text{norm}} = I - D^{-1}W \quad (26)$$

where D is a diagonal matrix and $d_{i,i} = \sum_j W_{i,j}$. One then computes the lowest N eigenvectors of the matrix and uses these as the encoding basis set, and finally runs the k-means algorithm searching for N clusters.

For intuition, we can interpret L as being a transition function of the consensus dynamics of the system. Assuming rows of W already sum to 1 for simplification:

$$\frac{dx}{dt} = -Lx \quad (27)$$

$$\frac{dx}{dt} = -(I - W)x = -x + \sum_j W_{i,j}x \quad (28)$$

where the fixed point is:

$$x = \sum_j W_{i,j}x \quad (29)$$

which is the weighted average of x (also showing that it has to be stable). The eigenvalues thus tell us how quickly consensus, or information, will be distributed. It can be shown that the graph has as many 0 eigenvalues as disconnected components; indeed, consensus would never occur in disconnected graph. The smaller the eigenvalues the slower the diffusion within that mode, and the similarity of the weights of the eigenvectors indicate which entries are similar. Indeed, for any vector f :

$$f'Lf = f'Df - f'Wf \quad (30)$$

$$= \sum_{i=1}^n d_i f_i^2 - \sum_{i,j=1}^n f_i f_j w_{ij} \quad (31)$$

$$= \frac{1}{2} \left(\sum_{i=1}^n d_i f_i^2 - 2 \sum_{i,j=1}^n f_i f_j w_{ij} + \sum_{j=1}^n d_j f_j^2 \right) \quad (32)$$

$$= \frac{1}{2} \sum_{i,j=1}^n w_{ij} (f_i - f_j)^2. \quad (33)$$

which is minimized if for large similarity w_{ij} , f_i and f_j are close.

4.2.2.2 Derivation

We first can show that if one wanted to use correlation as similarity metric, we can define a distance which depends on correlation (small when correlation is high and vice-versa) and which can be rewritten as the Euclidean distance between vectors. This then ensures that we can use this defined distance in a Gaussian kernel and that it remains p.s.d..

The correlation between vectors $\mathbf{x} \in R^N$ and $\mathbf{y} \in R^N$ being defined as:

$$\mathcal{R}(\mathbf{x}, \mathbf{y}) = \frac{\sum_i x_i \cdot y_i}{\sqrt{\sum_i x_i^2} \sqrt{\sum_i y_i^2}} \quad (34)$$

$$(35)$$

We assume already pre-normalized vector:

$$x_i = \frac{x_i}{\sqrt{\sum_i x_i^2}} \quad (36)$$

Then:

$$|\mathbf{x} - \mathbf{y}|_2^2 = 2 \sum_i (1 - x_i \cdot y_i) \quad (37)$$

$$= 2(1 - \sum_i (x_i \cdot y_i)) \quad (38)$$

$$= 2(1 - \mathcal{R}(\mathbf{x}, \mathbf{y})) \quad (39)$$

$$(40)$$

Thus, the euclidean distance between normalized vector defines a distance based on correlation.

In the same spirit, one can define a distance based on coherence and show that it can be written as a euclidean distance.

By definition, coherence is defined as:

$$\mathcal{C}(\mathbf{x}, \mathbf{y})(\omega) = \frac{|\sum_k \hat{x}_k(\omega) \overline{\hat{y}_k(\omega)}|^2}{\left(\sum_k \hat{x}_k(\omega) \overline{\hat{x}_k(\omega)}\right) \left(\sum_k \hat{y}_k(\omega) \overline{\hat{y}_k(\omega)}\right)} \quad (41)$$

$$(42)$$

We assume an already normalized spectrum to not carry the normalization term:

$$\sum_k \hat{x}_k(\omega) \overline{\hat{x}_k(\omega)} = 1 \quad (43)$$

Coherence can then be re-written as:

$$\mathcal{C}(\mathbf{x}, \mathbf{y})(\omega) = \left(\sum_k \hat{x}_k(\omega) \overline{\hat{y}_k(\omega)} \right) \left(\sum_l \hat{y}_l(\omega) \overline{\hat{x}_l(\omega)} \right) \quad (44)$$

$$= \left(\sum_k \sum_l \hat{x}_k(\omega) \overline{\hat{y}_k(\omega)} \hat{y}_l(\omega) \overline{\hat{x}_l(\omega)} \right) \quad (45)$$

$$= \text{vec} \left(\hat{X}(\omega) \overline{\hat{X}(\omega)}^\top \right)^\top \text{vec} \left(\hat{Y}(\omega) \overline{\hat{Y}(\omega)}^\top \right) \quad (46)$$

where $\hat{X}(\omega) \in R^k$ is a vector of all of the $\hat{x}_k(\omega)$. Thus, we define $Z(\omega) = \text{vec} \left(\hat{X}(\omega) \overline{\hat{X}(\omega)}^\top \right)$

and $W(\omega) = \text{vec} \left(\hat{Y}(\omega) \overline{\hat{Y}(\omega)}^\top \right)$

Thereby, we can define a distance based on coherence which can be written as a Euclidean distance between two vectors, thus guaranteeing that one can use this distance and have a Gram matrix that remains psd.

$$|Z(\omega) - W(\omega)|_2^2 = \sum_k \sum_l (\hat{x}_k(\omega) \overline{\hat{x}_l(\omega)})^2 + \sum_k \sum_l (\hat{y}_k(\omega) \overline{\hat{y}_l(\omega)})^2 - 2\mathcal{C}(\mathbf{x}, \mathbf{y}) \quad (47)$$

$$= 2(1 - \mathcal{C}(\mathbf{x}, \mathbf{y})) \quad (48)$$

Normalization

Typically coherence is averaged over frequencies:

$$\mathcal{C}(\mathbf{x}, \mathbf{y}) = \sum_{\omega} \mathcal{C}(\mathbf{x}, \mathbf{y})(\omega) \quad (49)$$

$$= \sum_{\omega} \frac{|S_{x,y}(\omega)|^2}{S_{xx}(\omega) S_{yy}(\omega)} \quad (50)$$

This definition weights each $\mathcal{C}(\omega)$ at each frequency irrespective of the power in that frequency band, this can be detrimental as bands of low power can be dominated by noise. To avoid this, one can normalize by the total power in all frequency bands:

$$\mathcal{C}(\mathbf{x}, \mathbf{y}) = \frac{\sum_{\omega} |S_{x,y}(\omega)|^2}{\sum_{\omega} S_{xx}(\omega) S_{yy}(\omega)} \quad (51)$$

which remains bounded between 0 and 1. In other words, one can normalize the entire vector $\hat{X} \in R^{K\Omega}$:


$$\hat{x}_k(\omega) \rightarrow \frac{\hat{x}_k(\omega)}{\sqrt{\sum_{\omega} \sum_k \hat{x}_k(\omega) \overline{\hat{x}_k(\omega)}}} \quad (52)$$

Application

We discuss in more details the results of coherence-based Spectral Clustering in [Section 5.3](#), [Fig. 16](#). As a motivating example as to the use of coherence vs. correlation as a metric, we run the algorithm with each metric on a single plane keeping the other parameters constant (τ , N :number of eigenmodes/clusters). We find that coherence based spectral clustering groups together more compactly tissues: e.g.: the proximal part of the nephric tube is grouped into one cluster despite their being significant time delay between dorsal and ventral portions as well as them exhibiting slightly differing waveforms ([Fig. 9A](#)). In contrast, correlation-based spectral clustering systematically separates them at the cost of finding other clusters in the data. This, in part, contributed to other clusters being more mixed ([Fig. 9B](#)). We note that these result depend on the number of clusters asked, and when larger numbers of clusters are searched for, both methods break up the nephric tube into disjoint sets, as indeed cells within a segment of kidney are more coherent within than outside of that segment. In passing, we note that the histogram of coherence coefficients is less uniformly distributed than those of correlation which can guide the choice of τ . Nonetheless both have a local peak at around 0.7 justifying comparing both methods with the same value of τ which we set to 0.3.

Cluster split

Coherence-based Spectral clustering clusters traveling waves, correlation tends to split.



Anatomical fidelity




Figure 9: **Spectral Clustering applied to WBI data.**

(A) Comparison of coherence vs correlation-based spectral clustering. Using correlation, the descending limb of the kidney is clustered into three distinct clusters (left), whereas when using coherence, the three regions are clustered into a single cluster (right). The average cluster activity is shown below the anatomical footprint of each cluster. In both runs with $n_{\text{clusters}} = 40$ and $\tau = 0.3$. The kidney displays traveling waves (Fig. 18) such that the time series of the kidney are relatively similar but time-lagged, and thus coherent but not necessarily that correlated.

(B) Anatomical fidelity of the clusters. The clusters using coherence tend to be more compact allowing for easier anatomical identification. Upon running correlation or coherence-based spectral clustering a group of cells around the aorta showing bistable dynamics are consistently identified. Through other experiments this cluster was shown to be the main sympathetic ganglion: through expansion microscopy we visualize the location of the sympathetic nervous system (green) relative to the vasculature (magenta) finding that the cell bodies hug the main vessels (both aorta and central veins) (right); in green, IHC against eGFP genetically driven by $th1+$ cells expressing eGFP, in magenta, IHC against RFP driven genetically by $flk1+$ (vasculature endothelium) cells, in orange, total protein stain. When using correlation-based clustering (left), the cluster is "mixed" with cells in other tissue types (e.g.: liver, epithelium). When using coherence is used, all other parameters constant, the spatial footprint is better localized.

4.3 Sequential Component Analysis

From molecules to whole organisms, the dynamics of natural living systems depart from "thermodynamic equilibrium" (Gnesotto et al., 2018). Statistically, the state trajectories produced by such systems are non-reversible, i.e. they do not have equal likelihood of occurring

in the reverse temporal direction even in stationary regimes.

Despite the prevalence and importance of sequential neural activity, there is a relative paucity of methods for exploring the spatio-temporal structure of irreversibility in multivariate time series. Here, we introduce Sequential Components Analysis (SCA), a simple yet effective and scalable method for doing this. SCA performs a systematic analysis of spatio-temporal covariances (all time pairs and unit pairs), and extracts the spatio-temporal modes of activity that contribute most to sequentiality. We highlight the main distinguishing features of the method using a toy example. To illustrate the method on real data but still within a setting where there is published prior knowledge about the dynamics within data, we apply it to a well-studied monkey M1 motor activity dataset as well as rat hippocampal data where we show that sequential features separate navigational memories better than mere principal components. The next step will be to apply SCA on WBI data which we leave for later work.

4.3.1 Sequentiality and space-time decomposition

We assume data samples, denoted as X , come as $T \times N$ matrices, where T is the number of time bins and N is the number of units (Fig. 10A, left). We say that a spatio-temporal process $X \in \mathbb{R}^{T \times N}$ is reversible (non-sequential), if, and only if $p(X_{ti}, X_{t'i'}) = p(X_{t'i}, X_{ti'})$ for any two distinct time points (t, t') and any two distinct units (i, i') . This definition generalizes the notion of “reversibility” (or “detailed balance”, which only applies to stationary processes) to the case of non-stationary processes. By negation, X is said to be “sequential” if it is not reversible.

To study the sequentiality of a spatio-temporal process, we first compute its space-time $TN \times TN$ covariance matrix. For ease of notation, we call the centered data X , too. The space-time covariance matrix $C \in \mathbb{R}^{TN \times TN}$ is defined as:

$$C = \left\langle \text{vec}(X) \text{vec}(X)^T \right\rangle. \quad (53)$$

where $\text{vec}(X)$ is the standard vectorization operator (vertical stacking of the columns of X) and $\langle \cdot \rangle$ denotes averaging over different spatio-temporal samples. While in general this matrix is prohibitively large to store or compute, we have developed a variant of the method that entirely bypasses the need to compute C directly (see below).

We then decompose C as the sum of two components: one that contains the symmetric parts of all $T \times T$ cross-covariance matrices between all pairs of units which we call $C^{(+)}$, and one that contains the anti-symmetric counterparts, which we call $C^{(-)}$ (Fig. 10A - center). The latter summarizes the spatio-temporal organization of sequentiality in the data, at least up to second-order moments (i.e. at least for a Gaussian process). Specifically, we define:

$$C^{(+)} = C + \sigma(C) \quad (54)$$

$$C^{(-)} = C - \sigma(C) \quad (55)$$

where, for any matrix $M \in \mathbb{R}^{NT \times NT}$, we use the notation $\sigma(M)$ to denote the matrix obtained by separately transposing each $T \times T$ block of M (there are N^2 of them). We note that $C^{(-)}$ is symmetric but not positive definite (e.g. its trace is zero).

We use these two components to define the degree of sequentiality of a spatio-temporal process as

$$s = \frac{\|C^{(-)}\|_F^2}{\|C^{(+)}\|_F^2} \quad (56)$$

where $\|\cdot\|_F^2$ is the squared Frobenius norm. This metric is upper-bounded by one, and can be used to compare the sequentiality of various datasets.

We show that the spatio-temporal structure of sequentiality can be discovered based on an analysis of $C^{(-)}$'s principal subspace. The methodology is straightforward: we seek the top singular vectors of $C^{(-)}$, and reshape them as $T \times N$ prototypical time series (Fig. 10A, right); these form a subspace of spatio-temporal activity patterns that capture the non-sequential aspects of the organization of space-time cross-covariances in the data.

Scaling the method and cross-validation For large N and/or large T , computing (or even storing) $C^{(-)}$ is intractable. To bypass having to compute $C^{(-)}$ directly, we derive an efficient computation of matrix-vector products of the form $C^{(-)}v$ for any TN -vector $v = \text{vec}(V)$:

$$C^{(-)}v = \text{vec}(\langle \text{Tr}[XV^T]X - XV^TX \rangle) \quad (57)$$

which can be estimated in $\mathcal{O}(\min(KNT^2, KTN^2))$ based on a number K of samples. Using this identity, it is possible to very efficiently compute the principal subspace of $C^{(-)}$ using either randomized SVD algorithms [Halko et al., 2011] or more standard iterative optimization methods for matrix factorization. Here, we directly minimize

$$\mathcal{C}(Z) = \|C^{(-)}C^{(-)T} - ZZ^T\|_F^2 \quad (58)$$

over matrix $Z \in \mathbb{R}^{NT \times r}$, where r is the desired rank; the gradient of $\mathcal{C}$ involves Eq. 57.

We have also developed a bootstrapping method (inspired by Machens, 2010) to estimate the noise floor in the singular values of $C^{(-)}$ – this allows us to determine how many of the top modes contribute significantly.

Extensions Our iterative method based on Eq. 58 allows constraints to be added such as sparsity. SCA with an ℓ_1 penalty on Z can demix sequences that occur in distinct populations (not shown here). We also show that the rotational dynamics model underlying jPCA (Churchland et al., 2012) assumes a particularly simple SCA decomposition: pure rotations in a plane spanned by two orthogonal vectors (x_1, x_2) (starting with a random phase in each trial) give rise to a $C^{(-)}$ of the form:

$$C^{(-)} = (x_1x_2^T - x_2x_1^T) \otimes (sc^T - cs^T) \quad (59)$$

where s and $c \in \mathbb{R}^T$ are a pair of sine and cosine temporal waveforms. In this case, the top two modes of $C^{(-)}$ correspond to rotational motion of the activity vector in the (x_1, x_2) plane. The equation above generalizes to multiple rotations in orthogonal spatial planes, in which case $C^{(-)}$ is the sum of multiple Kronecker products similar to Eq. 59.

4.3.2 Results

Illustrative toy example To illustrate the method, we construct a toy dataset comprised of one rank-2 sequential mode (pure rotation), overwhelmed by two non-sequential modes of much higher variance; then scaled by a transient envelope, plus Gaussian observation noise (Fig. 10B, left). Methods such as PCA can discard important sequential modes if these, for instance, ride on top of larger global fluctuations shared across all neurons. In contrast, SCA is impervious to non-sequential modes irrespective of their variance, and here successfully recovers the hidden sequence (Fig. 10B, right).

Sequences in M1 activity during reaching Movements require largely sequential activation of muscles, which presumably should be reflected in neural activity. We apply SCA to M1

activity (Churchland lab, $N = 218$ neurons) recorded in monkeys while they executed $K = 108$ different reaches (Fig. 1C, left). The data was first averaged across trials for each condition; thus, we analyzed across-conditions covariances. SCA here confirms that the data is highly sequential ($s = 0.48$), with several significant modes which, following reordering of the neurons by peak time in each mode, reveal clear sequences. While $C^{(-)}$'s singular vectors (seq. modes) are orthogonal, this does not in principle guarantee that they correspond to different sequences. In this case though, the sequences appear genuinely different: there is no apparent correlation in the neurons' rank order between pairs of modes (Fig. 1C, right).

Encoding of spatial memories in sequential subspace: hippocampus, navigational task

We apply SCA to unpublished rat hippocampal data recorded during a spatial memory task in a three-arm maze (Fig. 10D, left). The data consists of $N = 315$ neurons recorded over the 8-second delay period of the task ($K = 100$ trials), during which the rats have to run in a wheel before being allowed back into the maze for the next trial. It must then choose one of the three maze arms to collect a reward, before returning to the wheel. Critically, the rat must visit the arms in a specific order, and must therefore remember where he came from in the previous trial in order to decide where to go next. The data (courtesy of Brian Lustig, HHMI Janelia) was collected using silicone probes, binned in 20ms bins and filtered with a Gaussian kernel. We compared projections of the data into the top two SCA and PCA components (Fig. 10D, right). The sequential subspace better separates memories of recent navigation episodes, providing evidence that the brain may use sequential dynamics to encode such episodes in the hippocampus.




Figure 10: Please see text for full details. The SCA workflow (A) is applied to a toy dataset (B), to monkey M1 data (C), and to a delayed three-arm spatial memory task demonstrating that memories of recent navigation episodes are better separated in the top sequential subspace (D).

4.4 Non-reversible Kernels and GPFADS

4.4.1 Introduction

A common goal in the analysis of neural (and now physiological) data is to compress large population recordings into sets of interpretable, low-dimensional latent trajectories. This problem can be approached using Gaussian process (GP)-based methods which provide uncertainty quantification and principled model selection. However, standard GP priors do not distinguish between underlying dynamical processes and other forms of temporal autocorrelation. Here, we propose a new family of “dynamical” priors over trajectories, in the form of GP covariance functions that express a property shared by most dynamical systems: temporal non-reversibility. Non-reversibility is a universal signature of autonomous dynamical systems whose state trajectories follow consistent flow fields, such that any observed trajectory could not occur in reverse. Our new multi-output GP kernels can be used as drop-in replacements for standard kernels in multivariate regression, but also in latent variable models such as Gaussian process factor analysis (GPFA). We therefore introduce GPFADS (Gaussian Process Factor Analysis with Dynamical Structure), which models single-trial neural population activity using low-dimensional, non-reversible latent processes. Unlike previously proposed non-reversible multi-output kernels, ours admits a Kronecker factorization enabling fast and memory-efficient learning and inference. We apply GPFADS to synthetic data and show that it correctly recovers ground truth phase portraits. GPFADS also provides a probabilistic generalization of jPCA, a method originally developed for identifying latent rotational dynamics in neural data. When applied to monkey M1 neural recordings, GPFADS discovers latent trajectories with strong dynamical structure in the form of rotations.

4.4.2 Background: Learning latent dynamical models

The body have evolved as a rich dynamical system to control and coordinate the various and sometimes competing aspects of physiology. High-dimensional neural activity can often be efficiently recapitulated by lower dimensional latent dynamics (and likely the same can be said more generally for cellular activity), and multiple methods have been proposed over the years to tackle the challenge of extracting interpretable and actionable latent trajectories.

A first class of methods focuses on explicitly learning the transition function of an underlying dynamical system. These include parametric models such as linear dynamical systems (LDS) models (Roweis and Ghahramani, 1999; Churchland et al., 2012; Buesing et al., 2012a,b; Macke et al., 2011) and switching variants (Petreska et al., 2011; Linderman et al., 2017), probabilistic deep learning approaches such as LFADS (Pandarinath et al., 2018), as well as more flexible non-parametric models of the transition function and its uncertainty (Deisenroth and Rasmussen, 2011; Duncker et al., 2019). While appealing in principle, the latter methods do not allow exact inference, must combat pervasive local optima during training, and are computationally intensive. As such, they have yet to be more widely adopted in the field.

The second class of methods focuses on modeling the statistics of the latent processes directly, rather than learning a dynamical model for them. Such methods include Gaussian-process factor analysis (GPFA) and variants (Yu et al., 2009). Gaussian process (GP)-based methods are data efficient and have closed-form formulas allowing for uncertainty estimation and principled model selection (Rasmussen and Williams, 2006). Yet, these models fail to capture features of dynamical systems beyond basic smoothness properties, limiting our capacity to study the dynamics of brain computations.

We set out to bridge these two classes of models by imparting some notion of “dynamics” to GP-based models. A key property of autonomous dynamical systems is that they define a consistent mean flow field in state space, such that any segment of state-trajectory produced by the system is unlikely to be visited in the opposite direction (though this is not true of strongly input-driven, or partially observed systems). To capture this property in the Gaussian process framework, we introduce a measure of second-order nonreversibility and derive a new family of GP covariance functions for which this measure can be made arbitrarily large. These kernels can be derived from a variety of usual scalar stationary covariance functions, such as the squared-exponential kernel or the more expressive spectral mixture kernel (Wilson and Adams, 2013a). Conveniently, our nonreversible multi-output GP construction affords a specific Kronecker structure; we discuss how this property enables scalability to very large datasets. We validate these kernels on a regression problem where we show that non-reversible covariances yield better model fits than reversible ones for datasets originating from dynamical systems. We then introduce non-reversible kernels in GPFA, and call this variant GPFADS, Gaussian Process Factor Analysis with Dynamical Structure. We show how GPFADS allows demixing of dynamical processes from other high-variance latent distractors, even where demixing could not be performed by comparing lengthscales alone. Finally, we apply GPFADS to population recordings originating from the primary motor cortex of monkeys. We find that it discovers latent processes with clear rotational structure, consistent with earlier findings (Churchland et al., 2012).

4.4.3 Background: Gaussian Process Factor Analysis (GPFA)

Notation In the following, we use bold $\mathbf{x}$ for column vectors, and capital X for matrices whose elements we denote by x_{ij} . In any context where matrix X has been introduced, $\tilde{\mathbf{x}}$ is a shorthand notation for $\text{vec}(X^\top)$, where $\text{vec}(\cdot)$ is the operator which vertically stacks the columns of the matrix. The transpose is needed for consistency with the convention used in the rest of the paper, which requires that the rows be transposed and stacked vertically instead of columns. Finally, I_N denotes the $N \times N$ identity matrix, and $\mathbf{1}_N$ denotes the column vector whose N elements are all ones.

