UNIVERSITY OF CALIFORNIA
SANTA CRUZ

DATA-DRIVEN FILTRATION AND SEGMENTATION OF MESOSCALE NEURAL DYNAMICS

A dissertation submitted in partial satisfaction of the
requirements for the degree of

DOCTOR OF PHILOSOPHY

in

MOLECULAR, CELLULAR, AND DEVELOPMENTAL BIOLOGY
with an emphasis in STATISTICS

by

Sydney C. Weiser

September 2020

The Dissertation of Sydney C. Weiser
is approved:

Professor James Ackman, Chair

Professor Alexander Sher

Professor David Feldheim

Professor Yi Zuo

Quentin Williams
Interim Vice Provost and Dean of Graduate Studies

Table of Contents

List of Figures

Abstract

Data-driven filtration and segmentation of mesoscale neural dynamics

by

Sydney C. Weiser

The neocortex contains a constellation of sensory-motor regions whose functional interactions provide the basis of cognition and behavior. Simultaneously recording neuronal group activity across the cortical hemispheres is essential for understanding the nature of information flow across cerebral networks. Moreover, unbiased and robust methods for measuring functional interactions across the cortex within individual living subjects is critical for substantive tests of genetic and environmental factors that influence brain development and function. To this end, we image pan-neuronally expressed genetically encoded calcium indicators transcranially across the neocortex of unanesthetized, behaving mice throughout development. Recording from behaving mice produces a unique set of challenges, including optical and blood artifacts associated with movement. In addition, areal patterning of the cortex can vary among individuals, ages, and genotypes thus an unbiased, flexible workflow for video acquisition and analysis is necessary for producing high quality segmentations of functional structure across neocortex. To address these challenges, we have developed an eigendecomposition-based workflow that isolates hemodynamic and optical artifacts to recover underlying calcium activity patterns, and segments independent regions of the brain to create maps of functional units across the developing cortex. These unique, data-driven maps provide a reference for understanding developmental, genetic, as well as individual variation between functional units of the brain, and

provide a method for extracting optimized time courses from the cortical surface without the need for stimulation-based mapping or anatomical post-processing and alignment. In addition, we quantify the quality of separation of independent sources and use the resulting metrics as feedback to optimize our video acquisition parameters. The open source methods developed here are flexible enough to be utilized with various subject ages and genotypes, as well as numerous experimental configurations and computer architectures. Here we additionally present applications of these methods to anesthesia and development datasets. Overcoming these challenges opens the possibility of using these techniques to help address a number of key objectives in neuroscience, including the quantification of robust inter-areal dynamics between functional motifs across cerebral networks.

This work is dedicated to

my family,

for always supporting me throughout my adventures

my friends,

for encouraging me and believing in me

and to my partner,

who grounds me and motivates me to pursue my passions

Acknowledgments

First and foremost, I'd like to thank my advisor James Ackman, for encouraging me in my data exploration and analysis, supporting my requests for more computational power that many advisors would have rejected (like a dedicated computing node with a terabyte of RAM), and always challenging me to look for the bigger picture.

Brian Mullen was the best labmate I could have asked for. We worked closely to synchronize efforts on the data acquisition and analysis every step of the way, and was always there to discuss ideas, data, and to provide moral support through our graduate careers. He additionally conducted all of the experimental recordings for the data presented in this paper.

I'd also like to thank my mentee Desiderio Ascencio, for his work on these projects, as well as his support, energy, and enthusiasm for research. It's been great to watch you develop as a scientist, and I look forward to seeing where your future takes you.

I'd also like to acknowledge Alexander Sher, my undergraduate advisor and thesis committee member, for his early influence on me as an aspiring scientist. He taught me to always seek the thorough analysis of the data, and question the meaning of my results. I'd also like to thank the rest of my committee, for their insight and input into how to wrap my years of methods development into a single paper.

I'd additionally like to thank the Molecular, Cellular, and Developmental Biology department, for welcoming me, teaching me the fundamentals in techniques, as well as how to think about and communicate biological questions. I also sincerely appreciate the way everyone always believed that even with my limited previous backgorund in biology, I could succeed as

a biology PhD student. The UCSC MCD and Neuroscience community have also been a great resource, whether through seminars, meeting as friends, or discussing results over beers at NeuroClub, your support has meant a lot to me.

I'd also like to thank my friends and family, for all your moral support and encouragement throughout the years; your support has been invaluable. Thank you!

The contents of this thesis are currently being prepared for publication; the main chapter will be published as a methods paper, in conjuction with Brain's machine learning methods paper for automated component sorting. The development data was predomininatly Brian's work; he is planning on preparing and submitting the paper for publication later this year or early next year. We are also planning on publishing results on the anesthesia data as well, but research on this data is still ongoing.

Chapter 1

Introduction

1.1 Activity Patterns Instruct Circuit Maturation

The mammalian cerebral cortex is formed over an extended developmental time course, ranging from several months for a rodent to many years for a human [64]. Coupled with genetic programs, neural activity patterns play a critical role in this developmental process [66, 39, 37].

Cortical neurons exhibit correlated calcium transients even before cortical layers are finished developing [39]; first within the cortical column [38], then throughout local and global networks as circuits mature [3, 37]. This process has been shown to be critical for neural development by refining networks, both through anatomical and functional connections [3]. Networks are tuned gradually through this process, with increasingly fast and complex patterns [3], until the brain switches into a more adultlike state at eye opening [13, 34].

The refinement of neural circuits is highly plastic and dynamic, and continues throughout adulthood, shaping local and global networks that allow us to learn new skills and adapt to

new situations [25, 58]. However, the extended ontogenetic period for the neocortex, together with its latent plasticity for synaptic rewiring, also make the cortex especially susceptible to intrinsic or extrinsic factors that disrupt neuronal function [64]. Serious cortical malformations can arise when neurotransmission is disrupted, such as maternal use of epilepsy medication or in-utero exposure to pollutants [29, 47, 12, 59].

We are beginning to understand how activity patterns during development wire individual maps within primary sensory areas [3, 2, 50], yet we still know little about how this activity influences the cortex's functional networks. This aspect of brain development is critical for our understanding of neurodevelopmental disorders, since most mental disorders affect balance of local and global networks of the brain, rather than the function of a single area [61, 10, 7, 74].

In particular, the role of neural dynamics within localized functional networks are particularly understudied. Alzheimer's patients have altered regional connectivity, but little to no differences in global networks [7], while patients with schizophrenia have altered regional connectivity and global networks [48]. These studies have all been performed by observing differential localized clustering on the smallest resolution available in fMRI–the regional level [26]. It is very likely that sub-regional functional differences can also be detected, on the order of communications between cortical columns in sub-regional areas. However, these networks are too small to effectively study with the relatively low spatiotemporal resolution of fMRI.

Understanding the dynamics of the developing brain provide insight into how small perturbations in neural circuits can lead to drastically different local and global activity patterns, resulting in altered behavior [16]. These questions have been left unanswered due to

technological limitations; however, recent advances in bioengineering and computational science necessary now allow us to address these questions. Genetic and pharmocological tools allow easy modification of sensory inputs and dynamics within the brain. Recent advances in calcium imaging technologies allow us to monitor neural activity across populations in the whole neocortex with high enough spatiotemporal resolution to see subregional networks [11]. In addition, advancements in computational bandwith have let to practical improvements in the analysis of dense multivariate datasets, such as those found in high resolution scientific imaging.

We are combining the power and flexibility of mouse genetics with various engineering, computational, mathematical, and biological techniques to provide new insight into the developing local and global networks of the brain. Knowing how these networks interact during development is crucial for understanding how genetic or environmental factors can misregulate activity patterns and thus alter the functional maturation of circuits and networks [48]. This may lead to altered maturation of the cerebral cortex and ultimately the development of neurological disorders such as autism or schizophrenia [16, 61].

1.2 Mesoscale Calcium Imaging

We record developmental neural activity patterns at high spatiotemporal resolution, transcranially imaging fluorescence from a genetically encoded calcium indicator across the entire mouse neocortex. Calcium dynamics are necessary and sufficient for neurotransmission, and calcium indicators are being increasingly widely used in wide-field population imaging studies due to their high signal quality, and ease of targeted expression [11].

The genetically encoded calcium indicator, GCaMP6s, is expressed in all neurons under the control of the Snap25 promoter. This results in uniform cortical expression of the fluorescent transgene by three days after birth in mice [4, 60]. Calcium influx due to neuronal activity causes intracellular GCaMP6s to change into an active conformation that fluoresces green light [11]. Increases in action potentials lead to an additive increases in neural fluorescence; regions with brighter fluorescence correspond to an increase in neural activation frequency [11].

We expose and illuminate the cranium with blue wavelength light and capture fluoresced green light with a sCMOS camera. To observe the spatiotemporal properties of these population neural activity patterns, we crop the video to only neural tissue (neocortex, olfactory bulbs, and superior colliculus when visible), and compare the change in fluorescence over the mean fluorescence: $\Delta F/F$ (fig. 1.1).

We record these signals at very high resolution (2160x2560 px, $\sim$ 6.75um/px), and capture sub-regional mesoscale dynamics, with scale ranging from complex activation patterns in high-order circuits, to individual barrel activations, to whole cortical lobe activity patterns. During development, the skull is thin enough to clearly see these patterns transcranially, with

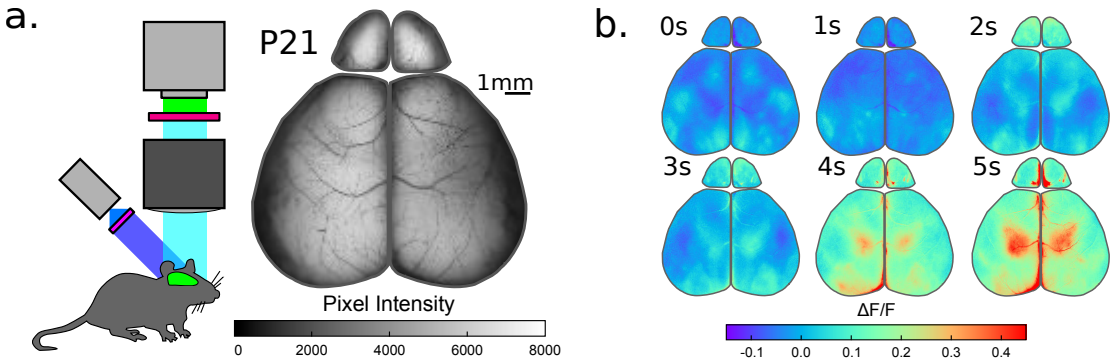

Figure 1.1: a) Recording schematic and fluorescence image of transcranial calcium imaging preparation, cropped to cortical regions of interest. b) Sample video montage of video frames after dF/F filtering.

very little diffusion.

Transcranial imaging leaves the skull and dura intact for a recording procedure that is minimally invasive, allowing us to record activity patterns that are highly sensitive to anesthetics [4]. Even under light doses of isoflurane, spontaneous visual signals are severely inhibited during development [4]. Thus, the mice must be unanesthetized during the recording process. Body or facial movements create large fluctuations in autofluorescence of the brain and blood vessels [41], which produce significant artifacts in the data. Furthermore, growing evidence suggests that even light anesthesia has a significant impact on brain activity patterns /cite{tan_spatial_2015}, further explored in sec. 3.1.

1.3 Separating Neural Signal from Artifacts

Historically, the most common way to reduce blood and movement associated artifacts in neuroscience experiments of all kinds (i.e. electrophysiology, intrinsic signal imaging, fluorescent dye imaging) is by preventing motion, often through induced motor paralysis or anesthesia [CITATIONS]. Inducing motor paralysis fundamentally changes the signals we're interested in, by inhibiting motor activity, and growing evidence suggests that network behavior can be very different under even light anesthesia [67, 63].