Latent variable models offer a parsimonious way of capturing statistical dependencies in multivariate time series. Gaussian process factor analysis (GPFA; Yu et al., 2009) is one such popular model used for simultaneous dimensionality reduction and denoising/smoothing of neural population recordings. Missing data are straightforward to handle, but for simplicity of exposition, we assume that observations $\mathbf{y}(t) \in \mathbb{R}^N$ are available for each of N variates at each of T time points. GPFA assumes that such observations arise as the noisy linear combination of a smaller set of M latent trajectories, $\mathbf{x}(t) = (x_1(t), \dots, x_M(t))^\top \in \mathbb{R}^M$, each modelled as an independent Gaussian process. Formally,

$$\begin{aligned} x_i(\cdot) &\sim \mathcal{GP}(0, k_i(\cdot, \cdot)) \\ \mathbf{y}(t) &\sim \mathcal{N}(\boldsymbol{\mu} + C\mathbf{x}(t), R) \end{aligned} \tag{60}$$

where $k_i(\cdot, \cdot)$ is the covariance function (or “kernel”) of the i^{th} latent GP. The model is trained by maximizing the log marginal likelihood $\mathcal{L}(\boldsymbol{\theta})$ w.r.t. the parameter vector $\boldsymbol{\theta}$, which comprises all kernel parameters (see below), a mean vector $\boldsymbol{\mu} \in \mathbb{R}^{N \times 1}$, a mixing matrix $C \in \mathbb{R}^{N \times M}$, and a diagonal matrix of private observation noise variances $R \in \mathbb{R}^{N \times N}$. For a data sample

$Y \in \mathbb{R}^{N \times T}$, the log marginal likelihood is proportional (up to an additive constant) to

$$\mathcal{L}(\theta, Y) \propto -\log |K_{yy}| - [\tilde{\mathbf{y}} - \boldsymbol{\mu} \otimes \mathbf{1}_T]^\top K_{yy}^{-1} [\tilde{\mathbf{y}} - \boldsymbol{\mu} \otimes \mathbf{1}_T] \quad (61)$$

$$\text{with } K_{yy} = (C \otimes I_T) K_{xx} (C^\top \otimes I_T) + (R \otimes I_T) \quad (62)$$

where $\tilde{\mathbf{y}} = \text{vec}(Y^\top)$, $K_{xx} \in \mathbb{R}^{MT \times MT}$ is the prior Gram matrix, and $\otimes$ denotes the Kronecker product. As the latents are a priori independent in this original formulation, K_{xx} is block diagonal with the i^{th} diagonal block corresponding to the $T \times T$ Gram matrix of latent x_i .

Given a particular observation $\tilde{\mathbf{y}} \in \mathbb{R}^{NT}$, the posterior mean and covariance over latent trajectories are given by:

$$\mathbb{E}(\tilde{\mathbf{x}}|\tilde{\mathbf{y}}) = K_{xx} (C^\top \otimes I_T) K_{yy}^{-1} [\tilde{\mathbf{y}} - \boldsymbol{\mu} \otimes \mathbf{1}_T] \quad (63)$$

$$\text{Cov}(\tilde{\mathbf{x}}|\tilde{\mathbf{y}}) = K_{xx} - K_{xx} (C^\top \otimes I_T) K_{yy}^{-1} (C \otimes I_T) K_{xx}. \quad (64)$$

Note that once $\tilde{\mathbf{v}} = K_{yy}^{-1} [\tilde{\mathbf{y}} - \boldsymbol{\mu} \otimes \mathbf{1}_T]$ is computed, the rest of the computation of the posterior mean can be sped up by using the Kronecker identity $(C^\top \otimes I_T) \tilde{\mathbf{v}} = \text{vec}(V^\top C)$. We further outline how the relevant quantities for inference and learning can be stably and efficiently computed for the original and our own model in [Appendix F.1](#). We also discuss a highly scalable implementation in [Appendix F.2](#).

4.4.4 Nonreversible Gaussian processes

A major limitation of the original GPFA model summarized in [Section 4.4.3](#) is the assumption that the latent processes are independent *a priori*. This in turn severely impairs the ability to extrapolate or look into any learned prior relationships between latents in search for dynamical structure (e.g. consistent phase lags between latents, delays, etc.).

In this paper, we introduce novel multi-output covariance functions aimed at expressing a key property of dynamical systems: that they produce state-space trajectories that follow lawful flow fields and are therefore temporally non-reversible. We begin by formalizing this idea of temporal non-reversibility for stationary GPs, before describing our construction of non-reversible multi-output GP kernels, which we then combine with GPFA, yielding GPFADS, Gaussian Process Factor Analysis with Dynamical Structure.

4.4.5 Quantifying non-reversibility and decomposing multi-output GP covariances

Consider a stationary zero-mean multi-output Gaussian process $\mathbf{x}(t) = (x_1(t), \dots, x_M(t))$ with covariance functions $k_{ij}(\tau) \triangleq \mathbb{E}[x_i(t)x_j(t+\tau)]$. We define $\mathbf{x}$ to be temporally reversible if, and only if, all pairwise cross-covariance functions have no odd part, i.e. $k_{ij}(\tau) = k_{ij}(-\tau)$ for all $i \neq j$ and $\tau \in \mathbb{R}$. This is equivalent to the condition that the spatial cross-covariance matrix $K(\tau) \triangleq \mathbb{E}[\mathbf{x}(t)\mathbf{x}(t+\tau)^\top]$ be symmetric for any lag τ . Thus, only multi-output GPs can be made non-reversible. To quantify departure from pure reversibility in a multi-output GP, we introduce the following measure of non-reversibility:

$$\zeta = \left(\frac{\int_{-\infty}^{\infty} \|K(\tau) - K(-\tau)\|_F^2 d\tau}{\int_{-\infty}^{\infty} \|K(\tau) + K(-\tau)\|_F^2 d\tau} \right)^{1/2} \quad (65)$$

where $\|\cdot\|_F$ denotes the Frobenius norm. In [Appendix A](#), we prove that $0 \leq \zeta \leq 1$. We note that, by this definition, any scalar (one-dimensional) GP is necessarily fully reversible ($\zeta = 0$).




Figure 11: **Sample trajectories from various planar GP kernels**, defined in Eq. 68. Five samples trajectories are shown over an interval of 5 time units for the following kernels, in this order: SE kernel $f(\tau) = \exp(-\tau^2/2)$ with $\alpha = 0$, SE kernel with $\alpha = 1$, cosine kernel $f(\tau) = \cos(0.15 \times 2\pi\tau)$ with $\alpha = 1$, and spectral mixture kernel $f(\tau) = \exp(-\tau^2/(2 \times 2.5^2)) \cos(0.06 \times 2\pi\tau)$ with $\alpha = 1$. Here we have set $\sigma_1 = \sigma_2$ and $\rho = 0$, resulting in spherical planar processes.

Our goal is to construct GP covariance functions that break reversibility. As a first step, we prove in Appendix B that any stationary M -output GP covariance admits a finite “Kronecker” decomposition:

$$K(\tau) = \sum_{\ell=1}^{n^+} \lambda_{\ell}^+ A_{\ell}^+ f_{\ell}^+(\tau) + \sum_{\ell=1}^{n^-} \lambda_{\ell}^- A_{\ell}^- f_{\ell}^-(\tau) \quad \text{with} \quad \begin{cases} \text{Tr}(A_{\ell}^{\pm} A_{\ell'}^{\pm\top}) = \delta_{\ell\ell'} \\ \int f_{\ell}^{\pm}(\tau) f_{\ell'}^{\pm}(\tau) d\tau = \delta_{\ell\ell'} \end{cases} \quad (66)$$

where $n^+ + n^- = M^2$ and $\lambda_{\ell}^{\pm} \geq 0$. In this decomposition, $\{A_{\ell}^{\pm}\}$ is a collection of $M \times M$ symmetric (+) or skew-symmetric (-) matrices, all orthonormal to each other in the sense expressed in Eq. 66. Similarly, $\{f_{\ell}^{\pm}\}$ is a matching set of orthonormal even (+) or odd (-) scalar functions. We assume without loss of generality that the weighting coefficients $\{\lambda_{\ell}^{\pm}\}$ are ordered by decreasing value within each (+) and (-) sets.

The decomposition in Eq. 66 is a “Kronecker” decomposition (Van Loan, 2000), because the Gram matrix instantiating $K(\tau)$ at a discrete set of time points is composed of a sum of Kronecker products: $K = \sum_{\ell} \lambda_{\ell}^+ A_{\ell}^+ \otimes F_{\ell}^+ + \sum_{\ell} \lambda_{\ell}^- A_{\ell}^- \otimes F_{\ell}^-$ with F_{ℓ}^+ and $F_{\ell}^- \in \mathbb{R}^{T \times T}$. This decomposition conveniently isolates terms that either strengthen (+) or break (-) reversibility. In particular, we show in Appendix B that the non-reversibility index of the process is related to the $\{\lambda_{\ell}^{\pm}\}$ coefficients through:

$$\zeta = \left(\frac{\sum (\lambda_{\ell}^-)^2}{\sum (\lambda_{\ell}^+)^2} \right)^{1/2}. \quad (67)$$

Thus, breaking reversibility requires the presence of skew-symmetric/odd terms. However, the decomposition does not immediately tell us how to *construct* such a non-reversible covariance function. Although one can show that the first term $A_1^+ f_1^+(\tau)$ must be positive definite, the addition of even a single $A_1^- f_1^-(\tau)$ odd term will not preserve positive definiteness in general, unless carefully specified. One of the main contributions of this work is to provide a constructive way of building sums of Kronecker products similar to Eq. 66, for which positive-definiteness is preserved while ζ can differ substantially from zero.

stationary kernel $f(\tau)$	Hilbert transform $\mathcal{H}[f](\tau)$
$\exp(-\tau^2/2)$	$2\pi^{-1/2}D(\tau/\sqrt{2})$
$\cos(\omega_0\tau)$	$\sin(\omega_0\tau)$
$\sin(\omega_0\tau)/(\omega_0\tau)$	$[1 - \cos(\omega_0\tau)]/(\omega_0\tau)$
$(1 + \tau^2)^{-1}$	$\tau(1 + \tau^2)^{-1}$
$\exp(- \tau)$	$[e^{-\tau}\text{Ei}(\tau) - e^{\tau}\text{Ei}(-\tau)]/\pi$
$\exp(-\tau^2/2)\cos(\omega_0\tau)$	$\exp(-\tau^2/2)\sin(\omega_0\tau) + \exp(-\omega_0^2/2)\text{Im } w((\tau + j\omega_0)/\sqrt{2})$

Table 1: **Hilbert transforms of usual scalar GP kernels.** Non-unit length-scales can be accommodated via a simple change of variable. Here, $D(\cdot)$ denotes the Dawson function, $\text{Ei}(\cdot)$ is the exponential integral, and $w(\cdot)$ is the Faddeeva function.

4.4.6 Planar non-reversible processes

To build intuition, we begin with a planar (two-output) process, $\mathbf{x}(t) = (x_1(t), x_2(t))^\top$, with zero mean and stationary matrix-valued covariance function $K(\cdot)$. If $x_1(t)$ and $x_2(t)$ are independent ($k_{ij}(\cdot) = \delta_{ij}f(\cdot)$) as in the original GPFA model (Yu et al., 2009), then the process is fully reversible. Consider, instead, the following construction:

$$K(\tau) = \underbrace{\begin{pmatrix} \sigma_1^2 & \sigma_1\sigma_2\rho \\ \sigma_1\sigma_2\rho & \sigma_2^2 \end{pmatrix}}_{A^+} f(\tau) + \alpha \underbrace{\begin{pmatrix} 0 & \sigma_1\sigma_2\sqrt{1-\rho^2} \\ -\sigma_1\sigma_2\sqrt{1-\rho^2} & 0 \end{pmatrix}}_{A^-} \mathcal{H}[f](\tau) \quad (68)$$

where $f(\cdot)$ is any scalar covariance function (an even function), $\mathcal{H}[f](\cdot)$ denotes its Hilbert transform (an odd function), $|\rho| \leq 1$ and $\alpha \in \mathbb{R}$. We show in Appendix C that Eq. 68 is a valid, positive semi-definite covariance, provided that $|\alpha| \leq 1$. Since $\mathcal{H}[f](0) = 0$, the first matrix on the r.h.s. parameterizes the instantaneous covariance $K(0)$ of the two processes (up to a positive scalar given by $f(0)$). Moreover, marginally, both $x_1(t)$ and $x_2(t)$ have temporal autocovariance function $f(\cdot)$.

Importantly x_1 and x_2 are now temporally correlated in such a way that reversibility is broken. In fact, Eq. 67 shows that $|\alpha|$ is related to the non-reversibility index ζ defined in Eq. 65 in the following way:

$$\zeta = |\alpha| \left[\frac{2(1-\rho^2)}{(\sigma_1/\sigma_2)^2 + (\sigma_2/\sigma_1)^2 + 2\rho^2} \right]^{1/2}, \quad (69)$$

which has a maximum of $|\alpha|$ when $\sigma_1 = \sigma_2$ and $\rho = 0$, i.e. for an instantaneously spherical process. Thus, Eq. 68 lets us construct planar GPs with arbitrary degrees of non-reversibility, with ζ ranging from 0 to 1.

For the construction of Eq. 68 to be of any use, one needs a practical way of evaluating the Hilbert transform of the marginal temporal covariance $f(\cdot)$. In Table 1, we provide a list of Hilbert transform pairs for several commonly used stationary GP kernels. Notably, we cover the case of the spectral mixture kernel (SM, last row of the table; Wilson and Adams, 2013a), which currently achieves state-of-the-art results in GP-based extrapolation for one-dimensional timeseries. Although some of the Hilbert transforms that we were able to derive

involve exotic functions, such as the Dawson and Faddeeva functions, these are readily available in most numerical programming environments. Moreover, they have analytical derivatives (Appendix D), such that they can easily be added to standard automatic differentiation software to enable automatic gradient computations for the model evidence e.g. in GP regression or GPFA (see below).

Fig. 11 illustrates the behavior of various spatially spherical planar GP kernels constructed from Eq. 68 with different kernels $f(\cdot)$. In cases where $\zeta = 1$, we emphasize that the time-reversed version of each of the samples shown (or indeed, of any subset thereof) has zero probability density under the prior from which it was drawn (Appendix C).

4.4.7 Fourier domain interpretation

To gain more insight into planar non-reversible GPs, we present an alternative construction of the process defined in Eq. 68, in the frequency domain. This construction can also serve as an alternative proof that Eq. 68 constitutes a valid GP covariance (see also Appendix C). We begin by noting that the Fourier transform of $\mathcal{H}[f]$ equals $-j \operatorname{sgn}(\omega) \hat{f}(\omega)$, where $\hat{f}(\omega)$ is the Fourier transform of f . In other words, $\mathcal{H}[f]$ is the real function that is phase shifted by $\pi/2$ away from f at all frequencies. Thus, using the Wiener-Khinchin theorem, the Fourier-domain equivalent of Eq. 68 is:

$$\mathbb{E} [\hat{x}_1(\omega) \overline{\hat{x}_1(\omega)}] = \sigma_1^2 \hat{f}(\omega), \quad \mathbb{E} [\hat{x}_2(\omega) \overline{\hat{x}_2(\omega)}] = \sigma_2^2 \hat{f}(\omega), \quad (70)$$

$$\mathbb{E} [\hat{x}_1(\omega) \overline{\hat{x}_2(\omega)}] = \sigma_1 \sigma_2 \hat{f}(\omega) [\rho - \alpha \sqrt{1 - \rho^2} j \operatorname{sgn}(\omega)], \quad (71)$$

where $\overline{\cdot}$ denote the complex conjugate and $\mathbb{E}[\cdot]$ denotes expectations w.r.t. the joint processes $(\hat{x}_1, \hat{x}_2)$ specified in the frequency domain. It is easy to verify that these (cross-)spectral densities can be achieved by sampling the two processes according to:

$$\hat{x}_1(\omega) = \sigma_1 \sqrt{\hat{f}(\omega)} [\hat{\varepsilon}_1(\omega) \sqrt{1 - \beta} + \hat{\eta}(\omega) \sqrt{\beta}] \quad (72)$$

$$\hat{x}_2(\omega) = \sigma_2 \sqrt{\hat{f}(\omega)} [\hat{\varepsilon}_2(\omega) \sqrt{1 - \beta} + \hat{\eta}(\omega) \exp(j \operatorname{sgn}(\omega) \varphi) \sqrt{\beta}] \quad (73)$$

where $\varepsilon_1(t)$, $\varepsilon_2(t)$ and $\eta(t)$ are independent white noise processes with unit variance, and β and φ obey the following parameter correspondance: $\rho = \beta \cos(\varphi)$ and $\alpha \sqrt{1 - \rho^2} = \beta \sin(\varphi)$. In other words, $x_1(t)$ and $x_2(t)$ are each entrained to a common latent process $\eta(t)$ with some degree of coherence β , and some frequency-independent phase lag (0 for x_1 without loss of generality, and φ for x_2). In particular, for an instantaneously uncorrelated joint process ($\rho = 0$), we have $\varphi = \pi/2$ and $\beta = \alpha$. Spatial correlations $\rho \neq 0$ can be introduced through phase shifts φ different from $\pi/2$.

4.4.8 Higher-dimensional non-reversible priors

The non-reversible planar processes described in Section 4.4.6 can be extended to M -output processes with $M > 2$ in several ways. Here, we focus on simple combinations of individual planes, but see Appendix E for potentially more flexible approaches. Specifically, we construct an M -output covariance function as a superposition of planar processes of the form of Eq. 68:

$$K(\tau) = \sum_{1 \leq i < j \leq M} A^{ij+} f_{ij}(\tau) + \alpha_{ij} A^{ij-} \mathcal{H}[f_{ij}](\tau) \quad (74)$$




Figure 12: **Non-reversible GP regression on state trajectories of the classical pendulum (top) and the Duffing oscillator (bottom).** (A) Five of the 20 planar trajectories $(x_1(t), x_2(t))$ used for training. Each bout of training data (color-coded) is of the same duration and contains several cycles, the exact number of cycles depending on the (conserved) Hamiltonian energy. (B) Marginal likelihood for 10 individual test trajectories, for the planar reversible SE kernel (x-axis, α pinned to zero) and its non-reversible counterpart (y-axis, α optimized to 0.99 for the pendulum, and 0.9 for the Duffing oscillator). (C) Posterior over $x_2(t)$ in each model, conditioned on the full time course of x_1 (dashed black) but only on the first time bin of x_2 . Ground truth $x_2(t)$ is in solid black.

with

$$A_{uv}^{ij+} = \sigma_{ij,1}^2 \delta_{ui} \delta_{vj} + \sigma_{ij,2}^2 \delta_{uj} \delta_{vi} + \sigma_{ij,1} \sigma_{ij,2} \rho_{ij} (\delta_{ui} \delta_{vj} + \delta_{uj} \delta_{vi}) \quad (\text{symm. PSD matrix}) \quad (75)$$

$$A_{uv}^{ij-} = \sigma_{ij,1} \sigma_{ij,2} \sqrt{1 - \rho_{ij}^2} (\delta_{ui} \delta_{vj} - \delta_{uj} \delta_{vi}) \quad (\text{skew-symm. matrix}) \quad (76)$$

and $|\alpha_{ij}| \leq 1$. Note that A_{ij}^+ and A_{ij}^- are defined in the same way as A^+ and A^- in Eq. 68, and involve only two of the latent dimensions (i and j). Thus, each term in the sum describes a covariance over a pair of dimensions. This sum of planar kernels is motivated by the general decomposition in Eq. 66, though it does not obey the orthogonality constraints therein. In our GPFADS experiments, we further truncate this sum to $M/2$ non-overlapping planes with no shared latent dimensions.

4.4.9 Experiments

In this section, we begin by demonstrating the utility of non-reversible GP priors for modeling time-series data produced by dynamical systems. We then go on to introduce such non-reversible priors in GPFA, and show that GPFADS recovers the Markov state of low-dimensional dynamical systems embedded in high-dimensional data. We also apply GPFADS to primary motor cortex data, where it automatically discovers rotational dynamics that have been shown to emerge during reaching movements (Churchland et al., 2012).




Figure 13: **GPFADS recovers the state trajectories of a dynamical system embedded in a high-dimensional ambient space.** GPFADS and GPFA latent trajectories (posterior mean) inferred from observations arising as a mixture of trajectories produced by the Van der Pol oscillator and a noisy process of equal variance, each embedded in 6D and added together with white noise (see text for details). For GPFADS, we show the two latent planes over which separate non-reversible planar priors were placed. For GPFA, we only show two arbitrarily chosen planes, but any other combination of latent dimensions resulted in similar unstructured trajectories.

4.4.9.1 Non-reversible GP priors better capture dynamics

Fig. 12 illustrates the relevance of non-reversible planar processes of the type of Eq. 68 for modelling multivariate time-series produced by dynamical systems. We simulated state trajectories of the classical pendulum ($\dot{x}_1 = x_2$ and $\dot{x}_2 = -\sin(x_1)$) as well as the Duffing oscillator ($\dot{x}_1 = x_2$ and $\dot{x}_2 = x_1 - x_1^3$), starting from random initial conditions (Fig. 12A). We then fitted a GP with kernel given by Eq. 68, either with α optimized as part of the fit (‘non-rev’), or pinned to zero (‘rev’). The non-reversible model consistently outperformed the reversible one on cross-validated marginal likelihood (Fig. 12B). Importantly, by optimizing the non-reversibility parameter α , the model learned to capture the phase relationship between x_1 and x_2 , resulting in much better extrapolations than for the reversible model (Fig. 12C). In particular, it was possible to accurately reconstruct x_2 by only conditioning on x_1 and the initial condition for x_2 .

4.4.9.2 GPFADS: recovering embedded latent dynamical systems

We now combine GPFA with the non-reversible multi-output priors introduced in Section 4.4.4, and call this combination GPFADS, Gaussian Process Factor Analysis with Dynamical Structure. In this section, we investigate the extent to which this extension of GPFA allows us to learn something about the dynamics that might underlie a set of multivariate time-series. We reason that noise tends to be more time-reversible than signal generated from a dynamical process, such that placing a non-reversible prior over latent trajectories might let us demix signal with dynamical structure from noise. To demonstrate this, we embedded a 2D dynamical system, the Van der Pol oscillator ($\dot{x}_1 = x_2$ and $\dot{x}_2 = (1 - x_1^2)x_2 - x_1$), into a higher dimensional ambient space ($N = 6$). We also included another (orthogonal) latent plane in which activity was drawn independently for each of the two dimensions from a GP with squared-exponential kernel. We matched the timescales of this reversible “distractor” process to the characteristic timescales of the Van der Pol oscillator.

We trained both GPFADS and GPFA with $M = 4$ latent dimensions on the same set of 50 trajectories, with the Van der Pol oscillator seeded with random initial states in each one. For GPFADS, we used the kernel described in Eq. 74 with all $f_{ij}(\cdot)$ set to the squared-exponential kernel (with independent hyperparameters), and with the sum over (i, j) planes restricted to $(1, 2)$ and $(3, 4)$ – i.e. two independent, orthogonal planes. For GPFA, we placed independent

squared-exponential priors on each of the 4 latent dimensions (Yu et al., 2009). We note that the two models had the same number of parameters: GPFA had two more timescales than GPFADS, but the latter model had two learnable non-reversibility parameters α_{12} and α_{34} .

We found that GPFADS successfully demixed the plane containing the oscillator from that containing the distractor process. Indeed, after training, one of the two latent planes was highly non-reversible ($|\alpha_{12}| = 0.88$, vs. $|\alpha_{34}| = 0.13$), and posterior trajectories in the non-reversible plane recovered the state trajectories of the Van der Pol oscillator (Fig. 13, left). In contrast, despite GPFA being able to correctly learn the various timescales in the latent processes, it failed to demix signal from noise, such that no clear dynamical picture emerged (Fig. 13, right). GPFA also performed worse than GPFADS based on the cross-validated marginal likelihood (not shown).

4.4.9.3 Uncovering rotational dynamics in M1

Collective neural activity in monkey and human primary motor cortex (M1) has been shown to embed strong rotational latent dynamics (Churchland et al., 2012; Pandarinath et al., 2018). The extraction of these dynamics has historically relied on a method called jPCA purposely designed to extract latent rotations wherever they exist. At the heart of such analysis is the desire to reveal dynamical structure in population activity, but one potential concern with jPCA is that it biases this search towards pure rotations, whereas dynamics could in fact be of other forms. With this in mind, we applied GPFADS to M1 population recordings performed in monkey during reaching (Fig. 14; Churchland et al., 2012), to investigate the extent to which rotational dynamics are revealed by a method which does not explicitly look for them, but only indirectly through a search for non-reversible behavior.