There exist a variety of methods that use a multi-wavelength collection system to estimate and correct for backscattering and other hemodynamic associated effects [28, 27, 42, 69]. We attempted multi-wavelength correction methods, but found that at the light levels necessary for good signal collection, it was nearly impossible to shield these alternating light sources from the young animal, who would grow increasingly agitated at the flashing from the alternating light sources. These techniques offer corrections of signal from the absorbance of GCaMP fluorescence by hemoglobin, but do not remove movement, recording artifacts, vessel contraction artifacts, or optical aberrations often present on the surface of recordings from young ($<$P5) mice.

Eigendecompositions have previously been used to identify and filter components of signal [34, 28, 4], and present a flexible method of filtering that is not hardware dependent, and can be applied to any video dataset regardless of the recording hardware or parameters.

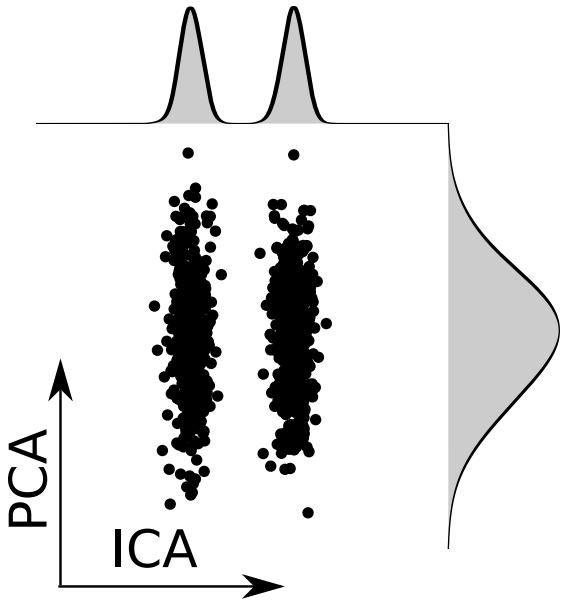

Figure 1.2: Different eigendecomposition algorithms project along different axes. PCA projects along the dimension of greatest variation (vertical), while ICA projects along the most non-Gaussian axis (horizontal).

1.4 Eigendecompositions for Artifact Removal

Eigendecompositions are a set of common signal processing tools that are used to identify underlying patterns in a dataset and transform a dataset from its original dimensions into a new set of dimensions. The original signal is then represented as a linear combination of these components. They are commonly used for data compression or blind source separation.

Principal Component Analysis (PCA) is the most widely known eigendecomposition. It iteratively removes eigenvectors from the data by projecting along the axis of maximal variation (fig. 1.2). Independent Component Analysis (ICA) is another eigendecomposition algorithm. Rather than finding the axis of maximal variation, ICA projects across the most non-Gaussian axis detected [32]. Since random variable are represented as Gaussian distributions,

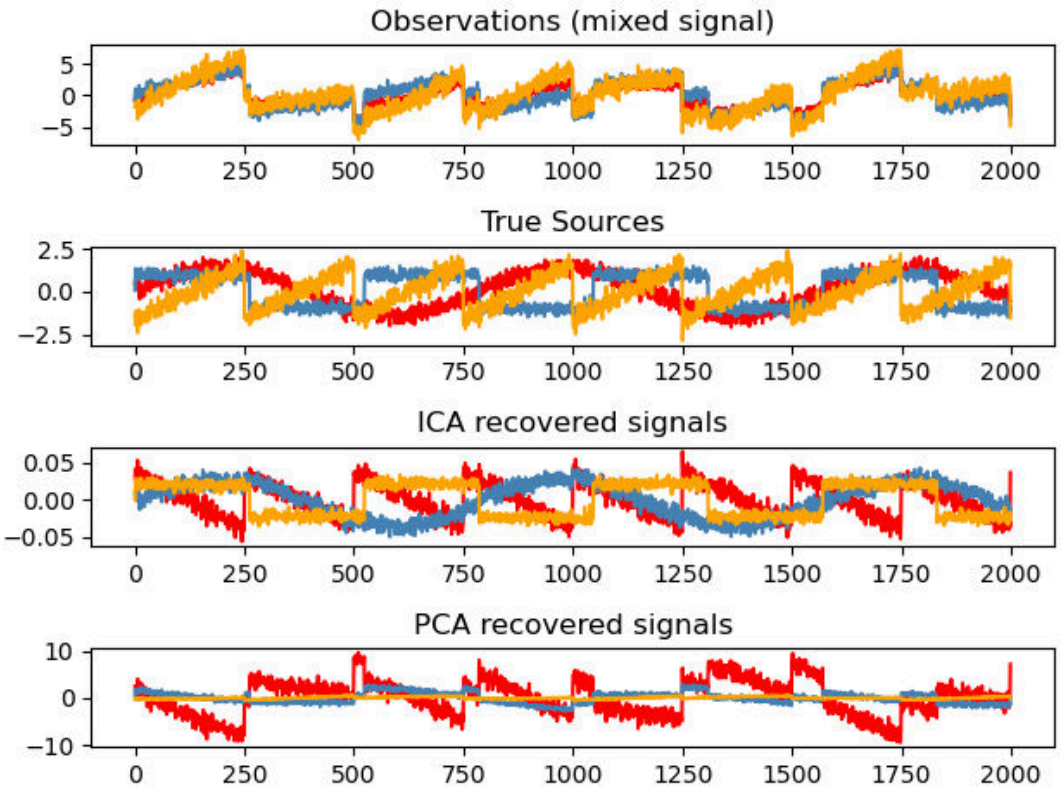

Figure 1.3: The ideal eigendecomposition ideal blind source separator unmixes observations into their underlying source signals. Observations (top) are a mixed sampling of different true source signals. In this case, we have a linear mixture of a sine, sawtooth, and square wave. Applying eigendecomposition blind source separators such as ICA and PCA result in compressed or identified signal source representations. Image source: [1].

this effectively isolates statistically independent sources.

ICA is commonly used as a temporal blind source separator to isolate various independent time series, such as mixed source audio recordings (fig. 1.3). Observations come from mixed signal sources (i.e. different instruments or voices). An eigendecomposition can be used to try to approximate these signals. In fig. 1.3, ICA recovers near approximations to the 3 source signals. One caveat to this approach is that the extracted sources are unordered, unscaled, and

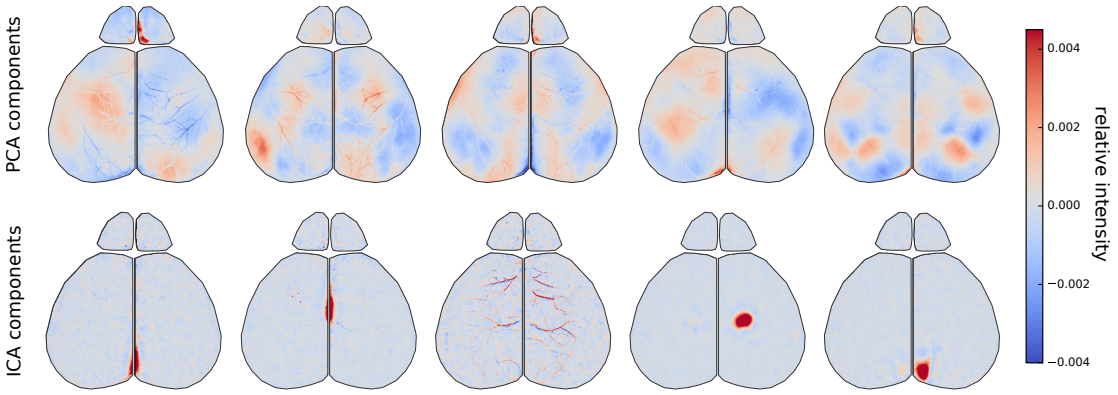

Figure 1.4: Example Components from a PCA decomposition on a P14 animal (top) demonstrate a mixed morphology of signal, with blood vessel effects. Example components from an ICA decomposition (bottom) demonstrate a more clean separation of signal and artifact.

can have inverted sign. In this case, the sawtooth signal is inverted, and the signals are all on vastly different scales from the original data. However, their relative scale is preserved, and when the sources are multiplied by the generated mixing matrix, the original data can be easily recovered in a lossless transformation. Notably, the original data can also be rebuilt from any combination of components. If the square wave represented an unwanted artifact, the original data could be rebuilt with only the sawtooth and sine components. To optimally filter the data, signal and artifact components must be fully separated such that no components represent a mixture of signal and artifact. When applied to video data, a spatial ICA decomposition can be applied to generate spatial components. These spatial components correspond to the different sources of signal, and every frame is represented as a linear combination of these different spatial signals.

Some calcium recording processing pipelines use spatial PCA decompositions for the identification and isolation of hemodynamics in recordings [34, 28]; however, since PCA is optimized for compressing data by pulling out the maximum amount of signal in as few

components as possible. This produces components that are a mixture of signal and artifact (fig. 1.4, top). Therefore, removing hemodynamic artifacts with Principal Component Analysis inevitably removes signal from the data as well. To reduce the intermixing of components, we have implemented an alternative ICA-based eigendecomposition (detailed in sec. 2.1). This assumption results in a more complete isolation of individual signals from an intermixed dataset, resulting in superior signal separation for artifact filtration (fig. 1.4, bottom).

ICA decomposition results in a more complete separation of neural signal and artifact than other decomposition methods. As a result, when the artifact components are removed from the video, filtration is achieved with less remaining artifacts, and with less alteration of neural signal. ICA is commonly used for fMRI and EEG data [17], but had not previously been applied successfully for mesoscale calcium imaging.

However, the lower spatiotemporal resolution of fMRI and data leads to different results for ICA segmentation. For EEG data, artifacts are often temporal, such as pulse readings contaminating electrical activation signals. [43] The scale of fMRI data results in components that correspond to intrinsic connectivity networks, rather than identification of individual areas, and artifacts that represent global artifact effects, rather than spatially localized effects [45].

In this sense, ICA segmentation can be a marker of signal quality, and is sensitive to the signal sources in a given dataset. Since EEG and fMRI datasets are relatively limited in their available spatial resolution, the effect of signal recording resolution on ICA decomposition of neural signal sources had not yet been studied.

1.5 Determining Recording Quality

Researchers have been recording calcium dynamics at frame rates ranging from 5-100Hz [54, 4, 69], often with little to no justification for recording at different spatial resolutions. In addition, recording resolution highly varies between different setups, but is often in the range of 512x512 (0.2 MP) to 256x256 (0.06 MP) pixels for the entire cortical surface, and is often further spatially binned [54, 4, 5]. Selection of spatial resolution is often dependent on the video observer's perceived quality of the image, rather than a quantified comparison of signal content.

We recorded all data at much higher resolution (2560x2160 px, 5.5MP; 10Hz). In sec. 2.2, we downsampled to determine how spatial and temporal resolution affected signal quality and source separation in our dataset. We believe this thorough approach to determine signal quality should be used to optimize neural signal recording on any recording system.

1.6 Spatial Reference in the Developing Cortex

For adult calcium imaging recordings, it is common to use sensory stimulation to identify specific regions in the neocortex, and align a reference map based on the location of these defined regions [70, 54, 5]. Even if these maps are reliable for the location of primary sensory areas, they often lack specificity in higher order areas, or even completely lack sub-regional divisions [56]. This is especially true in areas with a high degree of interconnectedness, with overlapping functionality, such as motor cortex [51]. Moreover, there is growing evidence that the shape and location of higher order regions vary from subject to subject [75, 19].