The data consisted of $N = 218$ neurons whose activity time-course was measured in each of 108 different movement conditions, and aligned temporally to the onset of movement. For each neuron, activity was averaged over hundreds of stereotypical repetitions of each movement, and further smoothed and ‘soft-normalized’ (Churchland et al., 2012). We fitted GPFADS with an increasing number of latent dimensions ($M = 2, 4, 6$; Fig. 14A-C). For each M , we used a non-reversible M -output GP kernel of the form described in Eq. 74, though we restricted the number of possible planes to $M/2$ independent planes with no shared dimensions. As C was not constrained to be orthogonal we fixed $\rho = 0$, as any prior spatial correlations in a given plane could in this case be absorbed by a rotation of the corresponding two columns of C . Due to the smoothing of neural activity at pre-processing stage (which we did not control), we found that fitting GPFA(DS) was prone to so-called Heywood cases where some diagonal elements of R in Eq. 60 converge to very small values if allowed to (Heywood, 1931; Martin and McDonald, 1975). To circumvent this, here we constrained $R \propto I$, but note that this issue would likely not arise in the analysis of single-trial, spiking data.

For $M = 2$, GPFADS learned a mixing matrix C that was near identical (up to a rotation) to the one learned by the original GPFA model (not shown). This is not surprising: for $M = 2$, a good fit for GPFADS and GPFA alike is likely to be one in which the two columns of C capture the most data variance, regardless of how non-reversible latent activity happens to be in the plane defined by these two columns. Nevertheless, we found that GPFADS learned a non-reversibility parameter $\alpha_{12} = 0.72$ in this case, indicating that dynamics were fairly non-reversible in this top subspace. Importantly, when GPFADS was fit with a larger latent dimension ($M = 4$ or 6 ; Fig. 14B-C), it cleanly segregated latent trajectories into strongly rotatory planes ($|\alpha_{ij}| \in \{0.94, 0.85, 0.95\}$) and planes that absorbed remaining fluctuations with less apparent




Figure 14: **Discovering latent dynamical structure in primary motor cortex with GPFADS.** (A-C) GPFADS latent trajectories (posterior mean) for $M = 2$ (A, one plane), $M = 4$ (B, two planes) and $M = 6$ (C, three planes), for all movement conditions. The value of $|\alpha_{ij}|$ in Eq. 74 for each of the $M/2$ planes is shown near each plot. All plots share the same scale, and planes are ordered from top to bottom by decreasing total variance in the learnt prior ($\sigma_i^2 + \sigma_j^2$). (D) Difference in marginal likelihood for each of 35 trajectories in the test set, between GPFADS (in which the non-reversibility parameters $\{\alpha_{ij}\}$ are learned) and a similar model where these parameters are pinned to zero. Colors refer to the 3 models (see title colors in A-C). Results were found to be highly consistent over independent splits of the 108 conditions into train and test sets.

dynamical structure ($|\alpha_{ij}| \in \{0.35, 0.59\}$). With increasing latent dimension, we found that allowing for non-reversibility in the prior yielded increasing benefits over an equivalent model where all $\{\alpha_{ij}\}$ parameters were set to zero (and the other parameters optimized as normal; Fig. 14D).

4.4.10 Discussion

A great challenge in neuroscience and whole-body systems physiology is to unravel the dynamical mechanisms that underlie bodily computations. As a first step to this, many data analysis methods focus on inferring latent processes which compactly summarize observations of cellular activity in various tasks. Gaussian process-based methods, such as GPFA (Yu et al., 2009), offer data-efficient ways of extracting such latent processes along with associated uncertainty. However, these methods fail to explicitly capture the dynamical nature of neural activity beyond basic smoothness properties. Here, we set out to impart some notion of “dynamics” to GP-based models, using temporal non-reversibility as a proxy for dynamics. We introduced a measure of second-order non-reversibility and derived a new family of GPs for which any sample has lower probability of occurring in reverse. We found that these priors outperform standard reversible ones on a number of datasets known to emanate from dynamical systems, including recordings from primary motor cortex.

An instance of a non-reversible kernel was introduced previously by Parra and Tobar (2017a), which extended the spectral mixture model to a multi-output covariance expressing a variety of

time delays and phase lags between dimensions (see also [Appendix G](#)). Here we have taken a different approach using the Hilbert transform, which – unlike [Parra and Tobar \(2017a\)](#)’s kernel – admits a Kronecker factorization enabling scalability to large datasets.

While temporal non-reversibility is an expected property of most dynamical systems in which state trajectories follow a lawful flow-field (unless they are strongly input driven), it is an incomplete characterization. In particular, the trajectories generated by our non-reversible GP models ([Fig. 11](#)) often cross over, which would not occur in an autonomous dynamical system where the flow would be entirely determined by the momentary state (unless the state is only partially observed). It would be interesting to explore non-reversible GP kernels that also express this complementary property — while the cosine kernel in [Table 1](#) satisfies both properties, it is unclear if a more general, less constraining form exists. Such models might be particularly well-suited for modeling population activity in M1 which is markedly “untangled” ([Russo et al., 2018](#)), i.e. lacks cross-overs.

Other non-probabilistic methods have been proposed for reducing the dimensionality of datasets whilst preserving “dynamical” characteristics. For example, Dynamical Components Analysis (DCA; [Clark et al., 2019](#)) seeks the lower-dimensional subspace that maximizes predictive information. The authors showed that DCA can successfully recover the Markov state of low-dimensional dynamical systems embedded in high dimensions, similarly to GPFADS in [Fig. 13](#). GPFADS is even more closely related to another dimensionality reduction technique which we have proposed previously, Sequential Components Analysis (SCA; [Rutten et al., 2020](#)), in which the low-dimensional projection is chosen to maximize our measure of second-order non-reversibility in [Eq. 65](#). However, in contrast to SCA, GPFADS does not explicitly seek to maximize this measure, but instead automatically learns the degree of non-reversibility (determined by the kernel hyper-parameters) that best explains the data.

5 Application: Studying distributed physiology using Whole Body Imaging

5.1 Introduction

Despite the knowledge that cells throughout an organism interact and that the function of a tissue is defined by that which it accomplishes for or by the rest of the body, many fields of biology are defined by the tissue they focus on (e.g.: neuroscience, nephrology, immunology) with studies primarily hovering around individual cell types (Buchman, 2003). This tissue-centric approach arose in part from methodological reasons. Whole Body Imaging (WBI) allows one to take on a more question-centric approach and search in a less bias way for the cellular players engaged in the process of interest and once identified, focus on these. This offers the advantage of not having to primarily rely on literature to select an initialization point for a study, however, the approach also comes with drawbacks. For example, the physiological responses to cold have been extensively studied (Castellani and Young, 2016), yet, upon exposing a sample to cold water, we found that one of the cell types with the largest Ca^{2+} transients in response to cold exposure were chondrocytes (Fig. 15). This ability of cartilage to directly sense and respond to cold has not yet been unreported in the literature to our knowledge and, with joint pain being known to be affected by cold (Fernandes et al., 2016), the observed high sensitivity of chondrocytes to thermal changes would motivate a new line of investigation. Whole-body ExM allows to follow up such observation and validate cell type identity (Fig. 15 D), discern morphology and later interrogate changes in gene expression profiles for instance; however, we note that the question of interpreting Ca^{2+} responses, inferring the downstream consequences, still remains to be answered on a case by case basis.

Interpreting Ca^{2+} fluctuations The above example, sensitivity of chondrocytes to cold, illustrates the ease of discovery through WBI, but also highlights that a) only one aspect of cellular dynamics is being measured, cytosolic Ca^{2+} concentrations, and, b) that the question of interpretation remains: "what does Ca^{2+} transients in cell X imply?". In this first presentation of WBI, we make use of genetically-encoded calcium indicators (GECIs), such as GCaMP to measure the fluctuations in cytosolic calcium concentrations. Calcium sensors have been extensively optimized and used for imaging neuronal activity, in which Ca^{2+} transients are largely used as a proxy for electrical activity. Indeed, upon membrane depolarization, voltage-gated channels open leading to Ca^{2+} influx. As such, Ca^{2+} waveforms in neurons are relatively stereotyped, timeseries are often deconvolved to infer spike-timing and neuronal Ca^{2+} transients are (very generally speaking) interpreted as a signature of likely synaptic release. While this thinking can be somewhat generalized to cell types which are excitable (i.e.: depolarizable), such as beta cells from the endocrine pancreas, interpreting Ca^{2+} transients from non-excitable cells calls for more consideration. Ca^{2+} is a universal cellular messenger governing many aspects of biology (e.g.: transcription regulation, exo/endocytosis, contraction, metabolism, cell turn-over, etc. (Clapham, 2007, Bagur and Hajnóczky, 2017)); this makes it a useful catch-all proxy for "cellular activity", however, its ubiquitousness in mediating intra-cellular events also make it less informative, a trade-off of the WBI approach. Thankfully, the wealth of research on the role of Ca^{2+} signalling and on the physiology of individual cell types provide context within which to interpret WBI observations. In the future, WBI could be extended to imaging other metabolite sensors and other aspects of physiology to provide more context in which to interpret cellular Ca^{2+} responses.

One illustration of how one can contextualize WBI observations and build testable hypothesis

would be for instance: continuing with the example of chondrocyte responding to cold stimuli, it has been shown that a) cold exposure causes joint pain-sensitivity in an arthritic mouse model in a TRPA1 dependent way (Fernandes et al., 2016), b) TRPA1 is highly Ca^{2+} permeable (Hu et al., 2021), c) TRPA1 was recently found to be expressed in chondrocytes (Nummenmaa et al., 2021). Together these papers provide one example of a testable mechanistic interpretation to explain the chondrocytic Ca^{2+} response to cold (e.g.: chondrocyte express TRPA1 channels which confer cold sensitivity and lead to intra-cellular Ca^{2+} flux upon activation, potentially inducing the secretion of chemical factors), as well as line of query as to whether sensitization to cold or other physiological responses to cold are mediated by factors secreted by chondrocytes (e.g.: is the rise in joint blood flow upon cold exposure caused by factors released from chondrocyte? If so, should we be thinking of chondrocytic tissue as much as an endocrine/paracrine gland as one a purely mechanical tissue?). This is one illustration of how despite the lack of specificity of Ca^{2+} , the wealth of existing literature can often provide context in which to interpret unexpected WBI observations and outline paths to build on and extend them.




Figure 15: **Finding stimulus-sensitive cells through whole body imaging.**

(A) Of all cells, chondrocyte display large calcium transients in response to acute cold stimuli. Left image shows a transverse plane of the head of a 7 d.p.f Tg(ubi:tTA;TRE:GCaMP7f) larval zebrafish. The parachordal cartilage forming the lateral wall of the cranium is delineated in white. A sequence of images (A1-4) shows progression of Ca^{2+} dynamics before and after cold water exposure.

(B) Average parachordal cartilage Ca^{2+} levels triggered on cold stimulus delivery showing time locked response and response attenuation over three trials. Trials were separated by 15 mins; at each trial a drop of 10°C water was added to the 2ml room temperature water bath.

(C) Comparison of single cell Ca^{2+} levels before and after stimulus delivery across trials and fish.

(D) Validation of cell type through Whole-Body ExM-FISH ([Appendix H.1](#)) confirming chondrocyte identity. The image shows the same anatomical region as in A (lateral anterior cranium). In green, total protein stain (NHS Alexa-488); in red, HCR probe against col2a1a. The col2a1a gene encodes the alpha-1 chain of type II collagen alpha-1 chain and is used as a marker for cartilage ([Dale and Topczewski, 2011](#)).

(E) Illustration of the location of imaging plane acquired in A. Cartilage in other ares of the fish showed similar responses.

5.2 Label-free tissue type identification assisted by 3D anatomical model

As the interpretation of Ca^{2+} dynamics is inherently tissue specific ([Section 5.1](#)), being able to assign identity to as many cells being imaged is critical to interpreting WBI observations as this enables one to leverage existing literature on that specific cell-type. The identification

of cells could either be done anatomically or based on functional activity. We found that in a static image of the sample, differential baseline expression of FPs and the texture of the tissue contained enough information to identify most of the major organs and many individual cell-types. These include: kidney (glomeruli, descending and ascending limbs, distal tubules), liver, gallbladder, endocrine and exocrine pancreas, cartilage, brain, spinal cord, muscle (Fig. 1A), as well as smaller populations such as the pineal gland, choroid plexus, and endothelium of blood vessels. Nevertheless, there are sections of the animals with less structure where identifying individual cell types is difficult. One such example is the space directly above the hindgut and below the notochord that in Fig. 1C is labeled "connective tissue". In this space, reside an amorphous group of cells likely containing a mixture of fibroblasts and perhaps immune cells. Such a terra-icognita could have been the case for much of the body, making the data much less valuable, thus highlighting the fortuitousness of the fact that major cell types are easily distinguishable. Tissue identification can be carried out and validated through comparisons with online anatomical references, transgenic lines labeling specific cell types and in-situ-hybridization experiments. We also relied heavily on whole body ExM datasets to be able to familiarize ourselves with the 3D visceral anatomy. In order to be able to summarize these data and communicate them to others, we built a simplified 3D anatomical model of the larval zebrafish. A sample was expanded using the Whole-Body IHC protocol, and the organs were hand annotated using Imaris. From this, binary masks were made and converted into meshes using Blender. The meshes were simplified and symmetrized when appropriate. This results in series of components representing various bodily compartments with accurate relative sizes and ratios which can be used as a means to familiarize oneself with the 3D geometry of the larval zebrafish, facilitating label-free tissue recognition, as well as a communication tool. (Fig. 1A2) shows a rendering of the model with a subset of the organs shown. The model contains at the moment about 50 components and can be readily expanded.

5.3 Spectra of dynamics

In addition to identifying cells based on their different morphology (Section 5.2), we found that a much more granular level of cell-lineage identification is achievable by considering the Ca^{2+} dynamics of cells: activity traces from cells across the body can be clustered based on their similarity, and location of cells assigned to cluster visualized in anatomical space. We found that the distribution of cells within a cluster, as well as the types of Ca^{2+} waveforms in many cases pointed to clear cell types and subdivided groups identified through anatomy alone. More specifically, after data registration and segmentation (Fig. 16 B), we run spectral clustering as described in section 2 on the timeseries data. Fig. 16C shows the spatial distribution of clusters in anatomical space for a plane; one can see, for example, one large cluster colored pink corresponding to an epaxial muscle group. As well as providing a means to better identify cell-types, clustering can be used to reduce data dimension. By volume, much of the body consists of muscle cells, which have high amplitude activity traces but which are also very low-dimensional. We found that it was important to use methods which are not excessively biased by cluster size, density, or variance explained, such as Spectral Clustering, to make sure not to end up with an overrepresentation of muscle tissue. Among the clusters we found there to be a range of types of dynamics: Fig. 16G shows the average population activity for some example clusters, selected manually to exemplify the presence of similar and dissimilar clusters. We go into a more detailed exposition of two of the clusters to illustrate how one can identify clusters and how ExM can be used to better interpret them.




Figure 16: Segmentation and clustering of WBI data can separate individual tissue types.

(A) After registration, WBI data is segmented into cellular-level functional units and then further clustered into tissue-level clusters. The top image shows a typical field of view of a WBI experiment. Multiple such planes are acquired spanning approximately half the width of the sample. The bottom image shows a high-magnification view of the proximal kidney.

(B) Segmentation. The data is segmented using the Voluseg pipeline (Section 2.5.1). The footprint of the extracted rois are shown in anatomical space. The area void of ROIs corresponds to ear which is mostly void of cells and therefore not segmented, as pixels are at background level and have little variance. The cellular-level segments corresponding to the proximal kidney are shown in the lower row.

(C) Clustering. The segmented data is clustered using Spectral Clustering as described in [Section 2.5.2](#). The footprint of the clusters are shown in anatomical space - same plane as in A (only a single plane is shown for conciseness). As in B, the functional clusters of the proximal kidney are shown at higher magnification showing spatial continuity of distinct clusters. Parameters can be altered to make the clustering more or less granular.

(D) Anatomical location of clusters shown in [Fig. 16G](#), color-matched, overlaid on anatomical reference. (For clarity, not all clusters are shown - black ones were omitted). Individual clusters are discussed in the text.

(E) Correlation matrix of all extracted clusters. Each row corresponds to one cluster; rows were sorted using linkage analysis, linkage shown.

(F) Illustration of the location of the plane (blue plane) shown in A and organs that it spans.

(G) Cluster time series. The mean activity of the clusters are shown for a subset of them. Some clusters appear very similar (e.g.: top row, last three columns), while others share little resemblance with any of the others.

As a first example, the cluster on [Fig. 16G](#), row2, column 2 (R2C2, dark green) shows distinct discrete burst-like events. Its anatomical footprint is shown in [Fig. 17A](#) and maps back to the gallbladder sphincter located in the most ventral part of the gallbladder. The gallbladder is shown in light blue in the [Fig. 17D](#) graphic with the location of sphincter shaded in dark blue. This sphincter consists of an annulus of smooth muscle cells and connects the gallbladder to the gut. This sphincter contracts rhythmically but is also observed to be sensitive to external cues showing a reduction in activity upon exposure to hypoxia or UV stress (not shown). This sphincter can be readily identified in ExM volumes, and labeling the autonomic nervous system using Whole-Body ExM-IHC allows us observe that a post-ganglionic vagal neuron lies directly next to the sphincter with projections wrapping around the sphincter ([Fig. 17B](#)). Having a) identified the gallbladder sphincter (as one example) and its spontaneous contractions through WBI, b) seen its possible innervation via ExM-IHC, and c) its sensitivity to external stimuli, one would next want to gain control over this node to be able to ask questions of causality. As one example of this, we can express channelrhodopsin (CoChR) in smooth muscle cells thereby being gaining optogenetic control over visceral muscle: ([Fig. 17C](#)) shows an image of a line expressing pan-cellularly jRGECO1b, a red calcium indicator along with CoChR in a subset of cells including smooth muscle cells. A dense green band is visible along the ventral border of the gallbladder where the sphincter lies.

The dual transgenic line (Tg(ubi:tTA;TRE:jRGECO1b)xTg(Isl1crest1:CoChR-eGFP)) then allows us to stimulate the sphincter (or any other areas containing smooth muscle: e.g.: gastrointestinal tract) using a 488nm laser and a DMD for periods of 2 seconds. ([Fig. 17E-F](#)) shows example trials. Carrying out optogenetic experiments in visceral tissue can be questionable, as the influx of extracellular Ca^{2+} can be difficult to interpret; smooth muscle is one tissue for which extra-cellular Ca^{2+} influx is the physiological means of activation and induces a measurable output: contraction; therefore, we propose it as an appropriate application of standard optogenetic tools.




Figure 17: Gallbladder identification, imaging and optogenetic activation through whole-body imaging and expansion microscopy.

(A) Anatomical footprint of gallbladder sphincter identified through unsupervised clustering (Fig. 16G, row 2, column 2) overlaid on anatomical reference located in region of gallbladder sphincter (top). Activity trace copied for reference (bottom).

(B) Visualization of gallbladder innervation through Whole-body ExM-IHC. Image showing expanded sample, with crop around gallbladder sphincter. In green, total protein stain, NHS-AF488. In magenta, IHC of eGFP expressed in cells of the autonomic nervous system by transgenic line Tg(phox2b:eGFP); arrow pointing to post-ganglionic neuron next to gall bladder duct.

(C) Optogenetic control over gallbladder sphincter. Image showing double transgenic Tg(ubi:tTA;TRE:jRGECO1b)xTg(Isl1crest1:CoChR-eGFP) expressing red calcium indicator pan-cellularly and opsin in smooth muscle cells. Large density of GFP+ cell observed in ventral portion of gallbladder.

(D) Illustration of relative location of gallbladder within animal and gallbladder sphincter (area shaded in dark blue within the organ shaded in light blue).

(E) Optogenetic induced contraction of gallbladder sphincter. Two example trials shown with sphincter in baseline state (left), and 5 seconds after stimulation (right). The bottom of the gallbladder becomes smaller, whilst the dorsal part remains largely unchanged, as expected from the density of smooth muscle cells concentrated in the ventral portion.

(F) Response to optogenetic activation of smooth muscle around the gallbladder sphincter. Trial triggered average (top). Full activity trace (bottom). We note that the trace only aims to capture "response" - ROI was drawn over sphincter - both Ca^{2+} changes and motion affect the change in fluorescence.

To contrast with the example above of the gallbladder sphincter cluster, which has a unique activity pattern not seen anywhere else in the body, row 1 has three clusters with similar activity (columns 2, 3, 4) which map to muscle, spinal cord neurons, and a population of cells lining the notochord, respectively. Given that spinal cord neurons induce muscle contraction, this is as expected. We return to the third population later on (Fig. 19C). The activity trace Fig. 16G R3C4 shows relatively regular oscillatory activity and maps to the outer border of the intestine where the enteric nervous system lies. The traces shown at Fig. 16G R5C2 and R5C4 map to two different portions of the proximal part of the kidney which appear active at similar times but with slightly different dynamics. Indeed, upon inspection we found that the kidney exhibits traveling Ca^{2+} activity along the nephric ducts on the order of minutes (Fig. 18A-D). The kidney in the larval zebrafish consists of a central midline glomeruli from which two tubules emerge which course ventrally before turning dorsally and then caudally (Fig. 18B). The major nephric segments in zebrafish have already been characterized molecularly in the literature delineating the major divisions (Wingert et al., 2007); here we recover the proximal groups (glomeruli, neck and proximal convoluted tubules) but also find that there appears to be a more granular functional division: (Fig. 18C) shows activity from subpopulations all considered to be with the "neck" of the kidney. The kidney can be readily identified in ExM volumes, and labeling motile ciliated cells using Whole-Body ExM-IHC lets us observe that the nephric duct is lined with cells that project long cilia with the lumen of the kidney (Fig. 17E). The kidney is responsible for filtering water and excreting waste product; while this could be thought of as a constant process, the observed activity suggests that this might happen in coordinated bouts perhaps facilitating the generation of ionic gradients which pull water through the system by osmosis or via active coordinated beating of cilia. Furthermore, oscillatory Ca^{2+} activity has also been reported in rat renal proximal tubular cells in similar frequency (prediocyte in minutes) suggesting that this may be a conserved phenotype (Szebényi et al., 2015).




Figure 18: **Traveling Ca^{2+} waves along the nephric ducts.**

(A) Sequential Ca^{2+} activity along nephric duct. Sequence of images of Ca^{2+} activity over the course of 10min showing propagation of Ca^{2+} activity down nephron.

(B): Illustration showing the tubular kidney which begins midline at the glomeruli and that branches into the left and right nephric tube which make a ventral loop before extending caudally. Inset shows direction of flow as a color gradient from orange to purple.

(C) Functionally defined kidney segments. Average Ca^{2+} activity along kidney segments (top). Raster showing cellular activity along the kidney ordered according anatomically and divided into clusters identified. The most proximal part (proximal tube, PCT1-PCT2) is the first to show activity and progresses downwards, down the distal limb (DL1-DL2), eventually progressing up in the ascending limb (AL). Clusters show differing waveforms with proximal regions having multiple beats, whereas the distal AL displays delayed and temporally more compact activity.

(D) Time to peak shown in anatomical space. Relative time to reach 50% of peak activity during an individual wave; two examples are shown. The time of travel is variable and lasts around 6 minutes.

(E) This same region of the kidney can be readily identified and studied with ExM. Whole fish ExM shows how the kidney tubule contains long motile-ciliated cells which extend within the lumen. The image shows a crop of kidney descending duct of a whole ExM-IHC sample of line labelling motile cilia. In green, total protein stain, NHS-AF488. In magenta, IHC of eGFP expressed in cells expressing *foxj1a* by transgenic line *Tg(foxj1a:eGFP)*.

Although most clusters map to distinctive anatomical locations and have characteristic dynamics by which they can be identified with reasonable confidence, other clusters (e.g.: Fig. 16G R2C4) map to areas where cell-types are either too intermixed (e.g.: gills) or light too scattering (e.g.: deep into the pharynx) and that have none distinctive dynamics remain still elusive, and more work needs to be done to identify them. Ultimately a model would be trained to be able to semiautomate this tissue-type assignment which for now remains manual.

5.4 Motor Coupled Visceral Activity

As well as studying within organ dynamics (Section 5.3), WBI allows one to search for couplings between tissues. As one such example to illustrate this, we can ask how the activity of cells across the body might change when there is skeletal-muscle activity. We note that induction of motor-coupled activity could arise directly as a consequence of muscle contraction through mechanosensation, or via sensing of molecules released by muscle cells, or indirectly, e.g., concurrent neural activation of muscle and non-muscle tissue, which we do not try to disentangle here. As an initial way to find cells whose activity changes upon motor activity within short-timescales, we first identify when the large epaxial muscles contract; we do so by averaging the traces within the identified "epaxial cluster" and running peak identification (Fig. 19A, row 3, top right). We then trigger the data on these timepoints, select a time-window of 20 seconds for instance, and compare average activity over this time-window before and after the contraction by running a paired t-test. Fig. 19 we show a few examples of tissue which exhibit strong acute (order of seconds) motor coupling consistently across samples.

Skin epithelium displayed large time-locked Ca^{2+} transients following muscle bursts (Fig. 19A), likely due to expressing mechanosensitive ion channels (Varadarajan et al., 2022). After initial activation of an epithelium cell, Ca^{2+} activity appeared to propagate radially ((Fig. 19A, row 2) likely due to the presence of gap junctions known to be expressed in epithelium (Moore et al., 2023). We note that not all epithelial cells respond and cells at the extremity of the fins tended to be the most responsive; this could be due to the fact that the ridge of the skin experiences a different force profile due to their contact with the agarose or from these outer border cells being inherently more sensitive.

A well-delineated cluster lining the ventral side of the spinal cord showed prolonged Ca^{2+} transients following large motor bursts (Fig. 19C). A literature search querying for cells with such anatomical distribution identified notochord sheath cells as candidate; we ran ExM-FISH against *col2a+*, a marker for such cells and found the anatomical distribution to match the functionally identified cluster (Fig. 19C, row 2, right). The ExM data also showed that there were no other cells at this boundary. In contrast to epithelium, a well-known, well-studied and easily recognizable cell type covering which is already known to have mechano-sensitive properties, notochord sheath cells are poorly characterized and would not necessarily be expected to have much motor-coupled activity. Moreover, this illustrates a cluster which is difficult to label by anatomy alone in a static WBI image, whilst functional clustering easily picks them out as the population has large frequent largely synchronous activity.

While all the above examples show an increase in Ca^{2+} levels after motor activity, other cells showed anticorrelated activity. A group of cells located on the dorsum of the brain displayed a synchronous suppression of Ca^{2+} levels at each motor bout followed by a small rebound (Fig. 19B). This group of cells had distinct anatomy corresponding to the choroid plexus. To confirm this observation, we imaged Ca^{2+} dynamics in a sparse line (*Tg(foxj1a:GCaMP7f)*) labeling choroid plexus epithelium and recovered the same dynamics (Fig. 19B, row 3). It has thus far not been reported to our knowledge that choroidal plexus activity is sensitive to movement. Choroidal epithelial cells were only recently imaged (Shipley et al., 2020). The choroid plexus is responsible for secreting cerebrospinal fluid and chemokines into the ventricles of the brain, thereby contributing to setting intra-cranial pressure. These cells are multiciliated and cilia are known to have mechanotransduction pathways. From such observations, we could hypothesize that movement which induces CSF flow is sensed through cilia by choroid plexus cells, and that this induces a drop in CSF Ca^{2+} level to inhibit cellular decision making processes which might act on the instantaneous increase in pressure caused by movement as if it were an increase in

baseline intra-cranial pressure and as a result undershooting in the amount of CSF production needed (as one example of interpretation of the observation).




Figure 19: Identification of motor coupled dynamics across the body.

(A) Calcium transients are observed across various bodily tissue types when motor actions are taken. Example of muscle activity (orange trace) with onsets of detected bouts (gray bars) used to compute triggered average activity of other tissues. Concurrent activity in skin epithelium cells (blue trace) which shows large transients at each swim bout (and at other times).

(B) Motor-triggered average of different tissues types. Skin epithelium, notochord sheath cell and ependymal cells show consistent time locked activity to motor activity. In contrast to these tissues, choroid plexus epithelium displays a systematic drop in baseline calcium followed by a rebound in response to motor activity. Muscle, skin and notochord sheath cells traces shown come directly from pan-cellular imaging data. Ependymal cell and choroidal cell responses were confirmed through cell-type specific imaging to rule out signal misallocation - data from co-imaging ($Tg(fox1ja:GCaMP7f)$ and $Tg(ubi:tTA;TRE:jRGECO1b)$) is shown).

(C) Functional tissues responses to motor activity. Examples of functional data from skin epithelium, ependymal cell and choroid plexus before and after motor activity. The magnitude of ependymal cell transient is correlated with strength of motor event - scatter plot of area under the curve of ependymal cell ($Tg(fox1ja:GCaMP7f)$) activity vs muscle activity (roi over muscle in $Tg(ubi:tTA;TRE:jRGECO1b)$).

(D) Identification of notochord sheath cells (NSC) and responsiveness to motor activity.

One cluster identified in [Section 2.5.2](#) localized to the boundary between notochord and spinal cord (column 2): overlay of cluster 4 on anatomical reference, sagittal plane. The cells within the cluster showed large synchronous bursts time-locked to motor activity (rastermap, column 3). A literature search for cells sharing this anatomical distribution pointed to notocord sheath cells. Using whole-body ExM-FISH [Section 3](#), we could localize these cells (column 4): ExM-FISH probing for *col2a* showing labeling of cell population around the notochord. The orientation of D4 is shown in the illustration (column 1): transverse plane. The lack of any remaining cells in this region goes towards confirming the identity of this motor coupled cluster.

5.5 Ependymal Coupled Dynamics

In [Section 5.4](#), we exemplified the use of WBI to recover known and uncover unknown couplings between tissue. These examples were all on the short (seconds) timescales. In this section, we illustrate how WBI also allows one to uncover and study interactions which happen at longer-timescales and how ExM can be instrumental in interpreting WBI results.

In a subset of imaged samples, we observed ultra-slow oscillations of periodicity $\sim 3\text{-}6\text{min}$ across the spinal cord and hindbrain ([Fig. 20A,C](#)). These oscillations persisted for extended periods of time ($>30\text{min}$). The oscillations were not a constant feature and could either emerge or disappear over the course of a multi-hour recording ([Fig. 20A](#)). Oscillations were only observed during periods of quiescence (low motor activity) during nocturnal recordings. The oscillations could not be localized to neurons or astrocytes, nor did imaging pan-neuronal or pan-astrocyte ever recover such dynamics. Upon further inspection, cell bodies exhibiting these ultraslow oscillations were found to be located below the central canal where ependymal cells are known to be located ([Fig. 20B](#)). Ependymal cells are a population derived from glia that line and define the ventricles of the brain. These cells control the production and flow of cerebrospinal fluid (CSF), which exchanges with the brain's interstitial fluid, thereby allowing for turn over and waste clearance ([Del Bigio, 2023](#)). To test whether the source of the oscillation was of ependymal origin, we generated two new transgenic lines: *Tg(fox1ja:mRuby)* and *Tg(fox1ja:GCaMP7f)* to be able to image this population of cells specifically. Imaging *Tg(fox1ja:GCaMP7f)* samples recovered the ultra-slow oscillations confirming the source of these signals. These temporal oscillations were also found to propagate in space; however, the anterior/posterior directionality was not consistent across the CNS. To better characterize this, we computed a linearized kymograph; that is we extracted a line at the base of the CNS and computed line integrals perpendicular to this line, resulting in a $1 \times N$ vector with N being the number of spatial bins used. Repeated over time, this results in an $N \times T$ matrix that can be further analyzed ([Fig. 20E](#)). Spatial and temporal smoothing was done to denoise the data. To identify spatial regions in which there was consistent activity propagation, we ran coherence-based kmeans clustering derived in [Section 4.2.1](#). Applying this algorithm to the data resulted in spatially continuous clusters ([Fig. 20H-I](#)). Noticeably, all clusters were of similar size of about $75\mu\text{m}$ which is approximately the size of a spinal segment. The directionality of the propagation differed between the clusters, with the anterior ones progressing anteriorly and the posterior ones posteriorly ([Fig. 20J](#)). Despite the zones of oscillations seeming to be initiated independently, we note that the frequency of oscillation within a fish is far more similar than across fish and that we never observed samples in which oscillations occurred in only part of the CNS; that is, this oscillatory state seems to be a global shared state and that frequency is also constrained to be similar within a sample.

Appearance of ultraslow oscillations across CNS during periods of sustained inactivity




Figure 20: Ultra slow propagating oscillations across ventricular system.

(A) Ultra-slow oscillations occurrence during periods of sustained quiescence (top). Example of power over 2.5h recording of muscle activity (orange, power in high frequency band in muscle ROI component;) and oscillatory activity (blue, power in low frequency range). Example of sample with cells initially showing no oscillatory activity, then low amplitude oscillations, and then going into persistent oscillations in Tg(ubi:tTA;TRE:GCaMP7f) line (bottom).

(B) Transverse section of Tg(ubi:tTA;TRE:GCaMP7f) sample at level of hindbrain and spinal cord. ROIs drawn at the location of cell bodies showing oscillatory activity. All identified ROIs are midline, ventral to spinal cord.

(C) Example traces of cells within the CNS in Tg(fox1ja:GCaMP7f) line showing ultra-slow oscillations.

(D) Example of the power spectrum of cells with oscillatory activity. Top: in black, average spectrum of oscillating cells across hindbrain and spinal cord. In blue, spectrum of individual cells. Within a sample all cells have peak power at similar frequency - in this shown sample around 3.3Hz. Bottom: average spectrum for another sample showing peak power at around 2.5Hz.

(E) Kymograph of ependymal cell activity imaged in Tg(foxl1a:GCaMP7f) line. Spinal chord and brain were linearized and kymograph computed. Distance is defined with 0 being the most anterior point (about 150 μ m into the hindbrain) and extending posteriorly. Traveling waves are visible though with discontinuities.

(G): Cross-correlation of traces in (E) sorted anatomically showing local peak at $\tau \approx 2.5$ min along whole population (left). Average population cross-correlation (right).

(F): Data projected onto the top two PC modes of the data colored according to time, showing cyclic activity.

(H) Cross-coherence matrix (absolute value) of data shown in C at 2.5Hz ordered anatomically. Matrix shows diagonal block structure. On the right, clustering results from running coherence-kmeans run on data as described in [Section 4.2.1](#), showing spatially continuous clusters.

(I) Anatomical reference image of Tg(foxl1a:GCaMP7f) with cluster limits identified from clustering results of matrix G. Cluster boundaries are found to be spaces at relatively regular interval with an average spatial width of around 70 μ m.

(J) Phase of cross-coherence of (G) shown for groups going from anterior to posterior. In most anterior regions (hindbrain), propagation appears to progress anteriorly, whereas more posteriorly (in the spinal chord), propagation of activity appears to travel caudally. Some groups (e.g.: group 3) show no clear phase difference.

The regularity of the spatial extent of the clusters led us to look for an anatomical basis underlying this structure. As the anatomy of ependymal cells in zebrafish has only been described crudely, with descriptions restricted to the superficial surface of the brain, we acquired a high-resolution anatomical reference of this cell population by using our whole body ExM-IHC protocol ([Section 3](#)) and the Tg(foxl1a:mRuby) line ([Fig. 4A](#)). The samples revealed multiple surprises: upon looking at the spinal cord, we found that ventral ependymal cells, reported to be compact cells in larval zebrafish ([D’Gama et al., 2021](#)), send projections directly ventrally in the form of a dense membranous sheath ([Fig. 21A](#)). Moreover, at regular intervals, these ventral protrusions split bilaterally to send large projections to each of the spinal nerves exiting the spinal cord ([Fig. 21H](#)), and once at the most lateral aspect of the spinal cord, these projections then fan out along the anterior-posterior axis of the spinal cord among the nerve tracts. This morphology was preserved in the hindbrain with greater span. The ventral ependymal cells project down ([Fig. 4B](#)), and at the level of nerve decussation expand in a dense membranous sheath that surrounds the motor nerve tracts. Apposition of ependymal cells and vasculature were also visible based on observations from total-protein stain. To better assess the relative geometry of these two systems, we crossed the ventricular line with a line labeling the vasculature (Tg(flk1:dsRED)) and expanded them jointly ([Fig. 4E](#)). The data showed that indeed the hindbrain ependymal cells projections extend to the very base of the brain where they make contact with the basilar artery of the brain ([Fig. 4G](#)) and then send smaller tree-like branches around the outer edge of the brain. The central arteries which branch off from the basilar artery course upwards up to the central canal, where they bifurcate laterally; along this trajectory, the central arteries are ensheathed at discrete points by these ventral ependymal projections ([Fig. 4H](#)). To further confirm this anatomy, ventral ependymal cells were located in an EM volume and the same morphology could be recovered as well as multiple contact sites with vascular endothelial cells (not shown).




Figure 21: Anatomical coupling between ependymal cells, motor circuits and vasculature. (A) Ependymal cell anatomy visualized *in-vivo* in new transgenic line (foxj1a:mRuby). (A1) Higher magnification view of A, showing spinal populations of ependymal cells fall into two groups, a ventral and dorsal population which make up the boundary of the central canal. (B) Ependymal cell anatomy visualized in expanded sample showing the same plane as in A. Here the midline ventral projections of the hindbrain ependymal cells can be seen to extend to the base of the brain whereas this is not clearly visible in A. (C) Ependymal cell projection at spinal segments. Sagittal section at the level of the spinal cord showing how the ventral population of cells extend bilaterally protrusions which reach out to most lateral extent of the spinal cord. The plane is marked by a dashed line in B. (D): Transverse section of A showing the regular projections of ependymal cells laterally forming discrete puncta located within the region where spinal nerves exit. The plane is marked by a dashed line in B.

In search for less traditional descriptions of ependymal cells, we found that researchers have identified similar cells in anamniote brain and spinal cord. The cells span the entire width of the brain or spinal cord, like radial glia which are progenitor cells in the CNS of developing

mammals, however with a soma contributing to the ependymal lining of the ventricle and long radial processes that abut the pial surface with endfeet-like structures. From this mixed identity, they coined the term "ependymo-radial glia" (ERG), reviewed here (Becker and Becker, 2015). They appear to have regenerative capacity giving rise to motor neurons in the adult (Reimer et al., 2008). Described in the spinal chord of zebrafish (Reimer et al., 2009) in the context of injury and tissue regeneration, the cells described share morphological aspects (Reimer et al., 2009 Fig 3A) (extending ventrally and laterally along the spinal cord), therefore we believe that this is the same population of cells. Furthermore, upon laser ablating a neuron in the spinal cord, we observe a propagating wave of Ca^{2+} activity in ependymal cells radiating out from the lesion site traveling at around 7 $\mu\text{m}/\text{second}$ - lesion responses of ependymal cells have typically been studied on the time-scale of hours and days - this shows that responses begin immediately. Mammals seem to have maintained a subpopulation of cells with such properties, coined tanycytes. Propagation of activity within the tanycyte population has been described and found to be ATP dependent (Frayling et al., 2011), optogenetic activation-induced propagation (Bolborea et al., 2020) as well as neural activity. Furthermore, studies have also begun highlighting the role of tanycytes in regulating the blood-brain barrier (Langlet et al., 2013) and the changes in brain permeability during sleep (Fultz et al., 2019). Our functional observations of ependymal cell activity showing "quiescence-coupled" ultra-slow oscillations observed only at night in conjunction with our anatomical observations of ependymal ensheathment of motor tracts and projections to vasculature suggests that a neuro-ependymal-vascular circuit exists in larval zebrafish and that, based on the above-cited literature, it is likely a conserved feature throughout much of the animal kingdom. This opens the possibility of studying this circuit in larval zebrafish which compared to other model organism is more optically accessible, facilitating the study of such a distributed circuit.

6 Conclusion

*Progress in science depends on new techniques, new discoveries and new ideas,
probably in that order.*
– Sydney Brenner

A long-standing aim of our species has been to understand what drives our own and others behavior, to learn how to predict our seemingly erratic choices, to learn how to sculpt and control them, ideally to reach better outcomes. The Ancient Greeks looked to the humors within bodily fluids to account for our wavering mental state; in the 19th century, the study of behavior became relegated to psychologists, to which neuroscientists joined them later on. In the last decades, there has been a resurgence in our awareness of the influence of bodily state in driving behavior and a rise in the number of studies dissecting the pathways connecting internal processes to the brain, the ultimate driver of external actions.

One persistent difficulty in studying physiology-driven behavior has been the challenge of simultaneously recording activity from cells across the body within a living organism. This Ph.D. has been an attempt at combining and adapting existing tools, outlined in [Section 1](#), to establish the larval zebrafish as one platform to study physiology-driven decision making. To this end, [Section 2](#) sets out the basic building blocks allowing for whole-body imaging (WBI) including transgenic lines, choice of microscope and reference channel, and basic data pre-processing pipeline to get from raw images to cellular resolution activity traces, along with means to efficiently annotate data with tissue identity. [Section 3](#) addresses two needs that arose from [Section 2](#): the need to a) have a strong anatomical prior to better interpret WBI data, and b) be able to query molecular expression profiles to validate and test WBI generated hypothesis. For this we develop a whole-body expansion microscopy pipeline compatible with IHC or FISH across the age range of interest (5-14 d.p.f.). In [Section 4](#) we consider the analysis of such data and propose the use of alternative metrics, such as coherence and non-time reversibility, to better capture features of WBI data to account for the multitude of time-scales, common occurrence of traveling waves, and the fact that dimensions of high-variance do not necessarily capture the most interesting dynamics. We derive models that use these metrics to guide dimensionality reduction and clustering. In the last chapter, [Section 5](#), we bring together the three preceding chapters and present a few illustrative examples of how only by combining them one can get the most out of the WBI platform. The last element that we only briefly touch upon in the thesis, but that is as critical is the inclusion of perturbation methods. As such, large datasets generated through WBI can be processed into a manageable and meaningful format ready for user analysis. We illustrate how the breadth of tissues imaged and their identifiability allows for a less biased screening of cellular players involved in a process of interest ([Section 5.1](#)) and how it can readily reveal unforeseen dynamics and surprising couplings between tissue types ([Section 5.5](#)). We exemplify how analysis of such functional data can reveal structure in the dynamics which point to some underlying physical constraint within the system being studied, and finally how whole-body expansion microscopy allows one to directly search for the molecular underpinning of such constraints ([Fig. 21](#)).

The next immediate frontier will be to add a) a more extensive suite of closed-loop perturbation tools to the setup to enable precise stimulus delivery (e.g.: using caged compounds) and optogenetic control over various cell types, b) increasing the rate and field of view recorded (e.g.: dual-plane dual-color ortho-view mesoscope), c) diversifying the type of signals recorded through the use of other sensors (e.g.: glycolytic sensor, cAMP sensor, etc.), d) combining within sample functional and molecular ExM imaging (this is so far done independently), e) scaling the number of probes imaged within a simple ExM sample, f) building a whole-body

open-source ExM molecular Atlas to reference and register WBI data to, g) automating cell-type identification, e) building tailored models to infer longer time range couplings, and f) finally adapting and extending this setup to an adult model organism such as *Danionella Cerebrum*. These are all in reach and relatively easy to accomplish; yet, despite their simplicity, this would allow us to probe organism-scale vertebrate biology like never before. While none of these propositions are radical in themselves, they will hopefully enable radical new discoveries, ideas and breakthroughs in understanding.

7 References

- Ahrens, M. B., Orger, M. B., Robson, D. N., Li, J. M., and Keller, P. J. (2013). Whole-brain functional imaging at cellular resolution using light-sheet microscopy. *Nature Methods*, 10(5):413–420.
- Anderson, J. L., Marí, A., Braasch, I., Amores, A., Hohenlohe, P., Batzel, P., and Postlethwait, J. H. (2012). Multiple sex-associated regions and a putative sex chromosome in zebrafish revealed by RAD mapping and population genomics. *PLoS ONE*, 7(7).
- Arthur, D. and Vassilvitskii, S. (2007). K-means++: The advantages of careful seeding. *Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms*, 07-09-January-2007:1027–1035.
- Bagur, R. and Hajnóczky, G. (2017). Intracellular Ca²⁺ Sensing: Its Role in Calcium Homeostasis and Signaling. *Molecular Cell*, 66(6):780–788.
- Becker, C. G. and Becker, T. (2015). Neuronal Regeneration from Ependymo-Radial Glial Cells: Cook, Little Pot, Cook! *Developmental Cell*, 32(4):516–527.
- Billman, G. E. (2020). Homeostasis: The Underappreciated and Far Too Often Ignored Central Organizing Principle of Physiology. *Frontiers in Physiology*, 11(March):1–12.
- Bolborea, M., Pollatzek, E., Benford, H., Sotelo-Hitschfeld, T., and Dale, N. (2020). Hypothalamic tanycytes generate acute hyperphagia through activation of the arcuate neuronal network. *Proceedings of the National Academy of Sciences of the United States of America*, 117(25):14473–14481.
- Britz, R., Conway, K. W., and Rüber, L. (2021). The emerging vertebrate model species for neurophysiological studies is *Danionella cerebrum*, new species (Teleostei: Cyprinidae). *Scientific Reports*, 11(1):1–11.
- Buchman, T. G. (2003). The community of the self. *International Congress Series*, 1255(C):3–5.
- Buesing, L., Macke, J. H., and Sahani, M. (2012a). Learning stable, regularised latent models of neural population dynamics. *Network: Computation in Neural Systems*, 23:24–47.
- Buesing, L., Macke, J. H., and Sahani, M. (2012b). Spectral learning of linear dynamics from generalised-linear observations with application to neural population data. In *Advances in neural information processing systems*, pages 1682–1690.
- Castellani, J. W. and Young, A. J. (2016). Human physiological responses to cold exposure: Acute responses and acclimatization to prolonged exposure. *Autonomic Neuroscience: Basic and Clinical*, 196:63–74.
- Chen, F., Tillberg, P. W., and Boyden, E. S. (2015). Expansion microscopy. *Science*, 347(6221):543–548.
- Chen, J., Xia, L., Bruchas, M. R., and Solnica-Krezel, L. (2017). Imaging early embryonic calcium activity with GCaMP6s transgenic zebrafish. *Developmental biology*, 430(2):385.