In addition, functional regions of the mouse brain are not well characterized during development. Stimulation-based mapping is not necessarily reliable, since sensory systems are not necessarily responsive to external stimuli throughout development. During this time, external sensory stimuli do not reliably provide functional input for developing primary or secondary visual and auditory cortical areas [8], and cannot be used for mapping regions. Even if these regions can be reliably mapped, there is no reference map established for the developing mouse cortex, so a full cortical map including higher order areas cannot be defined based on stimulation during development.

Furthermore, mice with genetic abnormalities may have different cortical maps. For example, Mouse lines such as Lmo4 [9], Fgf8, or Emx2 [20] where the cortical lobes are proportioned differently within the neocortex. In addition, the BTBR (Black and Tan Brachyury) is a mouse line that is commonly used for its lack of a corpus callosum and its autism-like phenotypes, also displays different cortical proportions [15, 65].

An alternative to functional mapping is histological tissue post-processing and imaging. Histological procedures can outline primary sensory areas, but do not provide information about the development of higher-order regions of the brain. Additionally, Histological mapping procedures are time-consuming to implement across a developmental timeline, and are difficult to precisely align to functional imaging. It is not clear how functional maps may be scaled differently in animals with different genetic or background abnormalities, and creating maps based on histology for every mouse line throughout development would be tedious, time consuming, difficult to align with functional recordings, and does not provide information about higher-order areas.

When reference cortical maps are used for video segmentation and signal extraction, improper alignment or misinformed regional boundaries can lead to a loss in dynamic range between signals across a regional border. By improperly averaging together signals from different sources, the extracted time series becomes intermixed, and is thus less useful for analyses, especially for spatial clustering.

Some researchers spatially downsample their videos until they can easily compute pixelwise correlations for various seeds across the cortex [70, 54], however these grid-based downsampling systems treat the cortex as a uniform sheet, and are not optimized for functional recordings either.

We know that most cortical processing happens within columnar structures, that vary in local and long-range connectivity. These columns are highly organized within localized neighborhoods, in nested hierarchical systems [52, 51]. Cortical neurons are not equally connected to their neighbors–the network shifts at conserved functional borders. For example, the

mouse forelimb and hind limb are neighbors in the sensory topological map, yet are not very functionally or anatomically linked for their close proximity [52, 53, 56].

Different recording methods offer insights into various levels of these hierarchical networks. Thus, to extract the most information from a recorded dataset, the level of parcellation must reflect the quality and sources present within the data. Thus we need a flexible method that respects functional boundaries of the cortex, is flexible to age, genotype, and individual variation. Previously, a method for creating optimized cortical maps for any functional video data was missing. In sec. 2.3, we present a processing pipeline using ICA that optimizes signal segmentation and functional mapping of the neocortex.

We aim to address these gaps in mesoscale imaging, and develop a computational method that filters the data (sec. 2.1), assesses input data quality (sec. 2.2), and spatially segments (sec. 2.3) the cortex into discrete regions for optimal time course extraction (sec. 2.4). To make these methods as flexible as possible and as applicable to both developmental and adult mesoscale imaging, the following experiments all use P21 mice. In sec. 3.2, we explore the application of these methods to developing mouse cortex.

Chapter 2

ICA-based Filtering and Segmentation

We have developed an algorithm that filters and compresses mesoscale imaging data using Independent Component Analysis. Metrics that summarize the quality and separation of input signal are automatically generated. These metrics can be used for various purposes, including determining the optimal spatiotemporal resolution or recording duration for data collection.

Additionally, a data-driven segmentation of the cortical surface is generated from the ICA decomposition. This method of segmentation optimizes the information represented in a reduced set of timecourses, while preserving the spatial structure and dynamic range of the original dataset.

2.1 ICA Filtration

A spatial ICA decomposition produces a series of spatial components and a mixing matrix, representative of the component's influence at each frame in the video (fig. 2.1). The components are sorted by influence over the video variance, and flipped so that they all represent positive effects (see methods: sec. 7). The independent components can be sorted into 3 major categories based on their spatiotemporal properties: signal components, artifact components, and noise components (not shown).

Signal components represent a distinct area of cortical tissue, which we refer to as it's cortical domain. The spatial morphology of these signal components can vary greatly. In spatial domain and eccentricity. Occasionally signal components can also contain a secondary or tertiary domain, with similar enough activation patterns to be identified as a single neural component. The second signal component appears to represent a higher order visual network, with multiple domains on the left hemisphere, and a small mirrored domain on the right hemisphere.

Artifact components can take many forms, including various blood vessels and arteries, movement artifacts, optical surface artifacts, etc. The top two artifact examples likely represent hemodynamics from the superior saggital sinus vein with the bottom artifact likely represents blood flow through the middle cerebral artery [73]. A very high resolution map of the vessel patterns can potentially be rebuilt from these components, with branching structures as small as $\sim 12\mu m$ in diameter (shown above in zoom panel). Noise components lack a spatial domain, and have little to no temporal structure. Noise sorting is further detailed in sec. 2.2. Signal and artifact components can be sorted manually in graphical user interface (fig. 5.1) or

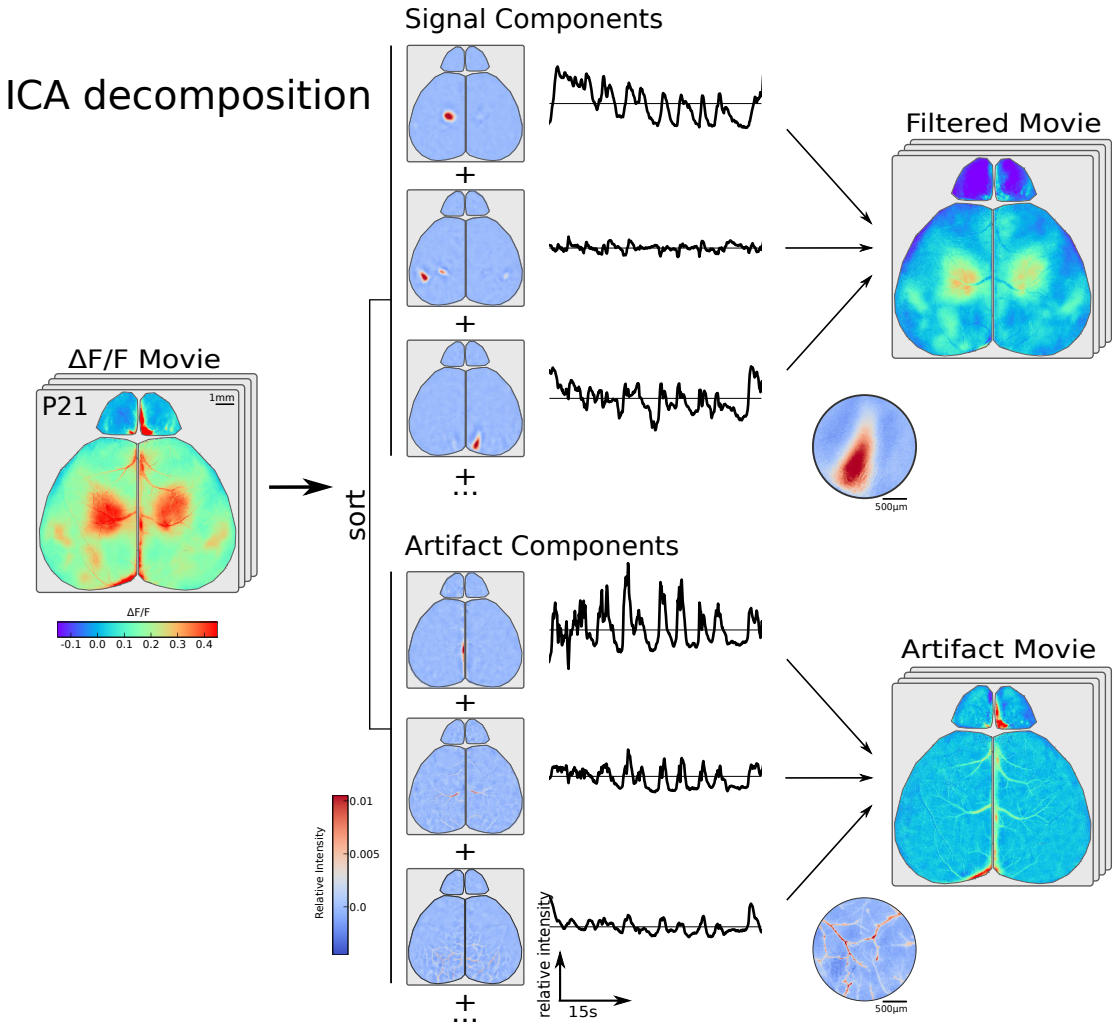

Figure 2.1: Independent Component Analysis (ICA) video decomposition schematic. A dF/F movie is decomposed into a series of statistically independent components that are either neural, artifact, or noise associated (not displayed). Each component has an associated time course from the ICA mixing matrix. Signal components can be rebuilt into a filtered movie. Alternatively, artifact components can be rebuilt into an artifact movie. Circular panels show higher resolution spatial structure in example components.

with a machine learning classifier (sec. 5.4).

Video data can be reconstructed using any combination of these components. In particular, a filtered video can be constructed by excluding all artifact components. The noise components can be additionally excluded for denoising video data with high spatial noise. For our analyses, we did not exclude ICA noise components when rebuilding. The artifact movie can also be reconstructed to verify that desired signal was not removed with the artifact filtration (see sec. 5.3).

2.2 Noise Sorting and Resolution Analysis

Non-noise (signal and artifact) components can be separated from noise components from their visual differences, as well as from their log temporal variances, or lag-1 autocorrelations. Non-noise components have spatial structure and a high lag-1 autocorrelation, corresponding to waves of calcium activity in a specific location. Conversely, noise components are highly dispersed across the cortex, and have a low lag-1 autocorrelation. The lag-1 autocorrelation metrics from these two groups are so polarized, that it is straightforward to separate these populations by their lag-1 autocorrelation alone (fig. 2.2a).

To automate this sorting process, a two-peak kernel density estimator (KDE) was fit to the histogram of lag-1 autocorrelation data. The KDE distribution is an easy way to summarize the two major peaks, as well as the minima between them, defined as the noise cutoff.
The locations of these peaks, and the minima between them is highly stable across our 8 P21 test recordings (see Methods–sec. 7–for animal details). We found the non-noise peak ($p1$) at an autocorrelation of 0.94 ± 0.01, and a noise peak ($p2$) at 0.13 ± 0.01. The central cutoff minima was slightly more variable, with an autocorrelation value of 0.61 ± 0.05. A high degree of separation between these peaks ($d_{p-p} = 0.82 \pm 0.01$; $p < 0.001$) suggests that the signal and noise signal sources were completely separated, and thus all signal sources were distinctly identified.