- 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., Looger, L. L., Svoboda, K., and Kim, D. S. (2013). Nihms489004-1. *Nature*, 499(7458):295–300.
- Choi, T. Y., Choi, T. I., Lee, Y. R., Choe, S. K., and Kim, C. H. (2021). Zebrafish as an animal model for biomedical research. *Experimental and Molecular Medicine*, 53(3):310–317.
- Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Foster, J. D., Nuyujukian, P., Ryu, S. I., and Shenoy, K. V. (2012). Neural population dynamics during reaching. *Nature*, 487(7405):51–56.
- Clapham, D. E. (2007). Calcium Signaling. *Cell*, 131(6):1047–1058.
- Clark, D., Livezey, J., and Bouchard, K. (2019). Unsupervised discovery of temporal structure in noisy data with dynamical components analysis. In *Advances in Neural Information Processing Systems*, pages 14267–14278.
- Crane, D. K., Gockenbach, M. S., and Roberts, M. J. (2020). Approximating the singular value expansion of a compact operator. *SIAM Journal on Numerical Analysis*, 58(2):1295–1318.
- Cutajar, K., Osborne, M., Cunningham, J., and Filippone, M. (2016). Preconditioning kernel matrices. In *International Conference on Machine Learning*, pages 2529–2538.
- Dale, R. M. and Topczewski, J. (2011). Identification of an evolutionarily conserved regulatory element of the zebrafish col2a1a gene. *Developmental Biology*, 357(2):518–531.
- Dana, H., Mohar, B., Sun, Y., Narayan, S., Gordus, A., Hasseman, J. P., Tsegaye, G., Holt, G. T., Hu, A., Walpita, D., Patel, R., Macklin, J. J., Bargmann, C. I., Ahrens, M. B., Schreiter, E. R., Jayaraman, V., Looger, L. L., Svoboda, K., and Kim, D. S. (2016). Sensitive red protein calcium indicators for imaging neural activity. *eLife*, 5(MARCH2016):1–24.
- Dana, H., Sun, Y., Mohar, B., Hulse, B. K., Kerlin, A. M., Hasseman, J. P., Tsegaye, G., Tsang, A., Wong, A., Patel, R., Macklin, J. J., Chen, Y., Konnerth, A., Jayaraman, V., Looger, L. L., Schreiter, E. R., Svoboda, K., and Kim, D. S. (2019). High-performance calcium sensors for imaging activity in neuronal populations and microcompartments. *Nature Methods*, 16(7):649–657.
- Dantzer, R. (2018). Neuroimmune interactions: From the brain to the immune system and vice versa. *Physiological Reviews*, 98(1):477–504.
- Davies, P. I. and Higham, N. J. (2005). Computing $f(A)B$ for matrix functions f . In *QCD and numerical analysis III*, pages 15–24. Springer.
- Deisenroth, M. and Rasmussen, C. E. (2011). Pilco: A model-based and data-efficient approach to policy search. In *Proceedings of the 28th International Conference on machine learning (ICML-11)*, pages 465–472.
- Del Bigio, M. R. (2023). History of research concerning the ependyma: a view from inside the human brain. *Frontiers in Cellular Neuroscience*, 17(January):1–13.
- D’Gama, P. P., Qiu, T., Cosacak, M. I., Rayamajhi, D., Konac, A., Hansen, J. N., Ringers, C., Acuña-Hinrichsen, F., Hui, S. P., Olstad, E. W., Chong, Y. L., Lim, C. K. A., Gupta, A., Ng, C. P., Nilges, B. S., Kashikar, N. D., Wachten, D., Liebl, D., Kikuchi, K., Kizil, C., Yaksi, E., Roy, S., and Jurisch-Yaksi, N. (2021). Diversity and function of motile ciliated cell types within ependymal lineages of the zebrafish brain. *Cell Reports*, 37(1):109775.

- Ding, Y., Vanselow, D. J., Yakovlev, M. A., Katz, S. R., Lin, A. Y., Clark, D. P., Vargas, P., Xin, X., Copper, J. E., Canfield, V. A., Ang, K. C., Wang, Y., Xiao, X., Carlo, F. D., Rossum, D. B., Riviere, P. L., and Cheng, K. C. (2019). Computational 3d histological phenotyping of whole zebrafish by x-ray histotomography. *eLife*, 8:1–28.
- Dong, K., Eriksson, D., Nickisch, H., Bindel, D., and Wilson, A. G. (2017). Scalable log determinants for Gaussian process kernel learning. In *Advances in Neural Information Processing Systems*, pages 6327–6337.
- Duncker, L., Bohner, G., Boussard, J., and Sahani, M. (2019). Learning interpretable continuous-time models of latent stochastic dynamical systems. *arXiv preprint arXiv:1902.04420*.
- Fernandes, E. S., Russell, F. A., Alawi, K. M., Sand, C., Liang, L., Salamon, R., Bodkin, J. V., Aubdool, A. A., Arno, M., Gentry, C., Smillie, S. J., Bevan, S., Keeble, J. E., Malcangio, M., and Brain, S. D. (2016). Environmental cold exposure increases blood flow and affects pain sensitivity in the knee joints of CFA-induced arthritic mice in a TRPA1-dependent manner. *Arthritis Research and Therapy*, 18(1):1–12.
- Filippone, M. and Engler, R. (2015). Enabling scalable stochastic gradient-based inference for gaussian processes by employing the Unbiased LLinear System SolvEr (ULISSE). *arXiv preprint arXiv:1501.05427*.
- Frayling, C., Britton, R., and Dale, N. (2011). ATP-mediated glucosensing by hypothalamic tanycytes. *Journal of Physiology*, 589(9):2275–2286.
- Fultz, N. E., Bonmassar, G., Setsompop, K., Stickgold, R. A., Rosen, B. R., Polimeni, J. R., and Lewis, L. D. (2019). Coupled electrophysiological, hemodynamic, and cerebrospinal fluid oscillations in human sleep. *Science*, 366(6465):628–631.
- Giles, M. B. (2008). Collected matrix derivative results for forward and reverse mode algorithmic differentiation. In *Advances in Automatic Differentiation*, pages 35–44. Springer.
- Gnesotto, F., Mura, F., Gladrow, J., and Broedersz, C. (2018). Broken detailed balance and non-equilibrium dynamics in living systems: a review. *Reports on Progress in Physics*, 81(6):066601.
- Heywood, H. (1931). On finite sequences of real numbers. *Proc. Royal Soc. London Series A*, 134:486–501.
- Horn, B. K. and Schunck, B. G. (1981). Determining Optical Flow.
- Horn, R. A. and Johnson, C. R. (2012). *Matrix analysis*. Cambridge university press.
- Horstick, E. J., Jordan, D. C., Bergeron, S. A., Tabor, K. M., Serpe, M., Feldman, B., and Burgess, H. A. (2015). Increased functional protein expression using nucleotide sequence features enriched in highly expressed genes in zebrafish. *Nucleic Acids Research*, 43(7):1–14.
- Hu, F., Song, X., and Long, D. (2021). Transient receptor potential ankyrin 1 and calcium: Interactions and association with disease (Review). *Experimental and Therapeutic Medicine*, 22(6):1–7.
- Izumi, T., Imai, J., Yamamoto, J., Kawana, Y., Endo, A., Sugawara, H., Kohata, M., Asai, Y., Takahashi, K., Kodama, S., Kaneko, K., Gao, J., Uno, K., Sawada, S., Kalinichenko, V. V., Ishigaki, Y., Yamada, T., and Katagiri, H. (2018). Vagus-macrophage-hepatocyte

- link promotes post-injury liver regeneration and whole-body survival through hepatic FoxM1 activation. *Nature Communications*, 9(1):1–13.
- Kawakami, K. (2005). Transposon tools and methods in zebrafish. *Developmental Dynamics*, 234(2):244–254.
- Keller, P. J., Schmidt, A. D., Wittbrodt, J., and Stelzer, E. H. (2008). Reconstruction of zebrafish early embryonic development by scanned light sheet microscopy. *Science*, 322(5904):1065–1069.
- 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.e5.
- Lam, P. Y., Mangos, S., Green, J. M., Reiser, J., and Huttenlocher, A. (2015). In vivo imaging and characterization of actin microridges. *PLoS ONE*, 10(1):1–18.
- Langlet, F., Mullier, A., Bouret, S. G., Prevot, V., and Dehouck, B. (2013). Tanycyte-like cells form a blood-cerebrospinal fluid barrier in the circumventricular organs of the mouse brain. *Journal of Comparative Neurology*, 521(15):3389–3405.
- Li, F., Long, Y., Xie, J., Ren, J., Zhou, T., Song, G., Li, Q., and Cui, Z. (2021a). Generation of gcamp6s-expressing zebrafish to monitor spatiotemporal dynamics of calcium signaling elicited by heat stress. *International Journal of Molecular Sciences*, 22(11).
- Li, Y., Jia, Z., Zhang, S., and He, X. (2021b). Progress in gene-editing technology of Zebrafish. *Biomolecules*, 11(9).
- Linderman, S., Johnson, M., Miller, A., Adams, R., Blei, D., and Paninski, L. (2017). Bayesian learning and inference in recurrent switching linear dynamical systems. In *Artificial Intelligence and Statistics*, pages 914–922.
- Lorincz, R., Emfinger, C. H., Walcher, A., Giolai, M., Krautgasser, C., Remedi, M. S., Nichols, C. G., and Meyer, D. (2018). In vivo monitoring of intracellular Ca²⁺ dynamics in the pancreatic β -cells of zebrafish embryos. *Islets*, 10(6):221–238.
- Macke, J. H., Büsing, L., Cunningham, J. P., Yu, B. M., Shenoy, K. V., and Sahani, M. (2011). Empirical models of spiking in neural populations. In Shawe-Taylor, J., Zemel, R. S., Bartlett, P., Pereira, F. C. N., and Weinberger, K. Q., editors, *Advances in Neural Information Processing Systems*, volume 24, pages 1350–1358.
- Maclaurin, D., Duvenaud, D., and Adams, R. P. (2015). Autograd: Effortless gradients in numpy. In *ICML 2015 AutoML Workshop*, volume 238.
- Martin, J. K. and McDonald, R. P. (1975). Bayesian estimation in unrestricted factor analysis: A treatment for Heywood cases. *Psychometrika*, 40:505–517.
- McDole, K., Guignard, L., Amat, F., Berger, A., Malandain, G., Royer, L. A., Turaga, S. C., Branson, K., and Keller, P. J. (2018). In Toto Imaging and Reconstruction of Post-Implantation Mouse Development at the Single-Cell Level. *Cell*, 175(3):859–876.e33.
- Miyawaki, A., Llopis, J., Heim, R., Michael McCaffery, J., Adams, J. A., Ikura, M., and Tsien, R. Y. (1997). Fluorescent indicators for Ca²⁺ based on green fluorescent proteins and calmodulin. *Nature*, 388(6645):882–887.

- Moore, J. L., Bhaskar, D., Gao, F., Matte-Martone, C., Du, S., Lathrop, E., Ganesan, S., Shao, L., Norris, R., Sanz, N. C., Annusver, K., Kasper, M., Cox, A., Hendry, C., Rieck, B., Krishnaswamy, S., and Greco, V. (2023). Cell cycle controls long-range calcium signaling in the regenerating epidermis. *Journal of Cell Biology*, 222(7).
- Mosimann, C., Kaufman, C. K., Li, P., Pugach, E. K., Tamplin, O. J., and Zon, L. I. (2011). Ubiquitous transgene expression and Cre-based recombination driven by the ubiquitin promoter in zebrafish. *Development*, 138(1):169–177.
- M'Saad, O., Shribak, M., and Bewersdorf, J. (2022). Unclearing Microscopy. *bioRxiv*, page 2022.11.29.518361.
- Mu, Y., Bennett, D. V., Rubinov, M., Narayan, S., Yang, C. T., Tanimoto, M., Mensh, B. D., Looger, L. L., and Ahrens, M. B. (2019). Glia Accumulate Evidence that Actions Are Futile and Suppress Unsuccessful Behavior. *Cell*, 178(1):27–43.e19.
- Niezink, N. M. D. and Van Der Schaft, A. J. (2011). Consensus in networked multi-agent systems. (August).
- Nummenmaa, E., Hämäläinen, M., Pemmari, A., Moilanen, L. J., Tuure, L., Nieminen, R. M., Moilanen, T., Vuolteenaho, K., and Moilanen, E. (2021). Transient receptor potential ankyrin 1 (TRPA1) is involved in upregulating interleukin-6 expression in osteoarthritic chondrocyte models. *International Journal of Molecular Sciences*, 22(1):1–14.
- Pandarínath, C., O'Shea, D. J., Collins, J., Jozefowicz, R., Stavisky, S. D., Kao, J. C., Trautmann, E. M., Kaufman, M. T., Ryu, S. I., Hochberg, L. R., et al. (2018). Inferring single-trial neural population dynamics using sequential auto-encoders. *Nature methods*, 15(10):805–815.
- Parra, G. and Tobar, F. (2017a). Spectral mixture kernels for multi-output Gaussian processes. In *Advances in Neural Information Processing Systems*, pages 6681–6690.
- Parra, G. and Tobar, F. (2017b). Spectral mixture kernels for multi-output Gaussian processes. In *Advances in Neural Information Processing Systems*, pages 6681–6690.
- Petreska, B., Byron, M. Y., Cunningham, J. P., Santhanam, G., Ryu, S. I., Shenoy, K. V., and Sahani, M. (2011). Dynamical segmentation of single trials from population neural data. In *Advances in neural information processing systems*, pages 756–764.
- Pozo-Morales, M., Garteizgogeoasca, I., Perazzolo, C., So, J., Shin, D., and Singh, S. P. (2023). In vivo imaging of calcium dynamics in zebrafish hepatocytes. *Hepatology*, 77(3):789–801.
- Randlett, O., Wee, C. L., Naumann, E. A., Nnaemeka, O., Schoppik, D., Fitzgerald, J. E., Portugues, R., Lacoste, A. M., Riegler, C., Engert, F., and Schier, A. F. (2015). Whole-brain activity mapping onto a zebrafish brain atlas. *Nature Methods*, 12(11):1039–1046.
- Rasmussen, C. E. and Williams, C. K. I. (2006). *Gaussian Processes for Machine Learning*. MIT press.
- Reimer, M. M., Kuscha, V., Wyatt, C., Sörensen, I., Frank, R. E., Knüwer, M., Becker, T., and Becker, C. G. (2009). Sonic hedgehog is a polarized signal for motor neuron regeneration in adult zebrafish. *Journal of Neuroscience*, 29(48):15073–15082.
- Reimer, M. M., Sörensen, I., Kuscha, V., Frank, R. E., Liu, C., Becker, C. G., and Becker, T. (2008). Motor neuron regeneration in adult zebrafish. *Journal of Neuroscience*, 28(34):8510–8516.

- Roosta-Khorasani, F. and Ascher, U. (2015). Improved bounds on sample size for implicit matrix trace estimators. *Foundations of Computational Mathematics*, 15:1187–1212.
- Roper, S. D. (2013). Taste buds as peripheral chemosensory processors. *Seminars in Cell and Developmental Biology*, 24(1):71–79.
- Rosas-Ballina, M., Olofsson, P. S., Ochani, M., Valdés-Ferrer, S. I., Levine, Y. A., Reardon, C., Tusche, M. W., Pavlov, V. A., Andersson, U., Chavan, S., Mak, T. W., and Tracey, K. J. (2011). Acetylcholine-synthesizing T cells relay neural signals in a vagus nerve circuit. *Science*, 334(6052):98–101.
- Roweis, S. and Ghahramani, Z. (1999). A unifying review of linear Gaussian models. *Neural computation*, 11:305–345.
- Russo, A. A., Bittner, S. R., Perkins, S. M., Seely, J. S., London, B. M., Lara, A. H., Miri, A., Marshall, N. J., Kohn, A., Jessell, T. M., et al. (2018). Motor cortex embeds muscle-like commands in an untangled population response. *Neuron*, 97:953–966.
- Rutten, V., Bernacchia, A., and Hennequin, G. (2020). Sequential components analysis. In *Cosyne, Denver, CO. T-22*.
- Shipley, F. B., Dani, N., Xu, H., Deister, C., Cui, J., Head, J. P., Sadegh, C., Fame, R. M., Shannon, M. L., Flores, V. I., Kishkovich, T., Jang, E., Klein, E. M., Goldey, G. J., He, K., Zhang, Y., Holtzman, M. J., Kirchhausen, T., Wyart, C., Moore, C. I., Andermann, M. L., and Lehtinen, M. K. (2020). Tracking Calcium Dynamics and Immune Surveillance at the Choroid Plexus Blood-Cerebrospinal Fluid Interface. *Neuron*, 108(4):623–639.e10.
- Sim, J., Park, C. E., Cho, I., Min, K., Eom, M., Han, S., Jeon, H., Cho, H.-J., Cho, E.-S., Kumar, A., Chong, Y., Kang, J. S., Piatkevich, K. D., Jung, E. E., Kang, D.-S., Kwon, S.-K., Kim, J., Yoon, K.-J., Lee, J.-S., Boyden, E. S., Yoon, Y.-G., Chang, J.-B., and Paulson, J. A. (2022). Nanoscale resolution imaging of the whole mouse embryos and larval zebrafish using expansion microscopy. *bioRxiv*, page 2021.05.18.443629.
- Sofroniew, N. J., Flickinger, D., King, J., and Svoboda, K. (2016). A large field of view two-photon mesoscope with subcellular resolution for in vivo imaging. *eLife*, 5(JUN2016):1–20.
- Steib, E., Tetley, R., Laine, R. F., Norris, D. P., Mao, Y., and Vermot, J. (2022). TissUExM enables quantitative ultrastructural analysis in whole vertebrate embryos by expansion microscopy. *Cell Reports Methods*, 2(10):100311.
- Stringer, C. and Pachitariu, M. (2019). Computational processing of neural recordings from calcium imaging data. *Current Opinion in Neurobiology*, 55:22–31.
- Szebényi, K., Füredi, A., Kolacsek, O., Csohány, R., Prókai, Á., Kis-Petik, K., Szabó, A., Bősze, Z., Bender, B., Tóvári, J., et al. (2015). Visualization of calcium dynamics in kidney proximal tubules. *Journal of the American Society of Nephrology*, 26(11):2731–2740.
- Tabor, K. M., Marquart, G. D., Hurt, C., Smith, T. S., Geoca, A. K., Bhandiwad, A. A., Subedi, A., Sinclair, J. L., Rose, H. M., Polys, N. F., and Burgess, H. A. (2019). Brain-wide cellular resolution imaging of cre transgenic zebrafish lines for functional circuit-mapping. *eLife*, 8:1–21.
- Van Loan, C. F. (2000). The ubiquitous kronecker product. *Journal of computational and applied mathematics*, 123(1-2):85–100.

- Varadarajan, S., Chumki, S. A., Stephenson, R. E., Misterovich, E. R., Wu, J. L., Dudley, C. E., Erofeev, I. S., Goryachev, A. B., and Miller, A. L. (2022). Mechanosensitive calcium flashes promote sustained RhoA activation during tight junction remodeling. *Journal of Cell Biology*, 221(4).
- Von Luxburg, U. (2007). A tutorial on spectral clustering. *Statistics and Computing*, 17(4):395–416.
- Wang, Y., Eddison, M., Fleishman, G., Weigert, M., Xu, S., Wang, T., Rokicki, K., Goina, C., Henry, F. E., Lemire, A. L., Schmidt, U., Yang, H., Svoboda, K., Myers, E. W., Saalfeld, S., Korff, W., Sternson, S. M., and Tillberg, P. W. (2021). EASI-FISH for thick tissue defines lateral hypothalamus spatio-molecular organization. *Cell*, 184(26):6361–6377.e24.
- White, R. M., Sessa, A., Burke, C., Bowman, T., LeBlanc, J., Ceol, C., Bourque, C., Dovey, M., Goessling, W., Burns, C. E., and Zon, L. I. (2008). Transparent Adult Zebrafish as a Tool for In Vivo Transplantation Analysis. *Cell Stem Cell*, 2(2):183–189.
- Wilson, A. and Adams, R. (2013a). Gaussian process kernels for pattern discovery and extrapolation. In *International conference on machine learning*, pages 1067–1075.
- Wilson, A. and Adams, R. (2013b). Gaussian process kernels for pattern discovery and extrapolation. In *International conference on machine learning*, pages 1067–1075.
- Wilson, A. and Nickisch, H. (2015). Kernel interpolation for scalable structured Gaussian processes (KISS-GP). In *International Conference on Machine Learning*, pages 1775–1784.
- Wingert, R. A., Selleck, R., Yu, J., Song, H. D., Chen, Z., Song, A., Zhou, Y., Thisse, B., Thisse, C., McMahon, A. P., and Davidson, A. J. (2007). The cdx genes and retinoic acid control the positioning and segmentation of the zebrafish pronephros. *PLoS Genetics*, 3(10):1922–1938.
- Wouk, A. (1965). Integral representation of the logarithm of matrices and operators. *Journal of Mathematical Analysis and Applications*, 11:131–138.
- Wu, Y., Ghitani, A., Christensen, R., Santella, A., Du, Z., Rondeau, G., Bao, Z., Colón-Ramos, D., and Shroff, H. (2011). Inverted selective plane illumination microscopy (iSPIM) enables coupled cell identity lineaging and neurodevelopmental imaging in *Caenorhabditis elegans*. *Proceedings of the National Academy of Sciences of the United States of America*, 108(43):17708–17713.
- Yang, Y. H. C., Kawakami, K., and Stainier, D. Y. (2018). A new mode of pancreatic islet innervation revealed by live imaging in zebrafish. *eLife*, 7:1–19.
- Yi-Wen, L. (2007). Interrenal Organogenesis in the Zebrafish Model. *Organogenesis*, 3(1):44–48.
- Yoshinari, N., Ando, K., Kudo, A., Kinoshita, M., and Kawakami, A. (2012). Colored medaka and zebrafish: Transgenics with ubiquitous and strong transgene expression driven by the medaka β -actin promoter. *Development Growth and Differentiation*, 54(9):818–828.
- Yu, B. M., Cunningham, J. P., Santhanam, G., Ryu, S. I., Shenoy, K. V., and Sahani, M. (2009). Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. In *Advances in neural information processing systems*, pages 1881–1888.

- Zaghloul, M. R. and Ali, A. N. (2012). Algorithm 916: computing the Faddeyeva and Voigt functions. *ACM Transactions on Mathematical Software (TOMS)*, 38(2):1–22.
- Zhang, X., Lei, B., Yuan, Y., Zhang, L., Hu, L., Jin, S., Kang, B., Liao, X., Sun, W., Xu, F., Zhong, Y., Hu, J., and Qi, H. (2020). Brain control of humoral immune responses amenable to behavioural modulation. *Nature*, 581(7807):204–208.

8 Appendix

Appendix — Non-reversible Gaussian processes for identifying latent dynamical structure in neural data

A Boundedness of the reversibility index

In this section, we prove that for a stationary multi-output GP, the reversibility index defined in [Eq. 65](#) is bounded between 0 and 1. Recall the definition of the (squared) reversibility index:

$$\zeta^2 = \frac{\int_{-\infty}^{\infty} \|K(\tau) - K(-\tau)\|_{\text{F}}^2 d\tau}{\int_{-\infty}^{\infty} \|K(\tau) + K(-\tau)\|_{\text{F}}^2 d\tau}. \quad (77)$$

Since $\|X\|_{\text{F}}^2 = \text{Tr}(XX^\top)$, we can expand this into:

$$\zeta^2 = \frac{\text{Tr} \int_{-\infty}^{\infty} K(\tau)K(\tau)^\top d\tau - \text{Tr} \int_{-\infty}^{\infty} K(\tau)K(\tau) d\tau}{\text{Tr} \int_{-\infty}^{\infty} K(\tau)K(\tau)^\top d\tau + \text{Tr} \int_{-\infty}^{\infty} K(\tau)K(\tau) d\tau}. \quad (78)$$

The first term in both the numerator and the denominator is non-negative because the integral of outer products $K(\tau)K(\tau)^\top$ is positive semi definite (PSD). We are left to show that the second term is non-negative too, which would then imply $0 \leq \zeta \leq 1$. From Parseval's theorem, we have that:

$$\int_{-\infty}^{\infty} K(\tau)K(\tau) d\tau = \int_{-\infty}^{\infty} \widehat{K}(\omega)\overline{\widehat{K}(\omega)} d\omega \quad (79)$$

where $\widehat{\cdot}$ denotes the Fourier transform (assumed to exist) and $\bar{\cdot}$ denotes the complex conjugate. Furthermore, from Cramér's theorem, since $K(\cdot)$ is a stationary cross-covariance function, $\widehat{K}(\omega)$ is Hermitian PSD, and so is $\widehat{\widehat{K}}(\omega)$, for any ω . Finally, for any two Hermitian PSD matrices, A and B, it can be shown that $\text{Tr}(AB) \geq 0$ (Theorem 4.3.53 in [Horn and Johnson, 2012](#)). Thus, we have shown that:

$$\text{Tr} \int_{-\infty}^{\infty} K(\tau)K(\tau)^\top d\tau \geq 0 \quad (80)$$

$$\text{Tr} \int_{-\infty}^{\infty} K(\tau)K(\tau) d\tau \geq 0 \quad (81)$$

from which we can conclude that: $0 \leq \zeta \leq 1$.