To test how this separation is affected by spatiotemporal resolution and video duration, we altered properties of the input video and observed its effects on the quality of signal separation through lag-1 autocorrelation distributions. Reducing the spatial resolution resulted

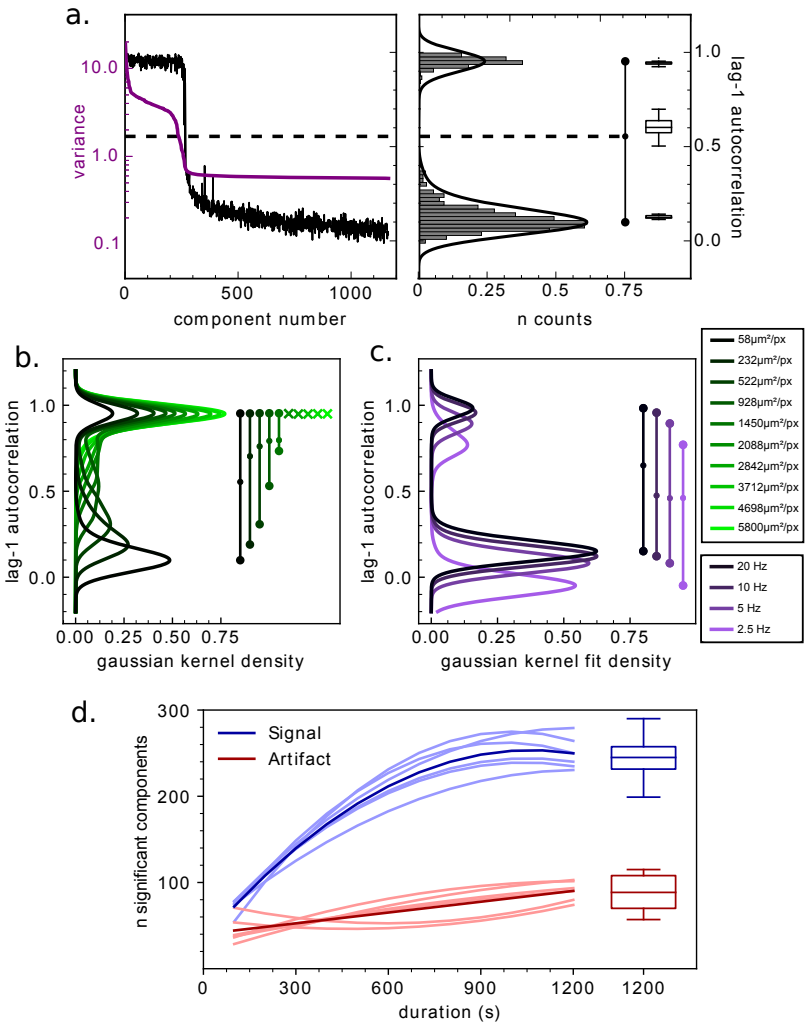

Figure 2.2: a) Distributions for lag-1 autocorrelation (black) and temporal variance (purple) are displayed for components 1-1200. A dotted line representing the cutoff determined from the distribution in the right panel. In the right panel, a horizontal histogram on the lag-1 autocorrelation with a two-peaked kernel density estimator (KDE) fit reveals a two peaked-histogram, summarized by a barbell line. Group data for each peak, as well as the central cutoff value is summarized by the boxplots on the right (n=16). b) 2-peaked KDE fits of horizontal histogram distributions under various spatial downsampling conditions, with barbell summary lines on the right. After spatial resolution decreases beyond 2088um^2/px, this two peak structure collapses, and an x denotes the primary histogram peak. c) 2-peaked KDE fits of horizontal histogram distributions under various temporal downsampling conditions, with barbell summary lines on the right. d) Component stabilization for different length video subsets of six 20-minute video samples. Individual thin lines show polynomial fit to signal or artifact components under each time condition. Thick lines denote the curve fit of the mean number of components in each category across these six experiments. The group distribution of components at 20 minutes is summarized by the boxplot on the right (n=16).

in a steady decrease in peak separation, until the dual peaked structure collapsed at a resolution of $2088\mu m^2$/px (fig. 2.2b).

Increasing the sampling rate above 10Hz showed little to no effect on the peak to peak distance ($\Delta_{p-p} < 0.01$), and a slight decrease in the autocorrelation of the primary peak ($\Delta_{p1} = 0.03$), but temporal downsampling below 10Hz resulted in a shifting of the signal and noise peaks ($\Delta_{p1} = 0.06$), and a reduction in the peak to peak distance ($\Delta_{p-p} = 0.02$). Since temporal dynamics of GCaMP6s are not expected to 8Hz (CITE), this result agrees with our original choice of temporal sampling at 10Hz. Together, this suggests that the separation quality of our captured dynamics are highly sensitive to spatial resolution, and not as sensitive to temporal resolution. We considered collecting spatial samples higher than our current resolution of $\sim 6.75\mu m$/px, but it was not possible with our current camera setup, and would result in extremely large datasets.

To determine the ideal duration of video collected, we calculated the number of significant signal and noise components for various video durations (for component sorting, see sec. 5.4). We found that a for ICA decompositions on an activity patterns from a P21 mouse, the number of noise and artifact components leveled off by 20 minutes of recording. Population analyses showed that this number was highly similar among P21 mice (n signal components: 244 ± 25.7; n artifact components: 87.2 ± 20.7). However, when we tested duration on other developmental time points, the required duration for stabilization varied (see sec. 3.2).

2.3 Generating Domain Maps

In addition to their applications for filtering, the components also are a rich source of information about spatial distributions of signal within the cortex. Components across the cortex show a wide diversity of spatial characteristics, and represent a detected independent unit of signal. We use the spatial domain footprints of each signal component to create a create a data-driven 'domain map' of the cortical surface by taking a maximum projection through each component layer (fig. 2.3a; for details, see methods:sec. 7). For analysis, 8 maps were created, with an average of 230 ± 14 detected domains. Domains were then sorted into regions, by a combination of network analysis, reference comparison, and manual sorting (see sec. 5.6 for additional information).

Domains did not exhibit uniform spatial characteristics across the neocortex. Different detected regions have different spatial characteristics such as area (ANOVA F=161.6, $p < 0.001$), as well as eccentricity (ANOVA F=47.4, $p < 0.001$). Generally, higher order and motor regions (R, V+, Ss, Mm, Ml) had larger domains than primary sensory areas (V1, A, Sc, Sb, S) (ANOVA F=558.4, $p < 0.001$), and also exhibited higher eccentricity (ANOVA F=199.6, $p < 0.001$).

To test the meaning of these maps, a series of comparisons were performed. Pairs of maps were overlayed on top of each other (fig. 2.3), and every domain was compared to its nearest domain in the comparison map (see methods for details: sec. 7). The Jaccard overlap was calculated for each of these domain pairs, and quantified for each pair of map comparisons. For a null hypothesis, randomly generated Voronoi maps were also compared (see methods for

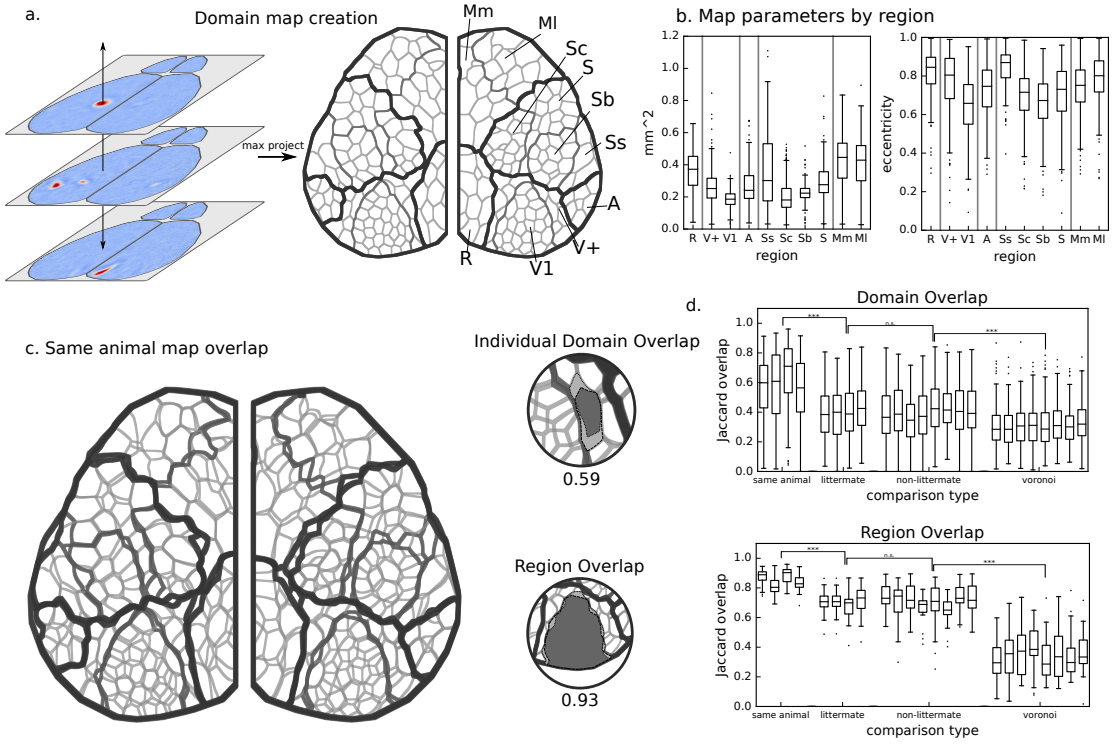

Figure 2.3: a) Schematic of domain map creation. A maximum projection is taken through each blurred signal component to form a domain map. Cortical domains are assigned to one of 10 identified regions, with region divisions denoted by the bold lines. b) Domain area and eccentricity by region. Population analysis of distribution of spatial characteristics individual domains within defined regions across multiple recordings (n=16). c) Example overlay of one domain map on another from the same animal. Individual domain or region overlap is calculated using the Jaccard index (intersect / union). d) Population analysis of the Jaccard index for domain and region overlap comparisons. Maps are generated from a different recording on the same animal, a littermate, a non-littermate, or a randomly generated voronoi map. Significance is calculated using a ANOVA two-way t tests with Holm-Sidak correction.

details: sec. 7).

Maps generated from different recordings from the same animal were found to be highly overlapping, and hence more similar (fig. 2.3, top; $p < 0.001$). There was no significant difference in comparisons between littermates vs non littermates. Non-littermate map comparisons were significantly more similar to each other than to voronoi maps ($p < 0.001$).

We additionally quantified whether detected regions were similar across map comparisons. We again found that comparison between maps from the same animal were highly similar (fig. 2.3, bottom; $p < 0.001$), no difference was found between littermates and non-littermates, and comparisons between different animals were significantly more similar than a comparison between a region map and a randomly generated voronoi map ($p < 0.001$). In summary, regions and domains are similar between recordings either in the same or on different animals, compared to a null map distribution.

2.4 Extracting Optimized Time Courses

At full resolution, there are approximately 1.5 million pixels along the surface of the cortex–far too many sources for most network analyses, which work best on 10-300 time series. We propose that these data-driven domain maps are an optimal method for extracting time courses from the cortical surface. Time series were extracted by averaging the filtered movie under each domain. This results in a series of $\sim$ 230 time series per video recording, representing a $\sim$ 6,500-fold reduction in size (fig. 2.4a).

To test how well the full filtered video was represented in these time series, we rebuilt 'mosaic movies', where each domain is represented by its mean extracted signal at any given time point (fig. 2.4b). Comparing the borders of the large higher order visual activation, one can see visually that the data appears more distorted in the voronoi and grid. To numerically compare whether this method of time course extraction was superior to alternate methods, we also compared mosaic movies rebuilt with either grid or voronoi maps.

The residuals between the mosaic movies and the filtered movies were compared to the total spatial variation in the filtered movie to quantify the amount of total signal represented by the extracted time courses (fig. 2.4c, left). In nearly every experiment, the optimized domain map performed better than any other time course extraction method, and accounted for 68 ± 1.2 % of the total spatial signal in the filtered video (n=8).

Domain maps generated from different videos from the same animal performed nearly as well as the optimized domain maps created from the video compared (fig. 2.4c, right). These maps performed significantly better ($p = 0.01$) than the grid maps, and much better than the

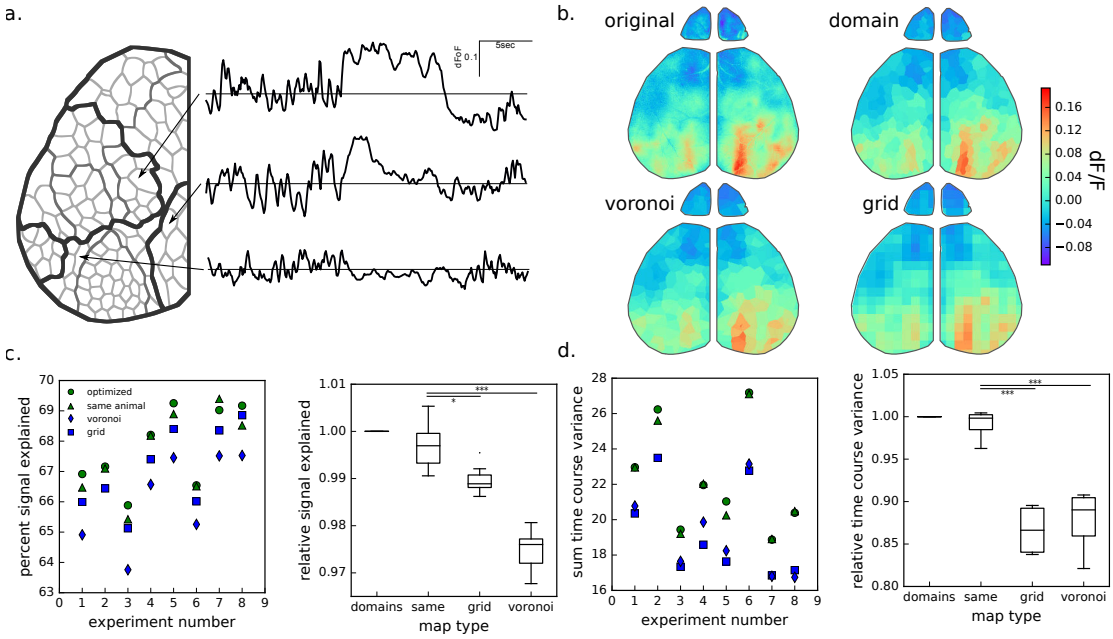

Figure 2.4: a) Schematic of mean time courses extracted from different domains generated from a domain map. Different dynamics are observed from time courses extracted from different domains. b) Example of mosaic movie frames rebuilt from just the information stored in mean domain time courses. The filtered movie is replicated in the upper left corner. The upper right image corresponds to the representation of the same video frame rebuilt from the extracted domain time series. The left and right bottom panels correspond to the representation of the same video frame rebuilt from the extracted domain time series from either voronoi or grid maps, respectively. c) Percent total signal of the filtered video represented by extracted time courses. Percent of overall video signal captured in domain maps was calculated from one video from each animal (green circle; n=8), and compared to a map generated from a separate video calculated from the same animal (green triangle). Percent total signal represented by time courses extracted from grid (blue square) or randomly generated (blue diamond) maps were compared as controls. In the right panel, the percent signal relative to the domain map percent signal was summarized in a box plot. d) Variation between time courses extracted with each map method was then quantified as a sum signal variation for each experiment. In the right panel, the sum signal variation for each comparison map relative to the optimized domain map sum signal variation was summarized in a box plot.

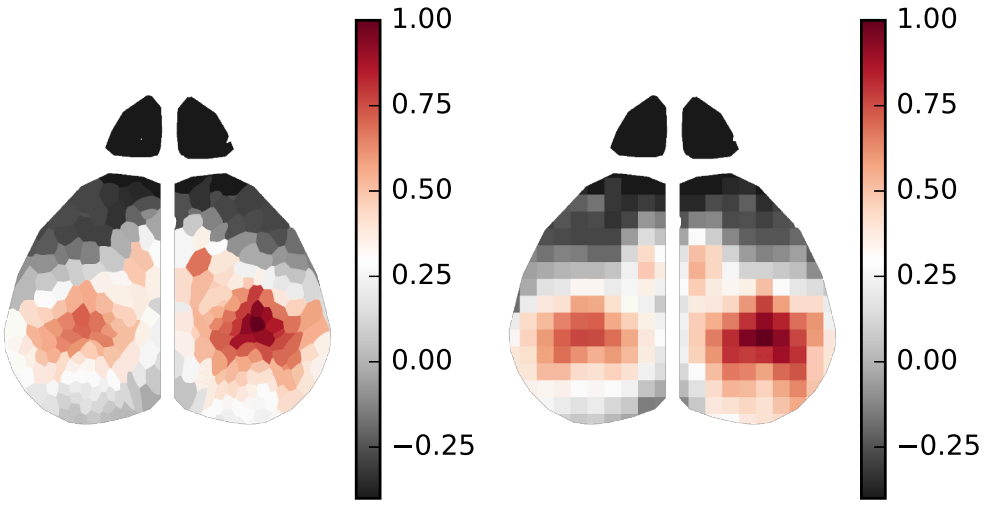

Figure 2.5: Example correlation maps generated from timecourses extracted from domain map segmentations (left) compared to grid segmentations (right).

voronoi maps ($p < 0.001$).

Compared to saving the full ICA compressed dataset, saving these extracted time courses and all associated metadata results in a file size of $\sim$ 100MB, for a $\sim$ 60 fold additional compression. One potential benefit to accounting for the underlying regions of the brain while extracting time courses is reducing the amount of times that an extracted mean signal is diluted by signal from a neighboring region. Properly restricting time series extraction to statistically independent units should enhance the dynamic range between extracted time series.

To test whether domain maps extracted time courses better extract the full range of variation in the cortical surface, we compared the total variation between time courses rebuilt under domain maps from the same video, same animal, or control grid and voronoi maps (fig. 2.4d, left). When normalized to the performance of the optimized domain map, domain maps from the same animal again had similar performance, but grid and voronoi maps per-

formed significantly worse ($p < 0.001$; fig. 2.4d, right). There is a $\sim 15\%$ reduction in signal variation in grid or voronoi maps compared to domain map extracted time courses.

The example shown in fig. 2.5 shows a single example of correlation maps generated from time series from either the domain map ([fig:correlation-example], left), or the grid map ([fig:correlation-example], right). The borders of sensory regions such as primary visual cortex (V1) and auditory cortex are visible in the domain map borders, but are blurred out in the grid representation. These clean borders for time series extraction are invaluable for network analyses.

2.5 Conclusions

We have shown that high resolution imaging of mesoscale cortical calcium dynamics combined with data-driven decomposition using ICA results in an optimized extraction of neural source signals. We demonstrate that these methods provide precise isolation and filtration of video artifacts due to movement, optical deformations, or blood vessel dynamics while recovering cortical source signals with minimal alteration.

Previous attempts at ICA filtration and segmentation in neuroscience have been less than ideal for identifying unmixed sources [21, 45, 14, 44, 55], likely due to the insufficient spatial sampling required for complete signal separation, and insufficient components extraction through ICA.

Here we report that signal separation from mesoscale calcium dynamics recorded across the cortical surface is the most complete at the highest spatial resolution tested (2560x2160 px; 5.5MP). Temporal resolution had less of an effect on ICA signal separation; we found that a 10Hz sampling rate was sufficient. These metrics for signal quality are automatically generated by our algorithm, and can be used to optimize signal collection on any given experimental setup (age, genotype, recording duration, resolution, etc). The number of components identified is highly stable after recording sufficient duration of dynamics, and provides a quick metric for spatial complexity of neural signal across the neocortex.

We further demonstrate an ICA-based method for using these components to perform a data-driven mapping of the captured cortical dynamics, resulting in a superior isolation of the various signal sources on the cortical surface. Using optimized signal extraction results in a

higher dynamic range of extracted time series, and thus would produce correlation maps and network analyses with enhanced separation between neighboring regions. Stimulation experiments or anatomical post-processing of brain tissue could provide insights on how well the sorted domain regions correspond with known functional units of the brain. If these maps align well to anatomical regions, they can additionally be used as a reference for in-vivo cortical mapping, with potential applications in live feedback or targeted injection experiments.

An additional benefit is the highly compressed data format. The original or video can be rebuilt with relatively little loss of information from a reduced set of ICA components, for a $\sim$ 10x reduction in file size. This results in a representation of a high resolution video that is much more manageable to work with on a local computer.

These methods are flexible for any age, genotype, or drug condition. This method of segmentation may provide insights into the underlying cortical anatomy; however further experiments are necessary for confirmation. We believe these methods can be used to generate metrics of spatial complexity. We further explore this through applied analyses in the following chapter.

Chapter 3

Applied Analyses

To test our artifact filtration and the usefulness of our spatial segmentation and signal extraction for generating useful analytics on a wide variety of datasets, we applied our methods to various conditions that alter spatiotemporal properties of cortical networks. First, we present the effects of isoflurane anesthesia on cortical networks. Then, we investigated spatial properties of the cortex during the first two weeks postnatally, from birth to eye opening.

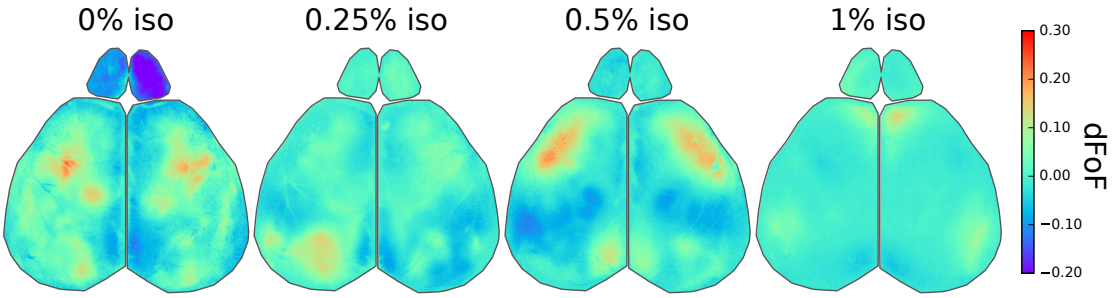

Figure 3.1: Example dFoF frames from videos sequentially recorded from the same animal under different isoflurane conditions.

3.1 Anesthesia

Surgical anesthesia in veterinary or clinical settings often use inhaled halogenated ethers, such as isoflurane, desflurane, or sevoflurane. Like many general anesthetics, the mechanism of action in vivo for isoflurane is not well understood. We do know that it is a positive allosteric modulator of GABA-A receptors [33, 35]. It has been shown that anesthesia affects information transfer between primary and higher order regions, and affects neural correlation patterns [71, 49, 72, 40].

This research has been conducted to analyze regional network changes using defined coordinate systems, but little has been done to analyze local circuit changes within functional units. To test whether our segmentation methods could identify differential spatial structure under different anesthesia conditions, we filtered videos, then applied our segmentation analysis to 3 different P21 mice under 4 increasing isoflurane levels (see methods:sec. 7).

Visually, videos recorded from animals under 0.25% isoflurane appeared fairly similar to videos recorded under 0% isoflurane fig. 3.1. However by 0.5% isoflurane, there seemed to be a loss of spatial complexity; patterns appeared highly symmetric. By 1% isoflurane, most

localized activity within the cortex was lost–there were some weak symmetric activations, such as the one shown in fig. 3.1.

Before doing any analysis, we first tested whether isoflurane affected the duration of time required for component stabilization, by performing a component stabilization test (fig. 3.2a), similar to the signal/artifact test applied in sec. 2.2. Surprisingly, we found that the addition of isoflurane did not seem to extend the duration of video data required for component stabilization. However, it did decrease the overall number of components detected. Even isoflurane concentrations as low as 0.25% had an impact on the number of components detected ($p < 0.01$).