B General Kronecker decomposition of stationary multi-output covariances

Here, we prove the existence of the decomposition of [Eq. 66](#). Let $K(\tau)$ be a matrix-valued (i.e. multi-output) cross-covariance function; we assume that each scalar cross-covariance function $K_{ij}(\cdot)$ is in L^2 . From the singular value decomposition theorem for compact operators in Hilbert spaces [Crane et al., 2020](#), there exist a basis set of M^2 functions $\{f_\ell(\cdot) \in L^2\}_{\ell=1}^{M^2}$, a matching set

of matrices $\{A_\ell \in \mathbb{R}^{M \times M}\}_{\ell=1}^{M^2}$, and a corresponding set of positive singular values $\{\lambda_\ell \in \mathbb{R}^+\}_{\ell=1}^{M^2}$ such that:

$$K(\tau) = \sum_{\ell=1}^{M^2} \lambda_\ell A_\ell \cdot f_\ell(\tau) \quad (82)$$

with the following orthonormality conditions:

$$\text{Tr}(A_\ell A_{\ell'}^\top) = \delta_{\ell\ell'} \quad (83)$$

$$\int f_\ell(\tau) f_{\ell'}(\tau) d\tau = \delta_{\ell\ell'}. \quad (84)$$

We will refer to Eq. 82 as a “generalized SVD”. We now show that the pairs $\{A_\ell, f_\ell\}$ are either symmetric/even, or skew-symmetric/odd, as stated in Section 4.4.5. Let $\tilde{\mathbf{k}}(\tau) = \text{vec}(K(\tau)) \in \mathbb{R}^{M^2}$ be the vectorized cross-covariance matrix at time lag τ . Let $P \in \{0, 1\}^{M^2 \times M^2}$ be the commutation matrix operating on the vectorized space of all M^2 pairs of outputs (i.e. the space of $\tilde{\mathbf{k}}$), which swaps index $(i + jM)$ with index $(j + iM)$. In other words, for any matrix $X \in \mathbb{R}^{M \times M}$, we have that $P \text{vec}(X) = \text{vec}(X^\top)$. Let $\mathcal{R}$ be the reflection (or ‘time-reversal’) operator in L^2 , such that $\mathcal{R}[f](\tau) = f(-\tau)$. It is easy to show that any multi-output cross-covariance function $K(\tau)$ obeys the symmetry $K(-\tau) = K(\tau)^\top$, which can also be written as:

$$\mathcal{R}[\tilde{\mathbf{k}}](\tau) = P\tilde{\mathbf{k}}(\tau) \quad (85)$$

where $\mathcal{R}[\cdot]$ is applied elementwise to the element functions in $\tilde{\mathbf{k}}$. In other words, reversing time and transposing space are two equivalent operations. The generalized SVD in Eq. 82 can be written in vectorized form as

$$\tilde{\mathbf{k}}(\tau) = \underbrace{(\tilde{\mathbf{a}}_1, \dots, \tilde{\mathbf{a}}_{M^2})}_U \underbrace{\text{diag}(\lambda_1, \dots, \lambda_{M^2})}_S \underbrace{(f_1(\tau), \dots, f_{M^2}(\tau))^\top}_{v(\tau)} \quad (86)$$

where $\tilde{\mathbf{a}}_\ell = \text{vec}(A_\ell)$. Thus, the symmetry of Eq. 85 can be re-expressed as

$$(PU)S(\mathcal{R}[\tilde{\mathbf{v}}](\tau)) = US\tilde{\mathbf{v}}(\tau). \quad (87)$$

Since the two permutation operators P and $\mathcal{R}$ preserve orthonormality (of matrices and square-integrable functions, respectively), the l.h.s. of Eq. 87 is a valid SVD for $K(\tau)$. Furthermore, when the singular values are distinct and kept in decreasing order, the generalized SVD of $\tilde{\mathbf{k}}(\tau)$ is unique up to a *simultaneous* change of sign in any pair $(\tilde{\mathbf{a}}_\ell, f_\ell(\cdot))$. Therefore, we must have that

$$PU = U\pm \quad (88)$$

$$\mathcal{R}[\tilde{\mathbf{v}}](\tau) = \pm\tilde{\mathbf{v}}(\tau) \quad (89)$$

where $\pm$ is a (shared) diagonal matrix composed of $+1$ and -1 elements only. Importantly, the matrix $\pm$ is the same in the two equations. From the spatial transpose meaning of P and time-reversal meaning of $\mathcal{R}$, Eqs. 88 and 89 therefore imply that for any ℓ :

- if $\pm_\ell = +1$, then $\tilde{\mathbf{a}}_\ell$ is invariant to the transposition operator P and $f_\ell(\cdot)$ is invariant to time reversal, implying that A_ℓ is a symmetric matrix and f_ℓ is an even function,
- if $\pm_\ell = -1$, then both $\tilde{\mathbf{a}}_\ell$ and f_ℓ have their sign flipped by spatial transposition and time reversal, respectively, implying that A_ℓ is a skew-symmetric matrix and f_ℓ is an odd function.

The decomposition of Eq. 66 follows from a simple, now justified renaming of symmetric/even and skew-symmetric/odd pairs as $\{A_\ell^\pm, f_\ell^\pm\}$. Moreover, we now see that the $+$ (resp. $-$) terms in Eq. 86 correspond to the generalized SVD of the symmetric (resp. antisymmetric) part of the covariance function, $K(\tau) + K(-\tau)$ (resp. $K(\tau) - K(-\tau)$) in Eq. 77. Given the known relationship between the squared Frobenius norm of a linear operator and the sum of its squared singular values, this shows that the non-reversibility index ζ can also be calculated using Eq. 67.

C Construction of valid non-reversible planar kernels

Here, we prove that Eq. 68 is a valid 2-output covariance function. We focus the proof on the purely spherical case, $\sigma_1 = \sigma_2 = 1$ and $\rho = 0$ – extension to arbitrary instantaneous covariances is straightforward but notationally cumbersome (and a more general proof was in fact already given in Section 4.4.7).

In the frequency domain, the Hilbert transform has a simple interpretation as a constant phase shift of $\pi/2$ at all frequencies; specifically, $\widehat{\mathcal{H}[f]}(\omega) = -j \cdot \text{sgn}(\omega) \widehat{f}(\omega)$, where $\widehat{\cdot}$ denotes the Fourier transform and $j^2 = -1$. As any scalar covariance function, f is even, and therefore $\widehat{f}$ is real and even. In contrast, $\widehat{\mathcal{H}[f]}$ is imaginary and odd. We now derive the eigendecomposition of $K(\cdot)$, and show that all eigenvalues are positive. Due to the hybrid nature of $K(\cdot)$ (discrete space, continuous time), an eigenfunction of $K(\cdot)$ is a time-varying 2-dimensional “vector” $[g_1(t), g_2(t)]^\top$ which satisfies

$$\int_{-\infty}^{\infty} K(\tau - t) \begin{bmatrix} g_1(\tau) \\ g_2(\tau) \end{bmatrix} d\tau = \lambda \begin{bmatrix} g_1(t) \\ g_2(t) \end{bmatrix} \quad (90)$$

or equivalently,

$$\int_{-\infty}^{\infty} \begin{bmatrix} f(\tau)g_1(\tau + t) + \alpha \mathcal{H}[f](\tau)g_2(\tau + t) \\ f(\tau)g_2(\tau + t) - \alpha \mathcal{H}[f](\tau)g_1(\tau + t) \end{bmatrix} d\tau = \lambda \begin{bmatrix} g_1(t) \\ g_2(t) \end{bmatrix}. \quad (91)$$

Knowing that the eigenfunctions of any stationary scalar kernel are the Fourier modes, $e^{j\omega t}$, we make the following ansatz for the eigenfunctions of $K(\cdot)$:

$$\mathbf{g}_\omega^\pm(t) = \begin{bmatrix} e^{j\omega t} \\ b^\pm e^{\pm j\omega t} \end{bmatrix} \quad \text{for some } b^\pm \in \mathbb{C}, \quad (92)$$

(with the understanding that the two $\pm$ signs are “tied”, i.e. they are either both $+$ or both $-$). Next, we note that for any scalar function h and $\omega \in \mathbb{R}$,

$$\int_{-\infty}^{\infty} h(\tau) e^{\pm j\omega(t+\tau)} d\tau = e^{\pm j\omega t} \int_{-\infty}^{\infty} h(\tau) e^{\pm j\omega\tau} d\tau \quad (93)$$

$$= e^{\pm j\omega t} \int_{-\infty}^{\infty} h(\tau) e^{-j(\mp\omega)\tau} d\tau \quad (94)$$

$$= e^{\pm j\omega t} \widehat{h}(\mp\omega) \quad (95)$$

where the last equality follows from the definition of the Fourier transform $\widehat{h}(\cdot)$. Thus, for $\mathbf{g}_\omega^\pm(t)$ to be an eigenfunction of $K(\cdot)$ with eigenvalue $\lambda_\omega^\pm$, $b^\pm$ must satisfy:

$$\widehat{f}(\omega) + b^\pm \alpha \widehat{\mathcal{H}[f]}(\mp\omega) = \lambda_\omega^\pm \quad (96)$$

$$-\alpha \widehat{\mathcal{H}[f]}(\mp\omega) + b^\pm \widehat{f}(\omega) = b^\pm \lambda_\omega^\pm \quad (97)$$

(easily obtained from Eq. 91 and some straightforward algebra). Inserting $\widehat{\mathcal{H}[f]}(\omega) = -j \cdot \text{sgn}(\omega) \widehat{f}(\omega)$, the above system of equations imply:

$$\begin{aligned} \text{either } b^\pm = +j &\rightarrow \lambda_\omega^\pm = (1 - \text{sgn}(\pm\omega)\alpha) \widehat{f}(\omega) \\ \text{or } b^\pm = -j &\rightarrow \lambda_\omega^\pm = (1 + \text{sgn}(\pm\omega)\alpha) \widehat{f}(\omega). \end{aligned} \quad (98)$$

Since f is itself a valid kernel whose real eigenvalues $\widehat{f}(\omega)$ are all strictly positive, we conclude that $K(\cdot)$ is a valid 2-output covariance function (i.e. all $\lambda_\omega^\pm \geq 0$) if, and only if, $|\alpha| \leq 1$.

This derivation also provides a key connection between our measure of second-order non-reversibility, ζ , and one’s intuitive understanding that a maximally non-reversible process should never “turn back on itself”. Specifically, we note that when $|\alpha| = 1$, Eq. 98 implies that half of the eigenvalues of $K(\cdot)$ are exactly zero. By inspection of the corresponding eigenfunctions in Eq. 91, we see that for any sample trajectory $(x_1(t), x_2(t))^\top$ from $K(\cdot)$, which must lie in the span of the eigenfunctions with non-zero eigenvalues, the time-reversed trajectory $(x_1(-t), x_2(-t))$ evolves in the span of the eigenfunctions with zero eigenvalues. Therefore, these time-reversed trajectories have zero probability density under our non-reversible prior when $|\alpha| = 1$; this extreme case corresponds to $\zeta = 1$, i.e. maximal non-reversibility (c.f. Eq. 69).

D Derivatives of special functions

Optimization of the hyperparameters requires evaluating the gradient of the marginal likelihood, which in turn requires gradients of all the functions involved in the multi-output kernel. We give details of the gradients for the functions listed in Table 1 and present them graphically in Fig. 22. These gradients are simple to evaluate numerically, enabling the addition of our non-reversible kernels to standard automatic differentiation systems.

Dawson function The Hilbert transform of the squared-exponential kernel involves the so-called Dawson function, given below along with its derivative:

$$D(\tau) \triangleq e^{-\tau^2} \int_0^\tau e^{s^2} ds \quad \frac{dD}{d\tau} = 1 - 2\tau D(\tau). \quad (99)$$

Exponential integral The Hilbert transform of the exponential kernel involves the exponential integral function:

$$\text{Ei}(\tau) \triangleq \int_{-\infty}^\tau \frac{e^s}{s} ds \quad \frac{d\text{Ei}}{d\tau} = \frac{e^\tau}{\tau}. \quad (100)$$

Faddeeva function The Hilbert transform of the spectral mixture kernel [Wilson and Adams, 2013b](#) involves the imaginary component of the Faddeeva function $w(\cdot)$ (related to the complex error function; [Zaghloul and Ali, 2012](#)). This is a new result that we derived which we were not able to find in the existing literature. The Faddeeva function is available in most numerical computing environments, and can be expressed as:

$$w(z) = V(x, y) + jL(x, y) \quad \text{with } z = x + jy \quad (101)$$




Figure 22: **Commonly used GP covariance functions and their Hilbert transforms.** Please refer to [Table 1](#) for details.

where $V(x, y)$ and $L(x, y)$ are the real and imaginary Voigt functions respectively. Gradients are given by:

$$\frac{\partial L(x, y)}{\partial x} = -\frac{\partial V(x, y)}{\partial y} = -2 \operatorname{Im}[zw(z)] + \frac{2}{\sqrt{\pi}} \quad (102)$$

$$\text{and} \quad \frac{\partial L(x, y)}{\partial y} = \frac{\partial V(x, y)}{\partial x} = -2 \operatorname{Re}[zw(z)]. \quad (103)$$

Evaluating these expressions only requires evaluating the Fadeeva function itself.

E Beyond non-reversible planar processes: construction of higher-dimensional non-reversible covariance functions

In [Section 4.4.8](#), we built non-reversible M-output GP covariances as superpositions of planar kernels, each associated with a (potentially) different scalar covariance function. Here, we introduce an alternative construction motivated by the following considerations of model degeneracies. In [Eq. 74](#), if (i) the marginal variances in each plane are all identical ($\sigma_{ij,1} = \sigma_{ij,2} = \sigma$), (ii) each planar process is spherical ($\rho_{ij} = 0$) and (iii) the scalar covariance functions $f_{ij}(\tau)$ are all the same f , then $K(\tau)$ is over-parametrized. Indeed, it can then be rewritten as

$$K(\tau) \propto \underbrace{I_M}_{A^+} f(\tau) + \underbrace{\begin{pmatrix} 0 & \alpha_{12} & \alpha_{13} & \cdots \\ -\alpha_{12} & 0 & \alpha_{23} & \cdots \\ -\alpha_{13} & -\alpha_{23} & 0 & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix}}_{A^-} \mathcal{H}[f](\tau) \quad (104)$$

Now, since A^- is antisymmetric, there exists a unitary transformation of the latent space in which this covariance becomes

$$K(\tau) \propto I_M f(\tau) + \begin{pmatrix} 0 & \omega_1 & 0 & 0 & \cdots \\ -\omega_1 & 0 & 0 & 0 & \cdots \\ 0 & 0 & 0 & \omega_2 & \cdots \\ 0 & 0 & -\omega_2 & 0 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix} \mathcal{H}[f](\tau) \quad (105)$$

where $\{\pm j\omega_1, \pm j\omega_2, \dots\}$ are the imaginary conjugate eigenvalues of A^- . In GPFADS, this unitary transformation can be absorbed in the mixing matrix C ([Eq. 60](#)). Thus, in

this configuration, our model does not truly possess $M(M-1)/2$ free parameters as the parametrization of Eq. 74 suggests, but only $M/2$, as Eq. 105 reveals. In other words, one can always rotate the latent space and directly parametrize a set of independent planes – in which case one must enforce $|\omega_i| < 1$ to ensure positive definiteness.

For these reasons, we also propose this alternative construction:

$$K(\tau) = \sum_{q=1}^Q U_q \left(\begin{bmatrix} A_{q1}^+ & 0 & \ddots \\ 0 & A_{q2}^+ & 0 \\ \ddots & 0 & \ddots \end{bmatrix} f_q(\tau) + \begin{bmatrix} \alpha_{q1} A_{q1}^- & 0 & \ddots \\ 0 & \alpha_{q2} A_{q2}^- & 0 \\ \ddots & 0 & \ddots \end{bmatrix} \mathcal{H}[f_q](\tau) \right) U_q^\top \quad (106)$$

with some scalar covariance functions $\{f_q(\cdot)\}$, unitary matrices $\{U_q\}$, non-reversibility parameters $|\alpha_{qi}| < 1$ associated with each of the $M/2$ planes, and 2×2 matrix blocks $\{A_{qi}^+, A_{qi}^-\}$. The latter are parameterized exactly as the symmetric and antisymmetric matrices A^+ and A^- in Eq. 68. We note that when such a kernel is used in GPFA, U_1 can be set to the identity matrix without loss of generality since it corresponds to a rotation of the latent space that can be absorbed in the mixing matrix C (Eq. 60).

F Implementation notes and scalability

F.1 Stable computation of the log marginal likelihood

The log marginal likelihood in GPFA(DS) can be computed in a stable way as follows. Recall its expression:

$$\mathcal{L}(\theta, Y) \propto -\log |K_{yy}| - (\tilde{\mathbf{y}} - \boldsymbol{\mu} \otimes \mathbf{1}_T)^\top K_{yy}^{-1} (\tilde{\mathbf{y}} - \boldsymbol{\mu} \otimes \mathbf{1}_T) \quad (107)$$

$$\text{with } K_{yy} = (C \otimes I_T) K_{xx} (C^\top \otimes I_T) + (R \otimes I_T) \quad (108)$$

where both K_{yy} and $\boldsymbol{\mu}$ depend on model parameters θ . A stable way to evaluate this is to begin with a Cholesky decomposition of $K_{xx} \in \mathbb{R}^{MT \times MT}$:

$$K_{xx} = LL^\top \quad (109)$$

and then use the Woodbury identity to rewrite the inverse of $K_{yy} \in \mathbb{R}^{NT \times NT}$ as:

$$K_{yy}^{-1} = (R^{-1} \otimes I_T) - (R^{-1}C \otimes I_T) (L^{-T}L^{-1} + C^\top R^{-1}C \otimes I_T)^{-1} (C^\top R^{-1} \otimes I_T), \quad (110)$$

which can be further transformed into:

$$K_{yy}^{-1} = (R^{-1/2} \otimes I_T) (I_N \otimes I_T - A^\top B^{-1}A) (R^{-1/2} \otimes I_T) \quad (\in \mathbb{R}^{NT \times NT}) \quad (111)$$

$$\text{where } A \equiv L^\top (C^\top R^{-1/2} \otimes I_T) \quad (\in \mathbb{R}^{MT \times NT}) \quad (112)$$

$$\text{and } B \equiv I_{DT} + AA^\top = I_{DT} + L^\top (C^\top R^{-1}C \otimes I_T) L \quad (\in \mathbb{R}^{MT \times MT}) \quad (113)$$

Using the matrix determinant lemma, $\log |K_{yy}^{-1}|$ can be simplified to:

$$\log |K_{yy}^{-1}| = \log |(R^{-1/2} \otimes I_T) [I_T \otimes I_N - A^\top B^{-1}A] (R^{-1/2} \otimes I_T)| \quad (114)$$

$$= 2 \log |(R^{-1/2} \otimes I_T)| + \log |I_N \otimes I_T - A^\top B^{-1}A| \quad (115)$$

$$= -T \log |R| + \log |B - AA^\top| + \log |B^{-1}| \quad (116)$$

$$= -T \log |R| + \log |I_{DT}| - \log |B| \quad (117)$$

$$= -T \log |R| - \log |B| \quad (118)$$

$$= -T \log |R| - 2 \log |W| \quad (119)$$

where $B = WW^\top$ is the Cholesky decomposition of matrix B . Thus, $\log |K_{yy}|$ is given by:

$$\log |K_{yy}| = T \log |R| + 2 \log |W| \quad (120)$$

where $\log |R| = \sum_i \log R_{ii}$ and $\log |W| = \sum_i \log W_{ii}$.

Finally, to compute products of the form $K_{yy}^{-1}v$, we will need to compute $B^{-1}v'$ for some v' . To do this stably and efficiently, we solve two successive triangular systems via back-substitution: we first solve $Wv'' = v'$ for v'' , and then solve $W^\top z = v''$ for z which returns $B^{-1}v'$.

This direct method of computing the marginal likelihood avoids loss of numerical precision by never explicitly computing any inverse. Moreover, it costs $\mathcal{O}((MT)^3)$, which is typically much less than the naive $\mathcal{O}((NT)^3)$. In [Appendix F.2](#), we show how this cost can be further reduced to $\mathcal{O}(TMN + MT \log T)$ to enable large scale applications.

F.2 Scalability to very large datasets

Here, we show how to reduce both the computational complexity and memory requirements of learning and inference in GP regression and GPFADS using the non-reversible kernels we have proposed. The methods outlined below hinge on the ability to perform very fast matrix-vector products with the Gram matrix, and contain a mix of well-known techniques and novel tricks to be published elsewhere. We first describe a set of methodological building blocks that can be used to scale up GP regression, and later explain how they apply to GPFADS too. Although the main text focuses on theoretical concepts and small-scale applications that do not make use of these acceleration techniques, we have implemented them to good effect.

Fast matrix-vector products with the Gram matrix

The full Gram matrix $K_{xx} \in \mathbb{R}^{MT \times MT}$ associated with the M -output covariance of [Eq. 106](#) for a specific set of T time bins is a sum of $2Q$ space-time Kronecker products of the form $A \otimes F$ where $A \in \mathbb{R}^{M \times M}$ and $F \in \mathbb{R}^{T \times T}$. This allows us to write efficient routines for matrix-vector multiplication with the Gram matrix, by exploiting the identity $(A \otimes F) \text{vec}(V) = \text{vec}(F^\top V A)$. When the time bins are regularly spaced on a grid, then F is a Toeplitz matrix that can be embedded in a circulant matrix (see below), enabling the computation of $F^\top V$ products in $\mathcal{O}(MT \log T)$. Thus, the complexity of a $K_{xx} \text{vec}(V)$ product can be reduced from the naive $\mathcal{O}(M^2 T^2)$ down to $\mathcal{O}(Q(M^2 T + MT \log(T)))$ (from now on, we will assume that $\log(T)$ dominates M , in which case this complexity simplifies to $\mathcal{O}(QMT \log(T))$). Importantly, as we will see below, this way of computing products allows us to lower the memory requirements by never storing the Gram matrix (not even any of its $T \times T$ blocks).

To compute fast $F^\top V$ products with any temporal Gram matrix $F \in \mathbb{R}^{T \times T}$ (e.g. associated with $f_q(\cdot)$ or $\mathcal{H}[f_q](\cdot)$ in [Eq. 106](#)), we use the fact that F is symmetric Toeplitz when the T time bins form a regular grid [Wilson and Nickisch, 2015](#):

$$F = \begin{bmatrix} f_0 & f_1 & \cdots & f_{T-2} & f_{T-1} \\ f_1 & f_0 & \cdots & f_{T-3} & f_{T-2} \\ \vdots & \vdots & \ddots & \vdots & \cdots \\ f_{T-2} & f_{T-3} & \cdots & f_0 & f_1 \\ f_{T-1} & f_{T-2} & \cdots & f_1 & f_0 \end{bmatrix} \quad (121)$$

One can embed this $T \times T$ Toeplitz matrix F into a $2(T-1) \times 2(T-1)$ circulant matrix F_c , every column being a cyclically shifted version of the previous:

$$F_c = \begin{bmatrix} f_0 & f_1 & \cdots & f_{T-2} & f_{T-1} & f_{T-2} & \cdots & f_1 \\ f_1 & f_0 & \cdots & f_{T-3} & f_{T-2} & f_{T-1} & \cdots & f_2 \\ \vdots & \vdots & \cdots & \vdots & \vdots & \vdots & \cdots & \vdots \\ f_{T-2} & f_{T-3} & \cdots & f_0 & f_1 & f_2 & \cdots & f_{T-1} \\ f_{T-1} & f_{T-2} & \cdots & f_1 & f_0 & f_1 & \cdots & f_{T-2} \\ f_{T-2} & f_{T-1} & \cdots & f_2 & f_1 & f_0 & \cdots & f_{T-3} \\ \vdots & \vdots & \cdots & \vdots & \vdots & \vdots & \cdots & \vdots \\ f_1 & f_2 & \cdots & f_{T-1} & f_{T-2} & f_{T-3} & \cdots & f_0 \end{bmatrix} \triangleq \begin{bmatrix} F & S \\ S^\top & F' \end{bmatrix}. \quad (122)$$

Notice that all the information is contained within the first column, which we denote by $\mathbf{f}$. This is the only part of the matrix that needs to be stored explicitly. It is well known that products with a circulant matrix can be performed in the Fourier domain as follows:

$$F_c \mathbf{z} = \text{DFT}^{-1} [\text{DFT}(\mathbf{f}) \odot \text{DFT}(\mathbf{z})] \quad (123)$$

where DFT refers to the discrete Fourier transform and $\odot$ denotes element-wise multiplication. Thus, a product $F\mathbf{v}$ can be computed by padding the vector $\mathbf{v}$ with zeroes to size $2(T-1)$, then computing

$$F_c \begin{bmatrix} \mathbf{v} \\ \mathbf{0} \end{bmatrix} = \begin{bmatrix} F\mathbf{v} \\ S^\top \mathbf{v} \end{bmatrix} \quad (124)$$

and simply discarding the lower $T-2$ elements. The DFT and inverse DFT of a $T \times 1$ vector can be computed efficiently in $\mathcal{O}(T \log T)$ using the Fast Fourier transform (FFT). Thus $C\mathbf{z}$ can be used to exactly compute $K\mathbf{w}$ for any $\mathbf{w}$ in $\mathcal{O}(T \log T)$ computations and $\mathcal{O}(T)$ memory.