The number of detected domains increased slightly at 0.25%, but was not statistically significant. Further increases in domains with higher isoflurane levels were all significant ($p < 0.001$). Number of domains detected for each isoflurane condition were as follows: (0%: 230 ± 14 domains; 0.25%: 208 ± 20 domains; 0.5%: 129 ± 10 domains; 1%: 67 ± 8 domains). The average size of a domain slightly increased between 0% and 0.25% isoflurane ($p < 0.001$). As isoflurane was further increased, the size of the average domain increased even more ($p < 0.001$; fig. 3.2b).

The domain maps generated for each of these conditions is displayed in (fig. 3.2c). By 1% isoflurane, the domain map was severely impacted, and large portions of the body core were not detected as belonging to any component at all. This suggests that these regions were rarely if ever active during the 20-minute video segment. Since these maps were so impacted by isoflurane, we applied the 0% map to the video under each condition for time series extraction for further analyses. In fig. 3.2d, we verified that this loss of structure was due to very little

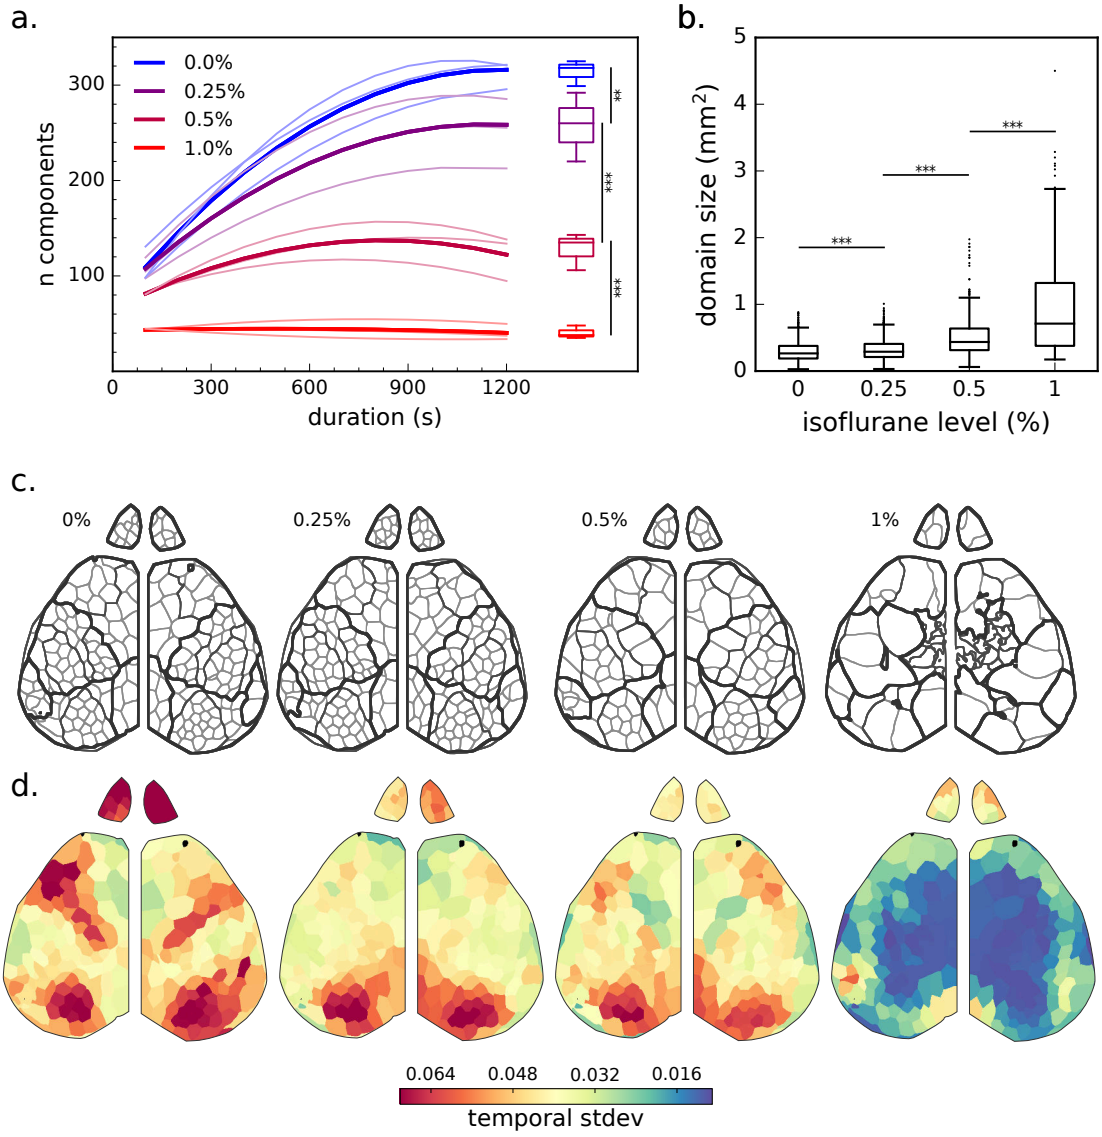

Figure 3.2: Spatial analysis of cortical activity under anesthesia. a) Component stabilization curves under various isoflurane conditions. Duration of the video is varied by 100s, and the number of significant components are compared. Light lines represent a 2nd degree polynomial approximation of the number of components detected for each animal under each condition. Thick lines represent a 2nd degree polynomial fit to all data points under each given isoflurane condition. Distribution of non-noise components for each isoflurane condition is summarized by the box plot on the right. b) Size of domain detected under each isoflurane condition. c) Example domain maps generated from a single animal under each isoflurane conditions. d) Example domain time series standard deviation for a single animal under each isoflurane condition, generated from 0% domain map.

signal over the somatosensory core area. As expected, The activity level, measured by the time series standard deviation from each domain, reported very little signal in somatosensory core areas.

To test how temporal properties differed under these isoflurane conditions, we extracted time series from cortical domains calculated from the 0% isoflurane video, and compared wavelet-based temporal metrics (see methods: sec. 7) for time series extracted under each of the 4 isoflurane conditions (fig. 3.3).

When we looked on the population level, we again found that the time series standard deviation significantly slighly decreased as with light anesthesia (fig. 3.3a; $p < 0.05$). As isoflurane further increased, the temporal standard deviation dropped even lower ($p < 0.001$). The range of frequencies with significant power for any given domain time series (fig. 3.3b) decreased under higher isoflurane ($p < 0.01$).

Under light isoflurane, domains varied greatly in their dominant frequency (fig. 3.3c), as measured by the maximum wavelet signal to noise ratio (see sec. 7). However under higher isoflurane levels, this diverse structure collapsed into a narrow range of frequencies (fig. 3.3c; $p < 0.01$). This suggests that in addition to its loss of spatial structure, the temporal dynamics across the cortex were also simplified.

We next performed a network analysis to further explore how the cortical functional structure is affected by isoflurane anesthesia. First, we generated network connectivity maps. Since the number of domains detected per region varied between animals, we calculated the percentage of highly correlated (wavelet coherence > 0.4) domains for each pair of region comparisons (fig. 3.3c).

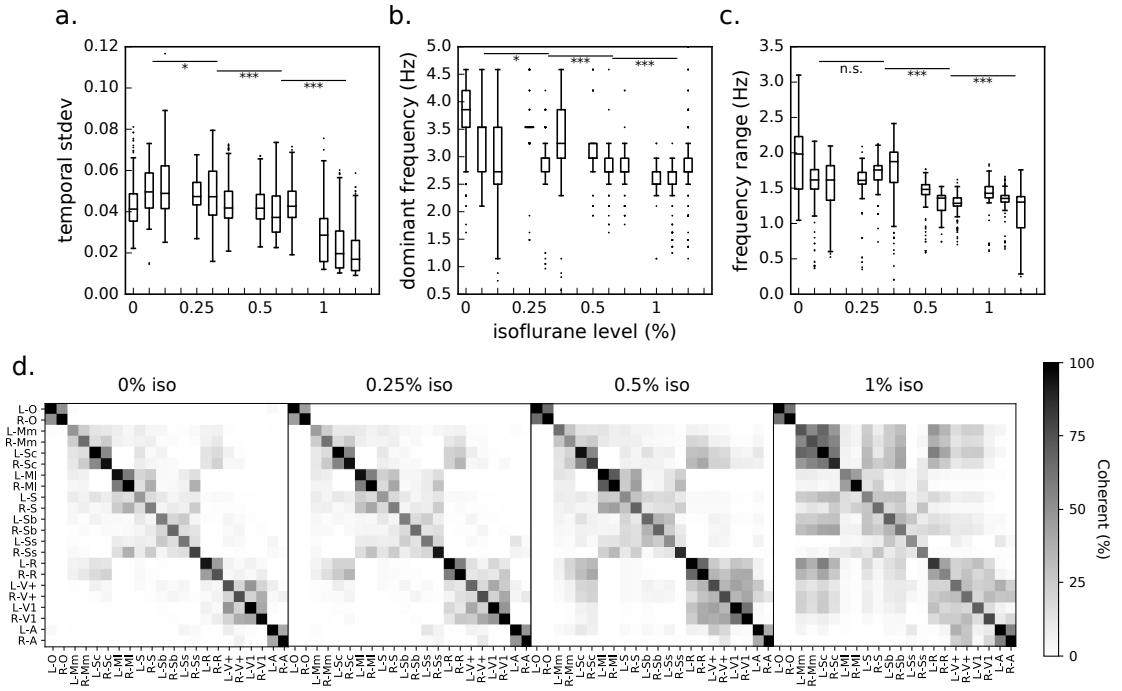

Figure 3.3: The effects of anesthesia on time series and network connectivity under isoflurane anesthesia. a) Boxplot of the temporal standard deviation of domain time series under different isoflurane conditions. Drug animals 1, 2, and 3 are plotted on the left, center, and right, respectively. Each boxplot represents the distribution of events detected from time series extracted from domains within a single animal. b) Boxplot of range of significant wavelet frequencies detected from domain time series under different isoflurane conditions. c) Boxplot of the dominant detected frequency in domain time series under different isoflurane conditions. d) Average intra-regional connectivity network for animals under each isoflurane condition. Each comparison represents the percent of domains within each region that were highly correlated (coherence > 0.4) with regions in the comparison target (n=3 animals per condition). e) Comparison of average network strength between different anesthesia levels. The difference was taken between network maps of increasing anestheia levels. Region comparisons in red represent an increase in connectivity as isoflurane was increased. Regions in blue represent a decrease in connectivity as isoflurane was increased.

Interdomain connectivity, represented by the diagonal under each isoflurane condition, was generally quite high, shown in the example network cross-coherence in fig. 3.3d. The overall coherence (as measured by mean percent coherent regions) increased as isoflurane levels increased (ANOVA F=13.70, $p < 0.01$).

The localized network structure as seen in the 0% isoflurane example became blurred and distributed globally across the cortex as anesthesia increased. Even though the anatomical structure has not been altered, the increase in coherence across multiple networks suggests that the 0% domain map is not an ideal way to represent the new spatiotemporal structure under isoflurane.

In conclusion, we found that even sub-surgical level of isoflurane anesthesia drastically affects spatial, temporal, and network interaction dynamics across the cerebral cortex. The number of identified ICA components, and the resulting domain map are reflective of the relative spatiotemporal complexity of the underlying neural signal. When temporal and spatial network dynamics simplify under increasing isoflurane, there is a reduction in unique signal sources. As a result, the ICA components and domain maps simplify as well, and reduce in number. Thus, even extremely light, sub-surgical levels of isoflurane reduced the spatial complexity of the cortical domain mosaic, showing that GABAergic neurotransmission shapes functional group dynamics in neocortex.

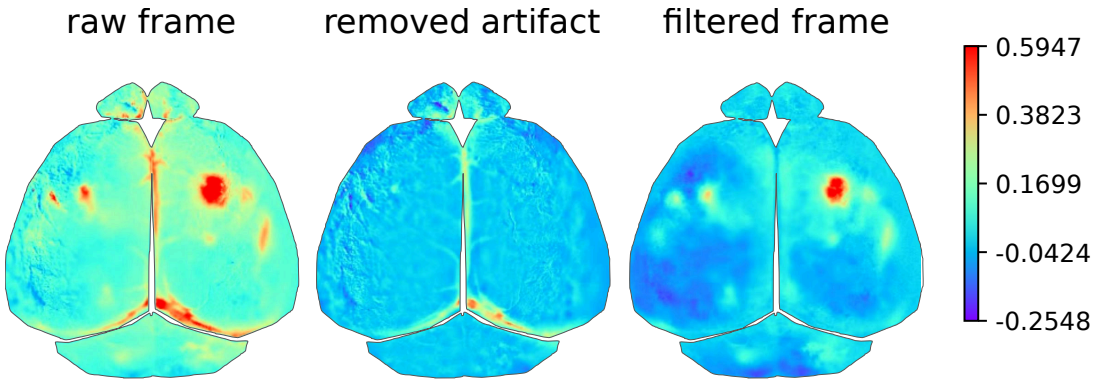

Figure 3.4: Filtering results for a P2 video recording. The raw dFoF video (left) is ICA filtered, to remove detected optical, surface blood, and hemodynamic artifacts. The center panel show what information was removed from the given frame by the ICA filter. The right panel shows the filtered video at the given frame. In this frame, the olfactory bulbs and superior colliculus are also shown, but were not used in this section's analysis.

3.2 Development

We also applied our filtering and segmentation methods on animals across early postnatal development. Filtering was highly successful for our young mice. Before the skull is developed enough to hold its shape under pressure, experiments on these young mice often have so many blood and optical surface artifacts that the data is unusable. In fig. 3.4, we provide an example frame of a filtering result on a P2 mouse recording. The filtering successfully removed hemodynamic, surface blood, and surface optical artifacts, thus rescuing the useability of this recording. Without this artifact filtering, we could not perform analyses on animals this young.

Similar to our P21 data, signal components could be identified throughout a range of developmental ages (fig. 3.5a). The spatial characteristics of these components varied somewhat, along with the number of multi-domains detected per signal component.

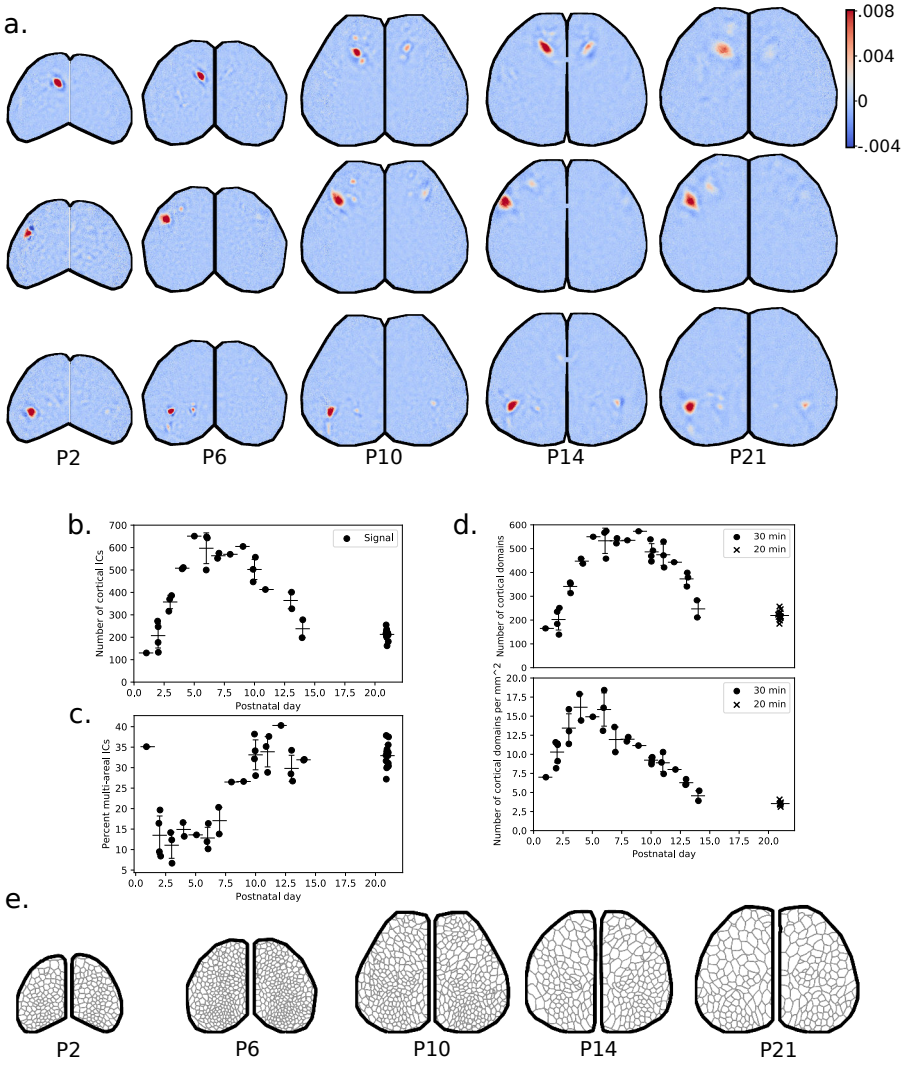

Figure 3.5: ICA metrics summarizing spatial changes across the cortex during development. a) Example ICA components extracted from different regions throughout select development time points (P2, P6, P10, P14, and P21). Each column corresponds to examples selected from a single animal. Rows correspond to similar components in distinct regions selected for comparison. b) Number of cortical signal components detected from 30-minute recordings from various mice throughout early postnatal development. Points are recordings from individual animals, with the group mean and standard deviation marked with a line and whiskers. c) Percentage signal components detected with multi-domains detected from 30-minute recordings from various mice throughout early postnatal development. d) the number of cortical domains extracted from ICA decompositions from 20 and 30 minute datasets across development. In the bottom panel, this relationship is adjusted to calculate the number of domains per mm^2. P21 data is presented as a 20-minute recording, since 30-minute decompositions were not possible. d) Example domain maps calculated from example animals across development. The same animals are used as presented in (a).

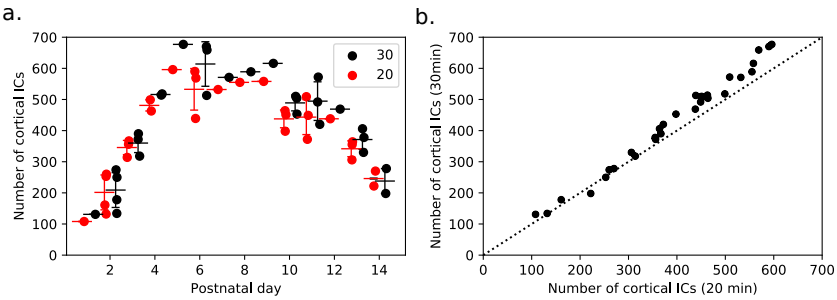

Figure 3.6: Determining the ideal length of recording for videos from different age mice. a) Number of cortical signal components detected in a 20-minute video (red) compared to the number of cortical signal components detected in a 30-minute video (black) for the first two weeks postnatally. b) Comparison of the number of signal components detected in a 20-minute video compared to the number of signal components detected in a 30-minute video. A dashed line represents the unity relationship, where the number of components are equal in each decomposition.

The number of cortical signal components detected throughout development rose sharply for the first postnatal week, before an inflection point around p6-p8, where the number of signal components decreases throughout the second postnatal week (fig. 3.5b). When we quantified the number of multi-domain signal components, there seem to be two distinct populations. In the first postnatal week, only $\sim$ 15% of components had multiple domains. In the second postnatal week, the number of multiple domain components nearly doubled. This property remained constant at 3 weeks postnatal (P21).

Since this number of resulting components was much higher, we tested the effect of increasing our video duration to 30 minutes. We found that between P5-P12, there was an increase in the number of detected cortical domains when video duration was increased (fig. 3.6). For this reason, we used 30-minute videos for the rest of the development experiements. P21 data contained too many samples to be processed as 30-minute segments, so the data was processed in 20-minute segments.

In addition to the number of signal components detected, we also quantified the density of domains identified by the domain map throughout development. This trend followed the same pattern as the number of independent components (fig. 3.5d, top). This trend remained even after correcting for cortical area, though the inflection point shifts slightly towards P5 or P6 (fig. 3.5d, bottom). Spatially, the cortical surface appears highly disorganized, with relatively uniform domains until P10, where the adult-like patterns and diversity in domains begin to take shape.

It is known that cortical projection neurons send out extensive projections throughout development, with maximal multi projections around P8 cite{fame_development_2011}. These projections are rapidly tuned through early postnatal development, and reach near adult-like circuitry by P12 [57]. Here, we find nearly identical patterning in our independent component and domain map metrics.

Using the number of identified independent components or domains provides a simple, unbiased metric for spatial complexity.

Since this metric is not dependent on any external references or maps, it is easy to extend this analysis to mice with altered genetic or environmental influences for which reference maps have not been defined. We aim to apply this analysis to other mice such as the acollasal mouse line BTBR [15] or genetically engineered Ephrin mutants [18].

Chapter 4

Conclusions

Here we have introduced data-driven methods for automated artifact filtering, quality analysis, spatial complexity analysis, domain segmentation, and optimized time series extraction from high resolution mesoscale calcium dynamics. These methods are flexible enough to be applied to any mice, regardless of age, drug and genotype conditions, without any manual parameter adjustment.

The automated extraction of time series allows for easy comparison of different conditions, without the need for external maps. These extracted time series are optimized for representing identified functional units in the cortical surface. By increasing the variance between extracted signals, correlation matrices thus have a higher dynamic range.

Our isoflurane experiments show that under light to moderate sub-surgical levels of anesthesia, the cortical surface undergoes a vast reduction in functional complexity.
Changes in temporal activation patterns explained some of these changes, but changes in spatial metrics were even more extreme, and reflective of the simplification of patterns evident from

visually observing functional group dynamics.

Intracortical projection, and activity-dependent circuit refinement and maturation are key features of early postnatal brain development. This process previously been quantified on a cellular level, but little work has been done studying the refinement of local *functional* circuits in the neocortex. To quantify this, a measure of spatial complexity is necessary. Here, we present the first attempt at quantify spatial functional complexity of the neocortex during development. These results line up with known changes in cellular and network behavior during development.

Now that we have produced a reference for how spatial complexity develops in the neocortex, we can compare how different drugs, environmental effects, or genetic differences affect this process. Better understanding of various processes controlling spatial complexity during development can give us a better idea of how the process of circuit maturation can have lasting effects on brain networks and dynamics. This could potentially lead us to a better understanding of brain disorders with altered circuit and network dynamics, such as autism and schizophrenia.

Chapter 5

Supplemental Information

5.1 Graphical User Interface

We developed a Tkinter-based graphical user interface (GUI) for browsing components and associated metrics and data. Since tkinter is installed natively on most operating systems, this does not require any additional installations. This interface was originally developed for browsing PCA decompositions, and could be also applied to any other eigendecomposition saved in a similar hdf5 formatted file.