Fast and low-memory evaluation of the marginal likelihood and its gradient

Evaluating and differentiating the marginal likelihood in GP regression involves solving linear systems of the form $K_{xx}\mathbf{z} = \mathbf{b}$ for $\mathbf{z}$, as well as computing $\log |K_{xx}|$ and its gradient. Here we describe a set of old and new approaches to performing these computations at scale.

Solving $K_{xx}\mathbf{z} = \mathbf{b}$ systems Computing the quadratic form in the GP log marginal likelihood, i.e. solving linear systems of the form $K_{xx}\mathbf{z} = \mathbf{b}$, can be done to numerical precision via (preconditioned) conjugate gradients (CG; Cutajar et al., 2016).

This iterative method only involves K_{xx} through matrix-vector products, which are fast (cf. above) and do not necessitate the explicit computation and storage of this large matrix. However, CG poses a problem when optimization is performed with the help of automatic differentiation (AD) software, an otherwise very useful way of obtaining gradients of the marginal likelihood. Naively backpropagating through every CG iteration incurs a memory cost proportional to the number of CG iterations, which is often prohibitive (in our case, the worst case would be MT iterations, bringing the memory requirements close to $\mathcal{O}(M^2T^2)$, i.e. the very cost of storing K_{xx} which we seek to avoid in the first place). We were not able to find ways to circumvent this problem in the existing literature.

To mitigate the memory requirements of CG, we went back to basic AD principles and derived a novel, constant-memory way of updating the adjoint of $\mathbf{b}$ and $\boldsymbol{\theta}$ from the adjoint of $\mathbf{z} =$

$K_{xx}(\boldsymbol{\theta})^{-1}\mathbf{x}$ obtained through CG iterations. We use the notation $\overline{X} = \partial\mathcal{L}/\partial X$ to denote the adjoint of X , where $\mathcal{L}$ is the loss (or objective) function of interest. First, recall the chain rule:

$$\overline{\theta}_i = \text{Tr} \left(\overline{K_{xx}}^\top \frac{\partial K_{xx}}{\partial \theta_i} \right). \quad (125)$$

Next, for any CG solve $\mathbf{z} = K_{xx}^{-1}\mathbf{b}$ in the forward pass, the adjoint of K_{xx} must be updated in the reverse pass according to $\overline{K_{xx}} \leftarrow \overline{K_{xx}} - K_{xx}^{-\top} \mathbf{z} \mathbf{z}^\top$ [Giles, 2008](#). Thus, the update for $\overline{\theta}_i$ (the gradient we need for training the model) is

$$\overline{\theta}_i \leftarrow \overline{\theta}_i + \text{Tr} \left[\mathbf{z} (-K_{xx}^{-1} \mathbf{z})^\top \frac{\partial K_{xx}}{\partial \theta_i} \right]. \quad (126)$$

At this stage, although $K_{xx}^{-1} \mathbf{z}$ can itself be computed in the reverse pass using CG without storing K_{xx} , [Eq. 126](#) suggests that one would still need to compute and store $\partial K_{xx}/\partial \theta_i$, a matrix that is just as large as K_{xx} . We reasoned that automatic differentiation libraries can be used in a slightly unconventional way to evaluate [Eq. 126](#) without ever explicitly representing neither K_{xx} nor its gradient, as long as a memory-efficient matrix-vector product routine is available (which is a premise for the use of CG, anyways). The key is to note that for a matrix-vector product $\mathbf{d} = K_{xx}(\boldsymbol{\theta})\mathbf{e}$, the adjoint of $\mathbf{d}$ is to be propagated back to that of $\boldsymbol{\theta}$ according to [Giles, 2008](#):

$$\overline{\theta}_i \leftarrow \overline{\theta}_i + \text{Tr} \left[\mathbf{e} \overline{\mathbf{d}}^\top \frac{\partial K_{xx}}{\partial \theta_i} \right]. \quad (127)$$

Thus, when a matrix-vector product with $K_{xx}(\boldsymbol{\theta})$ is computed under automatic differentiation, the reverse pass will automatically update the parameter vector $\boldsymbol{\theta}$ in the correct way given by [Eq. 127](#) but without representing $\partial K_{xx}/\partial \theta_i$ explicitly. Critically, [Eq. 127](#) has exactly the same form as [Eq. 126](#). By identifying terms, we see that implementing [Eq. 126](#) can be done by performing a dummy matrix-vector product $\mathbf{d} = K_{xx}(\boldsymbol{\theta})\mathbf{e}$ with $\mathbf{e} = \mathbf{z}$, and performing a reverse pass on this dummy operation, taking care of “manually” seeding it with the adjoint $\overline{\mathbf{d}} = -K_{xx}^{-1} \mathbf{z}$ where $\mathbf{z}$ is obtained as part of the primary reverse pass. Note that this requires hijacking the primary reverse-pass on the log marginal likelihood computation, which can be done relatively easily in modern AD software (e.g. autograd, [Maclaurin et al., 2015](#)).

Computing the log determinant and its gradient Exact computation of the log determinant is generally difficult for large Gram matrices. However, stochastic estimators exist for both $\log |K_{xx}|$ and its gradient that only require matrix-vector products with K_{xx} , again greatly alleviating the memory burden. To estimate the log determinant, we use the approach advocated by [Dong et al. \(2017\)](#) based on stochastic trace estimation [Filippone and Engler, 2015](#); [Roosta-Khorasani and Ascher, 2015](#):

$$\log |K_{xx}| = \text{Tr}(\log K_{xx}) \quad (128)$$

$$= \langle \boldsymbol{\xi}^\top \log(K_{xx}) \boldsymbol{\xi} \rangle_{\boldsymbol{\xi}} \quad (129)$$

where the expectation is over any spherical distribution $p(\boldsymbol{\xi})$ with covariance equal to the identity matrix (Hutchinson’s trace estimator). This expectation can be approximated with Monte Carlo samples. For computing $\log(K_{xx})\boldsymbol{\xi}$ products in [Eq. 129](#), we wish to exploit fast and memory-efficient $K_{xx}\mathbf{v}$ products. One way of doing this, which we have not seen used in the GP literature, is to use the following integral representation of the matrix logarithm [Wouk, 1965](#); [Davies and Higham, 2005](#):

$$\log K_{xx} = \int_0^1 (K_{xx} - I) [tK_{xx} + (1-t)I]^{-1} dt \quad (130)$$

Thus, we can estimate the log determinant of K_{xx} by drawing a set of P random vectors $\Xi \in \mathbb{R}^{MT \times P}$ and using any numerical quadrature method to compute

$$\log |K_{xx}| \approx \int_0^1 \text{Tr} [Z^\top (tK_{xx} + (1-t)I)^{-1} \Xi] dt \quad \text{with } Z = (K_{xx} - I)\Xi. \quad (131)$$

Here, the integrand can be computed via CG, based on matrix-vector products with K_{xx} . One can substantially speed up CG convergence by initializing CG iterations at a given t with a previously obtained solution at a neighbouring t' . Note that for ill-conditioned K_{xx} , solvers will typically need to do more work for t near 1 – this is where adaptive solvers can help greatly.

Trace estimation also applies to the gradient of $\log |K_{xx}|$ [Cutajar et al., 2016](#); [Dong et al., 2017](#):

$$\frac{\partial \log |K_{xx}|}{\partial \theta_i} = \text{Tr} \left[K_{xx}^{-\top} \frac{\partial K_{xx}}{\partial \theta_i} \right] \quad (132)$$

$$= \left\langle \xi^\top K_{xx}^{-\top} \frac{\partial K_{xx}}{\partial \theta_i} \xi \right\rangle_\xi \quad (133)$$

$$= \left\langle \text{Tr} \left[\xi (K_{xx}^{-1} \xi)^\top \frac{\partial K_{xx}}{\partial \theta_i} \right] \right\rangle_\xi \quad (134)$$

As in the computation of the quadratic form in [Eq. 126](#), [Eq. 134](#) seems to require an explicit representation of the large matrix $\partial K_{xx} / \partial \theta_i$. However, the same trick we introduced above can be used here too to exploit the existing machinery of AD software to evaluate [Eq. 134](#) without storing any large matrix. Specifically, note that [Eq. 134](#) has the exact same form as [Eq. 127](#), such that it suffices to compute a dummy matrix-vector product $d = K_{xx}e$ with $e = \xi$ and perform a dummy reverse pass seeded with $\bar{d} = K_{xx}^{-1}\xi$ (obtained via CG).

In sum, these tricks reduce the complexity of estimating the model evidence and its gradient from the naive $\mathcal{O}(M^3T^3)$ to $\mathcal{O}(QMT \log T)$ (potentially with a large constant pre-factor determined by the number of CG iterations, hence the importance of good preconditioning; [Cutajar et al., 2016](#)). Similarly, the memory requirements scale as $\mathcal{O}(M^2T)$, instead of $\mathcal{O}(M^2T^2)$.

Accelerating GPFADS

In GPFADS, since $K_{yy} = (C \otimes I_T)K_{xx}(C^\top \otimes I_T) + R \otimes I_T$, efficient products with K_{xx} also lead to efficient products with K_{yy} which do not necessitate explicit storage of K_{yy} . Therefore, the techniques described above apply to GPFADS directly. With $C \in \mathbb{R}^{M \times N}$, computing a $(C^\top \otimes I_T)v$ product costs $\mathcal{O}(MNT)$, and subsequent multiplication by K_{xx} costs $\mathcal{O}(QMT \log T)$ as detailed above. Thus, the cost of a $K_{yy}v$ product – which dominates the complexity of training the GP model and computing posteriors – is $\mathcal{O}(MNT + QMT \log T)$ overall. The memory requirement is $\mathcal{O}(NT + M^2T)$ (for small latent spaces, this is close to the cost of storing a data point in the first place).

G Relation to Parra and Tobar (2017)

We are only aware of one other paper introducing a non-reversible GP kernel [Parra and Tobar, 2017b](#). Here, we outline the key differences between their models and ours, and show how our construction affords both a cleaner handle on reversibility, and better scalability properties.

Parra and Tobar (2017b)’s model extended the spectral mixture model Wilson and Adams, 2013b to the multi-output setting. In their model, the cross-spectrum of any two variables i and j is given by:

$$\mathbb{E} \left[x_i(\omega) \overline{x_j(\omega)} \right] = \frac{w_{ij}}{2} \left(e^{-\frac{(\omega - \mu_{ij})^2}{2\sigma_{ij}} + j(\theta_{ij}\omega + \phi_{ij})} + e^{-\frac{(\omega - \mu_{ij})^2}{2\sigma_{ij}} + j(-\theta_{ij}\omega + \phi_{ij})} \right), \quad (135)$$

where ϕ_{ij} is a phase lag and θ_{ij} is a pure delay. The parameters μ_{ij} , σ_{ij} and w_{ij} are directly constrained by the marginal spectrum of the two variables (spectral mixture) ensuring that the resulting multi-output covariance function be positive definite.

In the time domain, this leads to the following real-valued cross-covariance function:

$$K_{ij}(\tau) = w_{ij} (2\pi\sigma_{ij}) \exp \left(-\frac{\sigma_{ij}}{2} (\tau + \theta_{ij})^2 \right) \cos((\tau + \theta_{ij})\mu_{ij} + \phi_{ij}). \quad (136)$$

A first notable difference between this kernel and ours lies in the ability – or lack thereof – to break reversibility at low frequencies. Indeed when there are no hard delays ($\theta_{ij} = 0$), the Parra and Tobar (2017b) model becomes fully reversible in the limit of a non-resonant process (one whose power spectrum has its maximum at $\omega = 0$, modeled by setting $\mu_{ij} = 0$ in Eq. 135), no matter the choice of phase lag parameters ϕ_{ij} . Indeed, the non-reversibility index ζ is tied to the number of oscillatory cycles that fit within the envelope of the autocovariance function (Fig. 22, right). In contrast, our construction retains non-reversibility in this limit, because the Hilbert transform implies a constant phase lag at all frequencies. For example, the non-reversible squared-exponential planar covariance shown in Fig. 11 has full non-reversibility despite a complete lack of marginal oscillations in each output.


The second major difference lies in the opportunity – or lack thereof – to scale learning and inference to very large datasets. As explained in Appendix F.2, all kernels in the family that we propose take the form of a sum of space/time-separable terms, implying a specific sum-of-Kronecker-products structure for the associated Gram matrices. In contrast, Parra and Tobar (2017b)’s model can never be expressed in Kronecker form, if reversibility is to be broken via the introduction of non-zero phase lags. In other words, to our knowledge, we have presented the first non-reversible multi-output kernel that can realistically support scalable learning and inference on very large datasets.

Appendix — Expansion Microscopy


H ExM Chemistry

Fixation Fixation is the first step of such protocols and is typically done with formaldehyde. This cross-links proteins, DNA and RNA with little specificity. Underfixation leads to excess protein loss; however, overfixation comes at the risk of attacking epitopes, which are then not detected during IHC. Fixation is typically performed at low temperature to ensure that PFA molecules have enough time to diffuse throughout the tissue before reacting too much.

Paraformaldehyde:



Paraformaldehyde polymer:




Triton X-100 is a surfactant, having hydrophilic and lipophilic groups. Intercalates within cell membranes, disrupting them, and leading to permeabilization. Empirically, incubating a sample with a small molecule dye leads to a core-shell staining pattern, whereas post-triton incubation, the fish are homogeneously stained.


Anchoring The first step in ExM is to bind the sample proteins to a moiety that can be incorporated in the future gel polymer. This is typically done via AcX, which is an acryloyl molecule bound to a hydroxysuccinimide ester (NHS) group.

The NHS groups react with primary amine groups, forming amide bonds, with the NHS group being released. This leaves primary amines within the tissue bound to an acryloyl group. Lysines are the only amino acid that contain an amine as a secondary group, but are common, so likely to be found in all proteins.

NHS ester reaction where $R1$ would be part of acryloyl in this context, and $R2$ the rest of lysine:



→




Monomer solution Next, we add the gel monomer solution. The monomer solution is composed of acrylamide monomer, sodium acrylate, and N,N-methylenebisacrylamide (Bis). Acrylamide is the base monomer. Sodium acrylate is a charged version of the former: Instead of having an amide group, it contains a carboxy group. At high salt concentrations, sodium comes to mask the negative charge on the oxygen. At low salt concentrations, unmasking occurs, leading to repulsions of negatively charged oxygens, which in turn leads to expansion. Bis is essentially two acrylamide molecules that have been cross-reacted with their amine groups. Bis has two free double carbon bonds, which can be incorporated into two distinct polymer chains thus enabling crosslinking of multiple polymer chains together.


Compounds and formulae:

- acrylamide monomer (AcAm)
- sodium acrylate (NaAc)
- N,N-methylenebisacrylamide (Bis)


Acrylamide:



Acrylate:



Acrylamide polymer:



Polymerization reaction

Initiation Initiation of polymerization is induced by the addition of the ammonium persulfate (APS) initiator, which forms free radicals by chemical decomposition ($(NH_4)_2S_2O_8 \rightarrow 2SO_4^{\cdot-} + 2NH_4^+$). The decomposition occurs spontaneously, albeit slowly; to accelerate the reaction rate, tetramethyl-ethylenediamine (TEMED) is added, which stabilizes the intermediate radical species formed during the decomposition of APS, thereby favoring the reaction. Free radicals react with the double carbon bond of the monomer and initiate a chain reaction. However, polymerization should not occur before the polymer reagents have diffused and equilibrated throughout the sample. To control the initiation time, the compound 4-hydroxy TEMPO (4HT) is also added. 4HT consumes free radicals and therefore adds delay. Oxygen also inhibits gel formation for the same reason. The low temperature also reduces the reaction rate.

Monomer incubation The sample is incubated in the solution at low temperature to allow all the compounds to diffuse within the sample but to minimize gel formation.

Gel formation The sample is then encased in a chamber of the desired end shape and brought to a higher temperature to induce the polymerization reaction (37 ° C for 1.5 to 2 hours).

Digestion The proteins of the samples are at this point conjugated to the polymer network that now exists within it. Little expansion has yet occurred, as the sample is still in a high-salt concentration solution. To enable homogeneous expansion, one now needs to disrupt existing bonds within and between proteins so that the gel can expand with minimal constraints. How this is achieved is crucial. To remove lipids and disrupt noncovalent bonds, detergent is added: Sodium dodecyl sulfate (SDS), a low molecular weight nonionic detergent widely used for membrane protein solubilization. TRIS buffer is added to maintain constant pH. To break peptide bonds, one approach is to use proteolytic digestion, such as with the enzyme proteinase K. The advantage of this is that it can be done at relatively low temperatures. However, it requires enzyme access to the tissue, and the enzyme is also somewhat biased toward which peptides it prefers to cleave. An alternative to this is disruption by high temperature. This denatures proteins (sodium dodecyl sulfate (SDS) is added to prevent them from clumping), and thus noncovalent bonds are broken. In addition to this, spontaneous hydrolysis occurs by breaking of peptide bonds. The advantage of this is that it does not require enzyme access; however, it also induces fragmentation of RNA molecules.

Total protein stains Total protein stains typically work with the same chemistry as AcX, that is, NHS click chemistry: Dyes are conjugated to NHS moieties that attack primary amines, thus labeling lysines. As the densities of this amino acid differ across the cellular compartments, this results in an informative histological stain. Dyes also have different properties, such as their hydrophobicity, which also affects the labeling propensity. Thiols are another group for which there exists click chemistry using maleimide moieties.

Example of dyes include:

- NHS-Alexa 488 (hydrophilic)
- ATTO-647N maleimide

H.1 Whole Fish ExM-FISH Protocol

Expansion-assisted iterative fluorescence *in situ* hybridization (EASI-FISH) in whole larval zebrafish.

Virginia Ruetten, Mark Eddison, Amy Hu, Misha Ahrens, Paul Tillberg
Janelia Research Campus, Ashburn, VA, USA.

Recently, a straightforward, robust, and reliable protocol that utilizes expansion microscopy (ExM) and hybridization chain reaction (HCR) amplification for multiplexed *in situ* hybridization for thick slices of mouse brain has been described (EASI-FISH, Wang et al., 2021). Here, we describe a modified version of this protocol optimized for whole larval zebrafish.

Expansion-assisted iterative fluorescence in situ hybridization (EASI-FISH) in whole larval zebrafish.	1
Protocol	1
Day 0: Fixation and Permeabilization	1
Day 1: Anchoring RNA and Protein	2
Day 2: Gelation and Proteinase K digestion	3
Day 3: Re-embedding	4
Day 4: Probe Hybridization	5
Day 5: Probe Washing	5
Day 5: Hybridization Chain Reaction (HCR)	6
Day 6: Mounting and Imaging	6
Reagents and Recipes	7
References	10

Protocol

Day 0: Fixation and Permeabilization

Solutions:

PBS (Fisher; BP24384)

4% PFA (Electron Microscopy Services; 15710) diluted in PBS.

PBST (nuclease-free 1x PBS with 0.5% and 0.1% Triton, PBS-T-0.5 and PBS-T-0.1 respectively).

70% Ethanol (in nuclease-free PBS)

RNAse Away (ThermoFisher; 7003)

1. Fix samples in 1 ml cold 4% PFA overnight @ 4°C on a shaker.
2. Rinse samples in 1 x 2 ml cold PBST-0.5.
3. Wash sample in 1 ml cold PBST- 4 x 15 min, on a shaker.
4. Rinse sample in 2 x 2 ml with cold 70% EtOH.
5. Store samples in 2 ml 70% EtOH @ 4°C for up to 6 months. Make sure samples are in a closed chamber to avoid evaporation.

Day 1: Anchoring RNA and Protein

Solutions and materials:

- 20mM MOPS Buffer (Fisher; BP308-100). Store @ RT.
 - Melphalan-X (Stock concentration: 2 mg/ml in anhydrous DMSO; working concentration: 1 mg/ml; see below for preparation instructions). Store @ -20°C.
 - Acryloyl-X (ThermoFisher; A20770, Stock concentration: 10 mg/ml in anhydrous DMSO; working concentration: 0.1 mg/ml). Store @ -20°C.
 - PBS-T (0.5% and 0.1% Triton in nuclease-free 1x PBS).
 - Poly-Lysine (Pelco; 18026, 1.6 ml + 3.2 ml Photo Flo-200 (EMS, 74257). Store @ 4°C.
 - 0.2 ml PCR tubes (USA Scientific; 1402-4700).
 - RNAse Away (Thermo Scientific; 7003).
 - Scotch Tape
 - Non-charged slide (Fisherbrand; 12-550-123)
 - Low melting point agarose (Thermo Fisher; 16520050). Store at RT.
1. Rehydrate the samples with 2 x 15 min wash in 500 ml PBS-T-0.1.
 2. (Optional) On a petri dish, cover each sample in a thin layer (less than 0.5 mm thick, otherwise monomer solutions will need longer to fully diffuse in) 0.7% low melting point agarose, orientate as desired (on side or dorsal side up) and leave to solidify. With a razor blade cut out a rectangle around the sample and delicately lift off from the petri dish. (Pre-embedding in agarose will ensure the sample has the desired orientation).
 3. Place each sample in a separate PCR tube.
 4. Incubate samples 1 x 30 min in 150 ml 20 mM MOPS Buffer.
 5. Thaw Melphalan-X and Acryloyl-X solutions.
 6. Dilute Melphalan-X stock 1:1 with MOPS Buffer.

7. Add 1:100 Acryloyl-X (10 mg/ml) to Melphalan-X working solution. Vortex, mix, and spin.
8. Remove all MOPS buffer from sample, leaving behind no more than 5 ul.
9. Add 100 ul of Melphalan-X/ Acryloyl-X solution to each PCR tube and gently mix.
10. Incubate overnight @ 37°C
11. Prepare 'sandwich' chambers for gelation the next day. For each fish, clean a non-charged slide with RNase Away. Layer 6 pieces of Scotch tape together (~11 if sample is pre-embedded in agarose) to create ~0.3 mm thick spacers. Cut two strips of spacer material, several mm wide each, and stick them to the slide ~15 mm apart to form the side walls of the gelation chamber.
12. Coat the glass surface of the gelation chamber with 2 ul poly-lysine using a P20 pipette tip. Air dry and repeat poly-lysine coating.

Day 2: Gelation and Proteinase K digestion

Solutions and materials:

- "Medium Density Gel" monomer solution (see materials) (Store @ - 20°C).
 - 10% Ammonium persulphate (APS) (Sigma; A3678) (Store @ - 20°C).
 - 10% N, N, N', N'-Tetramethylethylenediamine (TEMED) (Sigma; T22500) (Store @ - 20°C).
 - 0.5% 4-Hydroxy-TEMPO (4HT) (Sigma; 176141) (Store @ -20°C).
 - Proteinase K (NEB; P8107S, 800U/ml) (Store @ -20°C).
 - 50mM and 500mM Proteinase K Buffer (Store @ RT).
 - Nuclease-free 1x PBS
 - Small paint brush.
1. Place each sample in a separate well of a 24 well plate.
 2. Rinse 2 x 5 min with 500 ul PBS
 3. Incubate sample for 1hr in 500 ul PBS-T-0.5.
 4. Thaw monomer solution and 4HT, TEMED and APS. Vortex well and keep on ice.
 5. Mix monomer solution and 4HT, TEMED and APS at a ratio of 94:2:2:2 to produce the gelation solution. Vortex. Each chamber needs ~1 ml of gelation solution (100 ul to gel and 800 ul for preincubation..
 6. Place 24 well plate containing samples on ice and incubate each sample with 400 ul gelation solution 2 x 15 mins.

7. With fine tweezers, place fish in the gelation chamber, lying it on its side (or other preferred orientation). If the sample is agarose embedded, lay the agarose block flat.
8. Gently place a cover slip over the sample. If agarose embedded, the cover slip should lay flat on the agarose block.
9. Slowly pipette in 100 ul of gelation solution into each chamber. Avoid any air bubbles.
10. Once the fish is covered in gel solution, polymerize the gel @ 37°C for 1.5 hr.
11. Cool gels on the bench for a few minutes.
12. Take off the coverslip lid with a razor blade. Trim the gels into a rectangle.
13. Take off the gel from the slide and transfer to a 2 ml Eppendorf.
14. Prepare first digestion solution: 500mM ProK Buffer with 0.3% SDS, ProK Enzyme added 1:50 (750 ul per gel).
15. Incubate each gel with the first digestion solution @ 50°C for ~4 hrs.
16. Prepare second digestion solution: 50mM ProK Buffer with 1% SDS, ProK added 1:50 (1ml per gel)
17. Incubate each gel with the second digestion solution overnight @ 50°C.

Day 3: Re-embedding

Solutions and materials:

- PBS (Fisher; BP24384)
 - Gelation solution (made the previous day, see above)
1. Wash gels 3 x 10 mins in PBS.
 2. Trim gels
 3. Make gelation solution (as Day2, step 4).
 4. Incubate gels 1 x 30 min in 400 ul gelation solution at 4°C in 24 well plate (one gel per well).
 5. Make the sandwich chamber with ~11 layers of Scotch tape, one per sample (~20 layers are used if samples were pre-embedded in agarose).
 6. Place gel in chamber, cover with coverslip. Add 100 ul of fresh gelation solution to the chamber.
 7. Gel for 1.5 hr @ 37°C.
 8. Trim gel.
 9. Prepare first digestion solution: 500mM ProK Buffer, 0.3% SDS, ProK Enzyme added 1:50 (1ml per gel)
 10. Incubate each gel with the first digestion solution @ 50°C for ~4 hrs.

11. Prepare second digestion solution: 50mM ProK Buffer, 1% SDS, ProK added 1:50 (1ml per gel)
12. Incubate each gel with the second digestion solution overnight @ 50°C.
13. *Fish should look relatively clear at this point. If not, do another overnight digestion.*

Day 4: Probe Hybridization

Solutions:

- Hybridization Buffer (Molecular Instruments).
 - Probe Wash Buffer (Molecular Instruments).
 - HCR probe sets (from Molecular Instruments, see <https://www.molecularinstruments.com/hcr-rnafish-products>). Store @ -20°C)
 - DAPI (Sigma; D9534) diluted in PBS (Fisher; BP24384) at 500 ng/ml.
1. Wash samples 4 x 15 min with 1 ml PBS.
 2. Thaw and mix Hybridization and Probe Wash buffer.
 3. Incubate each gel in 500 ul Hybridization Buffer for 30 min @ 37°C.
 4. Dilute probes 5 ul in 500 ul (10 nM final concentration) Hybridization Buffer per gel. Vortex. Incubate @ 37°C.
 5. Incubate gel with probes overnight @ 37°C. No shaking is necessary.
 6. Place Probe Wash Buffer and PBS @ 37°C in preparation for washing the next day.

Day 5: Probe Washing

Solutions:

- Probe Wash Buffer @ 37°C.
 - PBS (Fisher; BP24384) @ 37°C.
1. Wash each gel 3 x 30 min 750 ul Probe Wash Buffer @ 37°C.
 2. Wash each gel 3 x 30 min 1 ml PBS @ 37°C.
 3. Wash each gel 3 x 1 hr 1 ml PBS @ 37°C.
 4. Wash gels in 1.5 ml of PBS @ RT overnight.

Day 5: Hybridization Chain Reaction (HCR)

(for principals see <https://www.molecularinstruments.com/hcr-rnafish>)

Solutions:

- Amplification Buffer (Molecular Instruments; Store @ 4°C.)
 - Fluorescent Hairpins (448 and 546 from Molecular Instruments, or conjugated to 669, see protocol below. Store @ -20°C).
 - 5x SSCT (5x SSC, 0.1% Tween in Nuclease Free Water). Store @ RT.
 - 0.5x SSCT (0.5x SSC, 0.1% Tween in Nuclease Free Water). Store @ RT.
 - DAPI diluted in PBS (5ug/ml) (labels RNA in the cytoplasm)
1. Incubate each gel in 500 ul Amplification Buffer for at least 30 min @ RT.
 2. Snap cool hairpins with PCR machine @ 95°C. for 90 sec and cool @ RT for 30 min.
 3. For each fluor mix, add hairpins h1 and h2 @ 1:50 in 400 ul Amplification Buffer (1x per gel). Vortex.
 4. Incubate each gel with hairpins for 4 hr @ RT in the dark.
 5. Wash gels 2 x 20 min in 750 ul 5X SSCT @ RT.
 6. Wash gels 2 x 40 min in 1 ml 0.5X SSCT @ RT.

Day 6: Mounting and Imaging

Solutions and materials:

- Sample holders for Z.1 Microscope (files attached).
 - 8 mm circular coverslips (CS-8R, Warner Instruments, 64-0701).
1. Add a poly-lysine coat to the area of the Z1 sample holder where gel will be mounted.
 2. Mount gel, sample side up.
 3. Image on a Zeiss Z1 Lightsheet Microscope. Allow 45 mins (with gel in place) for the system to equilibrate before acquiring multi-tile acquisitions.
 4. After imaging, gels can be teased off the holder with a paint brush.
 5. For long term storage, keep gels in 5X SSC @ 4°C.

DNase treatment for Stripping Hairpins (and genomic DNA).

Solutions:

- RNase-Free DNase1: (Qiagen, 79254)* Note: Do **not** use RDD buffer supplied.

- DNase1 buffer: (10 mM Tris-HCL pH8.0; 2.5 mM MgCl₂; 0.5 mM CaCl₂)
1. Incubate gel for 30 min in 1 ml of DNase1 Buffer @ 37°C
 2. Add 450 ul of DNase Buffer to 50 ml DNase1. Mix. (increase to 750 ul if gel is not fully covered).
 3. Incubate gel in DNase1 for 2 hr @ 37°C.
 4. Wash 4 x 15 min with 1 ml PBS.
 5. For multiplexing with alternative probes go to Day 4, step 2.

Reagents and Recipes

200mM MOPS Buffer (10X Stock):

1046.5 mg in 25 ml in NFW, pH to 7.7 with 10N NaOH. Store @ -20°C.

Melphalan stock (2.5 mg/ml) (Cayman Chemicals; 16665).

Dissolve 2.5 mg per ml in anhydrous DMSO (Invitrogen; D12345). To dissolve, heat to 37°C and vortex vigorously and place on a shaker. May take an hour to dissolve. Aliquot in 800 ul batches. Store in a desiccated environment @ -20°C.

Acryloyl-X (AcX) stock (10 mg/ml) (Thermo Fisher; A20770).

Dissolve in anhydrous DMSO. Aliquot in 20 ul batches. Any extra aliquot in 5 ul batches (for extra-AcX step). Store in a desiccated environment @ -20°C. Don't re-use AcX after thawing.

Melphalan-X (2 mg/ml):

Combine an equal concentration of Acryloyl-X (10 mg/ml) and Melphalan (2.5 mg/ml) (1-part AcX to 4-parts Melphalan (i.e., 200 ul:800 ul). Incubate overnight @ RT with shaking. Store in 50 ul aliquots in a desiccated environment @ -20°C. Use at 1 mg/ml by 1:1 dilution in 20 mM MOPS.

Medium Density Gel monomer solution components:

4 M Sodium Acrylate* – made from Acrylic Acid.

40% Acrylamide (Bio-Rad; 1610140)

2% N,N Methylene-Bis-acrylamide (Bio-Rad; 1610142)

10x PBS (ThermoFisher; AM9625)

Nuclease Free Water (NFW) (ThermoFisher; AM9932)

4 M Sodium Acrylate stock solution

1. In a fume hood, place 5.5 ml of acrylic acid into a 50 ml tube. Place the tube in a room temperature water bath (e.g. a beaker).
2. Add 4.5 ml water.
3. Add 7.2 ml 10 M NaOH gradually to prevent excessive heating and evaporation/boiling.
4. Remove tube from hood (at this point most of the acrylic acid has been converted to non-volatile sodium acrylate).
5. Add 1 M NaOH (nominally 1 ml) gradually until the pH is between 7.5-8 using a pH meter, at RT. Do not use pH test strips.
6. Add water up to a final volume of 20 ml.

Note: acrylic acid has a pKa of 4.76--at pH 7.75 this solution has about 4 mM remaining buffering capacity

For 9.4 ml Medium Density gel (not including APS +TEMED + 4HT):

4 M Sodium Acrylate	2500 µl
40% Acrylamide	1250 µl
2% N,N MethylBisacrylamide	500 µl
10x PBS	1000 µl
Nuclease Free Water	4150 µl

10% APS: 100 mg APS + 900 ul H₂O

10% TEMED: 100 ul TEMED + 900 ul H₂O

0.5% 4HT: 5 mg 4HT + 995 ul H₂O

Store APS, TEMED and 4HT in 150 ul aliquots in PCR tubes @ -20°C.

50 mM ProK + SDS Buffer (50 ml)

50 mM Tris-HCL-pH8, 1 mM EDTA, 0.5% TritonX, 50 mM NaCl, 1% SDS.

2.5 ml 1M Tris-HCL > 50 mM

2.5 ml 10% Triton > 0.5%

0.5 ml 5M NaCl > 50 mM (or 5 ml for 500 mM)

100 µl 0.5M EDTA > 1 mM

5 ml 10% SDS > 1% SDS (or 1.5 ml for 0.3%)

Nuclease Free Water up to 50 ml

RNase-Free DNase1

Add 550 ul DNAase1 Buffer to DNase1 powder. Mix. Aliquot 50 ul per PCR tube, store @ -20°C.

DNase1 buffer (50 ml)

500 ul 1M Tris-HCL > 10 mM
125 ul 1M MgCl₂ > 2.5 mM
25 ul 1M CaCl₂ > 0.5 mM
49.350 ml Nuclease Free Water

HCR Wash Buffers (50 ml)

5xSSCT:

12.5 ml 20x RNase free-SSC (ThermoFisher; AM 9763)
500 ul 10% Tween
37 ml Nuclease Free Water

0.5xSSCT:

1.25 ml 20x RNase free-SSC
500 ul 10% Tween
48.25 ml Nuclease Free Water

JF-669 conjugation to unlabelled hairpins

(Hairpins are amine modified, JF-669 has an NHS ester group).

Reagents

JF-669, SE (Tocris; 6420, Store @ -20°C)
1 nmol unlabelled amine-modified hairpins (~70mer, Molecular Instruments; Store @ -20°C).
Acetonitrile, anhydrous (Thermofisher; 042311-K7).
0.1M Sodium Bicarbonate pH 8-9.
Anhydrous DMSO (Invitrogen; D12345).
QIAquick Nucleotide Removal Kit (250) (Qiagen; 28306).
0.5 ml screw-cap microcentrifuge tubes, Amber (USA Scientific; #1405-9707).

1. Turn on speed vacuum (eg. Thermofisher; SPD120) and defrost dye @ RT for 30 min.
2. Resuspend 2 mgs of JF-669, SE in 630 µl Acetonitrile, mix and vortex, and aliquot 30 µl into labeled skirted 0.5 ml screw-cap centrifuge tubes > each will contain 0.1 mg of dye.

3. Evaporate acetonitrile in speed vacuum (in organic solvent mode) for 45 min. Store @ -20°C.
4. In two 1.5 ml Eppendorf tubes, evaporate 5 µl (500 pmol/10 µg) of unlabelled hairpins h1 and h2 using a speed vac, in aqueous mode, for 30 min. If you use 10 µl (1 nmol) use two tubes per hairpin. Check the solution has fully evaporated.
5. Add 3 µl of 0.1M Sodium Bicarbonate pH 8-9 to each evaporated hairpin. Mix.
6. Add 2 µl of anhydrous DMSO to 0.1 mg of dye. Mix.
7. Add 2 µl dye mix (100 mg) to each 3 µl of hairpin (10 µg). Mix. It will change colour.
8. Leave hairpin-dye mixture to react overnight @ RT in the dark.
9. The next morning, add 5 µl nuclease-free water to bring to 10 µl.
10. Remove excess dye with a QIAquick Nucleotide removal kit (add 100 µl PN1).
11. Elute dye-oligo conjugate in 50 µl nuclease-free water.
12. Check the hairpin concentration on a spectrophotometer (eg. NanoDrop One, ThermoScientific) and dilute to 60 ng/µl.
13. Separately store hairpins h1-669 and h2-669 in 25 µl aliquots in a PCR tube @ -20°C.
14. Test conjugation by HCR.

References

Wang et al., EASI-FISH for thick tissue defines lateral hypothalamus spatio-molecular organization. *Cell* 184 (2021) 6361-6377.e24. DOI: [10.1016/j.cell.2021.11.024](https://doi.org/10.1016/j.cell.2021.11.024)

H.2 Whole Fish ExM-IHC Protocol

Expansion-assisted immunohistochemistry (IHC) in whole larval zebrafish.

Virginia Ruetten, Mark Eddison, Amy Hu, Misha Ahrens, Paul Tillberg
Janelia Research Campus, Ashburn, VA, USA.

Expansion-assisted immunohistochemistry (IHC) in whole larval zebrafish.	1
Protocol	1
Day 0: Fixation and Permeabilization	1
Day 1: Bleaching and Immunohistochemistry (primaries)	2
Day 4: Immunohistochemistry (secondaires)	2
Day 6:	3
Protein Anchoring	3
Chamber preparation	3
Gelation and Disruption	4
Day 7:	5
Re-embedding	5
Chamber preparation	5
Second Gelation	6
Staining	6
Day 9: Mounting and Imaging	6

Protocol

Day 0: Fixation and Permeabilization

Solutions and materials

- PBS (Fisher; BP24384)
- 16% PFA (Electron Microscopy Services; 15710) diluted to 4% in PBS.
- PBST (1x PBS with 0.5% or 0.1% Triton, PBST-0.5 and PBST-0.1 respectively).

Protocol

1. Euthanize samples with an overdose of MS-222 (tricaine).
2. Fix samples in 1 ml cold 4% PFA overnight @ 4°C on a shaker.
3. Rinse samples in 1 x 2 ml cold PBST-0.5.

Day 1: Bleaching and Immunohistochemistry (primaries)

Solutions and materials

- PBS (Fisher; BP24384)
- PBST-0.5
- Goat Serum
- Na Azide
- Triton-100
- Blocking Buffer (5% goat serum, 0.5% Triton-100 in 1x PBS, 0.01% Na Azide)
- Bleaching solution (180 μ L in 1M KOH, 60 μ L of H₂O₂ in 1 ml of H₂O)
- Primary antibodies

Protocol

1. Prepare fresh bleaching solution (180 μ L in 1M KOH, 60 μ L of H₂O₂ in 1 ml of H₂O).
2. Incubate samples in 500 μ L of bleaching solution until bubbling is observed. Replace the bleaching solution with fresh bleaching solution. Monitor the samples until pigmentation disappears (~2min).
3. Rinse the samples twice with PBST-0.5.
4. Incubate the samples in Blocking Buffer for 3 hours at room temperature (RT).
5. Incubate the samples in primary antibodies diluted 1:100 in Blocking Buffer for 3 days at RT on a shaker. Make sure NaAzide is included to prevent microbial growth.
6. Wash sample in 1 ml cold PBST-0.5 (4 x 15 min), on a shaker.

Day 4: Immunohistochemistry (secondaires)

Solutions and materials

- Blocking Buffer (5% goat serum, 0.5% Triton-100 in 1x PBS, 0.01% Na Azide)
- Secondary antibodies

Protocol

1. Wash the samples in Blocking Buffer (3 x 2 hours) on a shaker.
2. Incubate the samples in secondary antibodies (1:100) in Blocking Buffer for 2 days at RT on a shaker.

Day 6:

Protein Anchoring

Solutions and materials

- Acryloyl-X (ThermoFisher; A20770, Stock concentration: 10 mg/ml in anhydrous DMSO; working concentration: 0.1 mg/ml). Store @ -20°C.
- PBS-T (0.5% or 0.1% Triton in 1x PBS).
- Low melting point agarose (Thermo Fisher; 16520050). Store at RT.

Protocol

1. Wash the samples in Blocking Buffer (2x 2 hours, at RT). For specimens that have not been stained with antibodies, permeabilize for at least 5hr in Blocking Buffer.
2. Wash the samples in 1x PBS and transfer to fresh 1x PBS.
3. (Optional) On a petri dish, cover each sample in a thin layer 0.7% low melting point agarose, orientate as desired (on side or dorsal side up) and leave to solidify. With a razor blade cut out a rectangle around the sample and delicately lift off from petri dish. (Pre-embedding in agarose will ensure sample has the desired orientation).
4. Prepare Acryloyl-X solution (working solution: 20µg/ml, dilution 1:500 in PBS) . Do not prepare in advance.
5. Incubate samples in Acryloyl-X solution (1x 1hr at RT).

Chamber preparation

Solutions and materials

- Poly-Lysine (Pelco; 18026, 1.6 ml + 3.2 ml Photoflo-200 (EMS, 74257). Store @ 4°C.
- 0.2 ml PCR tubes (USA Scientific; 1402-4700).
- Non-charged slide (Fisherbrand; 12-550-123)

Protocol

1. Prepare chambers for gelation, one for each sample. Layer 6 pieces of Scotch tape together (~11 if sample is pre-embedded in agarose) to create ~0.3 mm thick spacers. Cut two strips of spacer material, several mm wide each, and stick them to the slide ~15 mm apart to form the side walls of the gelation chamber.
2. Coat the glass surface of the gelation chamber with 2 ul poly-lysine using a P20 pipette tip. Air dry and repeat poly-lysine coating.

Gelation and Disruption

Solutions and materials

- First monomer solution (see materials) (Store @ - 20°C).
- 10% Ammonium persulphate (APS) (Sigma; A3678) (Store @ - 20°C).
- 10% N, N, N', N'-Tetramethylethylenediamine (TEMED) (Sigma; T22500) (Store @ - 20°C).
- 0.5% 4-Hydroxy-TEMPO (4HT) (Sigma; 176141) (Store @ -20°C).
- 1x PBS
- Small paint brush.
- Disruption Buffer (see materials).

Protocol

1. Rinse samples with PBS.
2. Thaw first gelation monomer solution as well as 4HT, TEMED and APS. Vortex well and keep on ice.
3. Mix first monomer solution and 4HT, TEMED and APS at a ratio of 94:2:2:2 to produce first gelation solution. Vortex. Each chamber needs ~1ml of gelation solution (100 ul to gel and 800 ul for preincubation).
4. Incubate samples in first gelation solution on ice (3x 10mins) on a shaker.
5. With fine tweezers, place fish in the gelation chamber, lying it on its side (or other preferred orientation). If agarose embedded, lay the agarose block flat.
6. Gently place a cover slip over the sample. If agarose embedded, the cover slip should lay flat on the agarose block.
7. Slowly pipette in 100 ul of gelation solution into each chamber. Avoid any air bubbles.
8. Once fish is covered in gel solution, place the chambers at 37°C for 1.5 hr to 2hr to let the gel polymerize.
9. Cool gels on the bench for a few minutes.
10. Take off coverslip lid with a razor blade. Trim the gels into a rectangle.
11. Take off gel from slide and transfer to a 2 ml Eppendorf.
12. Incubate the samples in Disruption Buffer at 100°C overnight.

Day 7:

Re-embedding

Solutions and materials

- Second monomer solution (see materials) (Store @ - 20°C).
- 10% Ammonium persulphate (APS) (Sigma; A3678) (Store @ - 20°C).
- 10% N, N, N', N'-Tetramethylethylenediamine (TEMED) (Sigma; T22500) (Store @ - 20°C).
- 0.5% 4-Hydroxy-TEMPO (4HT) (Sigma; 176141) (Store @ -20°C).
- 1x PBS

Protocol

1. Wash the samples in 1x PBS (3x 20mins).
2. Prepare second monomer solution.
3. Mix monomer solution and 4HT, TEMED and APS at a ratio of 94:2:2:2 to produce second gelation solution. Vortex. Each chamber needs ~1ml of gelation solution (100 ul to gel and 800 ul for preincubation).
4. Incubate samples in second gelation solution on ice (3x 10mins) on a shaker.

Chamber preparation

Solutions and materials

- Poly-Lysine (Pelco; 18026, 1.6 ml + 3.2 ml Photoflo-200 (EMS, 74257). Store @ 4°C.
- 0.2 ml PCR tubes (USA Scientific; 1402-4700).
- Non-charged slide (Fisherbrand; 12-550-123)

Protocol

1. Prepare chambers for gelation, one for each sample. Layer ~11 pieces of Scotch tape together (~20 if sample is pre-embedded in agarose). Cut two strips of spacer material, several mm wide each, and stick them to the slide ~15 mm apart to form the side walls of the gelation chamber.
2. Coat the glass surface of the gelation chamber with 2 ul poly-lysine using a P20 pipette tip. Air dry and repeat poly-lysine coating.

Second Gelation

Protocol

1. Carefully place each gel flat in the gelation chamber.
2. Slowly pipette in 100 ul of gelation solution into each chamber. Avoid any air bubbles.
3. Place the chambers at 37°C for 1.5 hr to 2hr to let the gel polymerize.
4. Cool gels on the bench for a few minutes.
5. Take off coverslip lid with a razor blade.
6. Trim the gels into a rectangle, removing any second gel material that formed outside the first gel.
7. Take off gel from slide and transfer to a 2 ml Eppendorf.

Staining

1. Incubate gels in Alexa488-NHS and Atto647N-maleimide dye (1:1000 each in PBS) for 1hr at RT.
2. Wash in PBS (3x 30mins) on shaker.

Day 9: Mounting and Imaging

Solutions and materials

- Sample holders for Z.1 Microscope (files attached).
 - 8 mm circular coverslips (CS-8R, Warner Instruments, 64-0701).
1. Add poly-lysine coat to the area of the Z1 sample holder where gel will be mounted.
 2. Mount gel, sample side up.
 3. Image on a Zeiss Z1 Lightsheet Microscope. Allow 45 mins (with gel in place) for system to equilibrate before acquiring multi-tile acquisitions.
 4. After imaging, gels can be teased off holder with a paint brush.
 5. For long term storage, keep gels in 5X SSC @ 4°C.

9 Research paper declaration

During the course of this thesis the paper "[Non-reversible Gaussian processes for identifying latent dynamical structure in neural data](#)" was published at NeurIPS 2020. The authors of the manuscript as they appear in the publication are: Virginia Ruetten, Alberto Bernacchia, Maneesh Sahani, Guillaume Hennequin. Guillaume Hennequin and Maneesh Sahani supervised the work. Alberto Bernacchia defined the original metric of non-reversibility and proved it was bounded, as well as gave feedback throughout the time of the work. The work was peer reviewed. No earlier version of the manuscript was uploaded to a preprint server. The content of the work is layed out in [Section 4](#).