The GUI is primarily designed for manual selection or validation of component sorting (fig. 5.1a). Pages of components are presented 15 at a time, and the user is prompted to select components to toggle them between signal (cool/warm colormap) or artifact (black/white colormap). In addition to this first page, six other pages are available for viewing various properties of the dataset. The second page displays the component, timecourse, wavelet decomposition, and frequency power for any given component (fig. 5.1b).

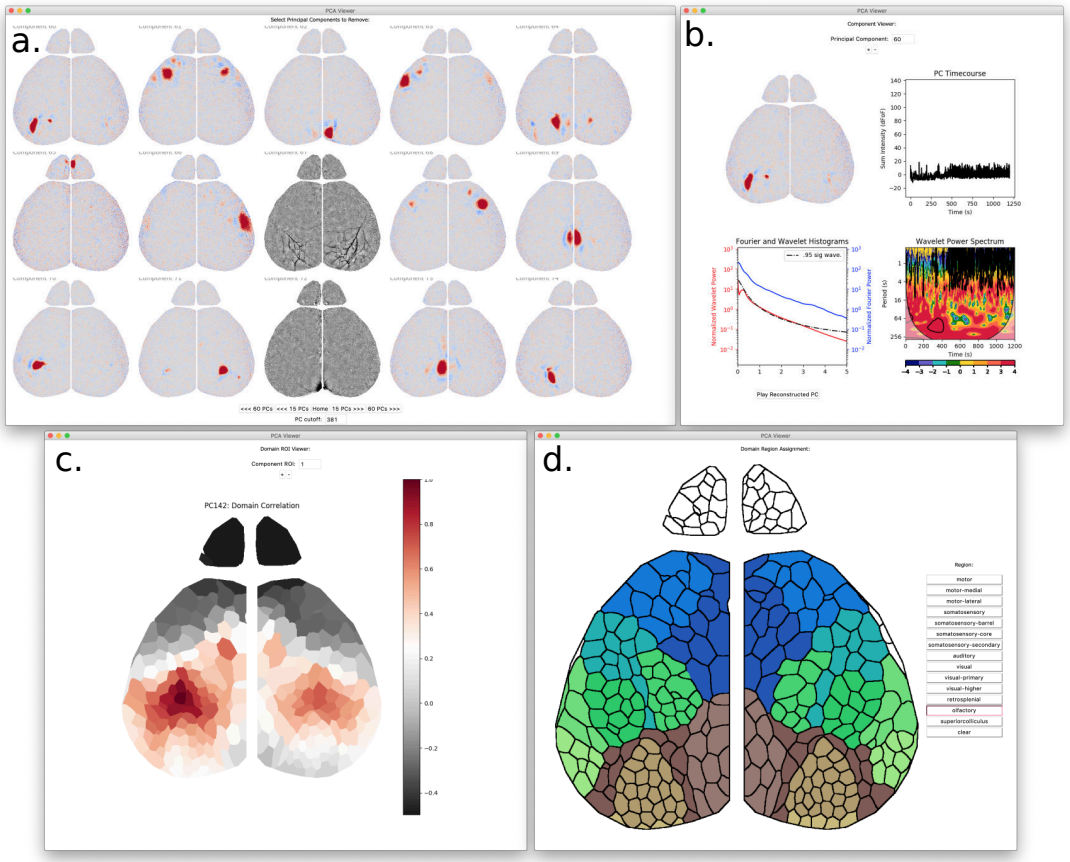

Figure 5.1: A Tkinter-based graphical user interface (GUI) for browsing independent components, and associated data. a) 15 independent components, order 60-74 by variance. Components displayed in grey are selected as artifact either manually or using a machine learning classifier. A click on the display for any given component manually toggles its classification as either signal or artifact associated. Components colored in the cool/warm colormap are signal associated. Components colored in the black/white colormap are artifact associated. Buttons on the bottom panel control GUI movement through the dataset. The text panel at the bottom displays where the index for the signal/noise cutoff. b) The component viewer displays additional temporal metrics about any given component. The top controls allow movement through the dataset by manual scrolling with (+/-) buttons, up/down keys, or through typing a desired component in the text box. PC timecourse displays the mixing matrix timecourse extracted by the IC operation for the given components. The Wavelet power spectrum is displayed in the bottom right, and an integrated wavelet or fourier representation is available on the bottom left. 0.96 significance as estimated by the null hypothesis is displayed as a dot-dash line. c) The domain map correlation page shows the pearson cross correlation value between a selected seed domain and every other domain detected on the cortical surface. The seed domain can be changed through the arrow keys, the (+/-) buttons, or by clicking on a different domain on the displayed domain map. d) The Component region assignment page allows manual region assignment for each domain. After the region is selected from the menu on the right, each domain clicked on the domain map is assigned to that region.

The third page displays the variance of each component, compared to an estimate of the noise floor. The fourth page displays an old method of segmentation where each domain's gaussian cutoff binarized spatial footprint is displayed as an overlapping histogram. This page is only available if the binarized domains are calculated and saved.

The fifth page shows the calculated domain maps as described in sec. 2.3. Selecting any region displays a pearson timecourse cross correlation map between the selected domain and every other domain on the cortical surface (fig. 5.1c).

The sixth page shows a domain map with a selection tool for assigning region identification to each identified domain (fig. 5.1d). The seventh page shows a tool displaying how components were sorted into noise and non-noise groups based on the lag-1 autocorrelation, as described in sec. 2.2.

When the GUI is exited with the save ('s') command, all updates to region assignment or artifact assignment are saved to the data hdf5 file. To save storage space, all calculations happen within the python instance running the GUI, rather than being stored in the hdf5 file. We found a laptop with 16GB of RAM was more than enough to run and operate the GUI.

5.2 Mean Filtration

ICA decomposition must take place after frame mean subtraction, and removes only spatial artifacts, not intensity artifacts that affect the entire frame. Therefore, when blood flow artifacts are so intense that they increase the overall frame intensity significantly, additional filtration is required. These mean intensity artifacts, when left unfiltered, can misrepresent the underlying signal.

This artifact could be due to tissue specific changes in optical paths during periods of high blood flow–either by proteins obscuring or scattering optical fluorescence from deeper GFP, or due to endogenous fluorescence of other proteins that are known to bleach quickly [76]. Visually, these artifacts are characterized by a sudden, intense increase in frame mean, followed by a slow decay in intensity.

To determine the proper frequency cutoff for filtration, we compared our GCaMP mean time series frequency distributions to a control transgenic-GFP fluorescence that was not linked to neural activity (fig. 5.2). The significant frequencies shared between the control data were between 0-0.5Hz, and we used this as our cutoff for mean filtration. After removal of these frequencies, frames that had previously exhibited these mean intensity artifacts were significantly improved (fig. 5.2, bottom right). This results in a more accurate representation of just the neural signal without the optical artifacts from the change in tissue.

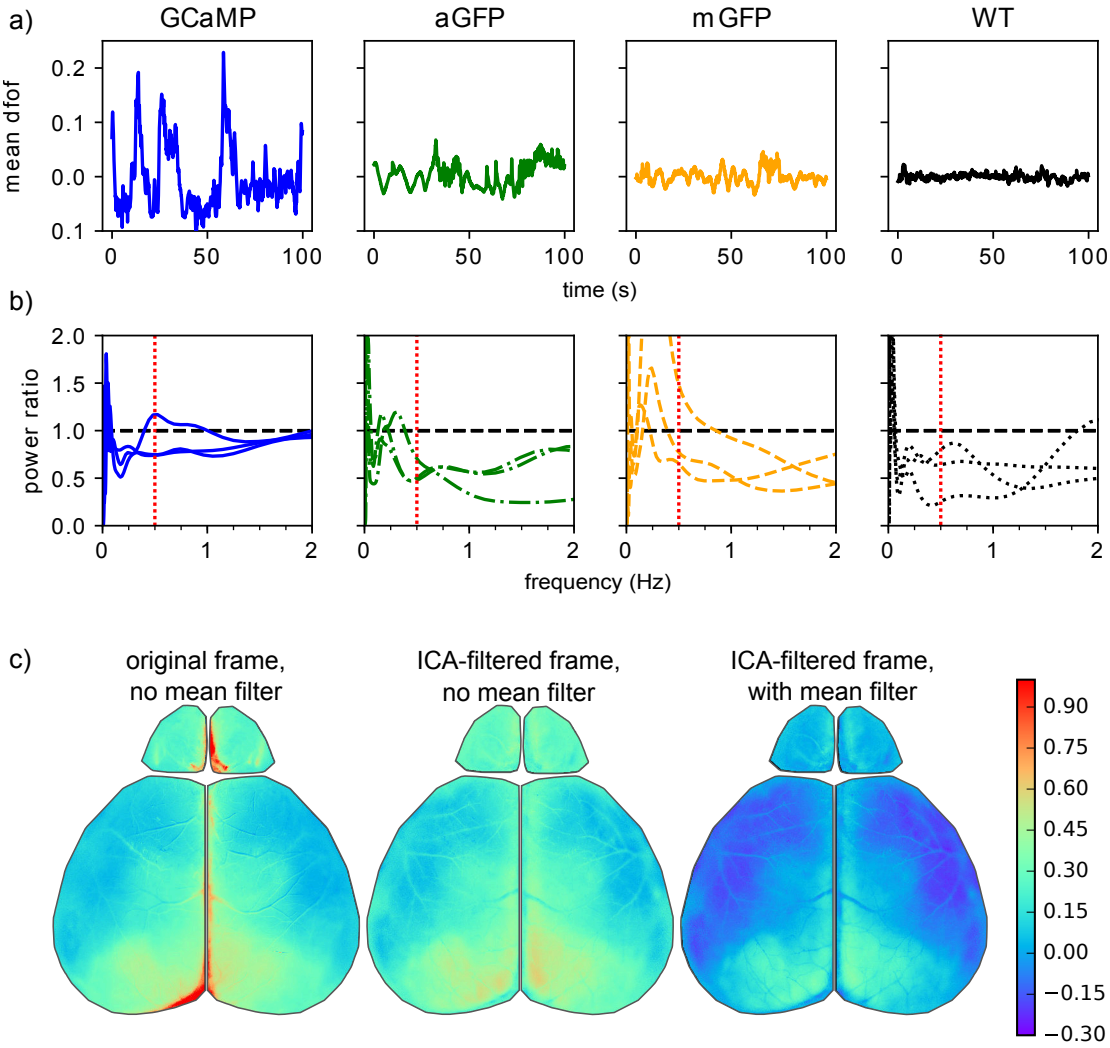

Figure 5.2: The mean time series must be extracted in order to perform ICA. To reduce large intensity artifacts, the mean time series must also be filtered. a) 100s examples of mean time series signals from P21 animals with various genotypes: our experimental mice, mice from a GFP-astrocyte line, mice from a GFP-microglia line, and wild type mice lacking any GFP on a Bl/6 background. b) The wavelet power ratio for time series from each genotype (n=3 per category). Since the endogenous fluorescent GFP controls (2 and 3) had significant power spectra up until 0.5Hz that the GCaMP data lacked, a 0-0.5Hz high pass wavelet filter was applied to the data. The 0.5Hz cutoff line is marked as a dotted red line. c) Example of the effect of mean filtering. Left panel is the original frame of the movie with an unfiltered mean. The widespread activation across the cortical surface is likely due to blood autofluorescence. In the center panel, the mean is re-added, but intensity across brain regions is slightly elevated relative to the original frame because of the removal of the intense blood vessel contraction artifact while retaining the same mean. After filtering, the signal intensity across brain regions is more localized to individual domains. This likely represents the amount of underlying signal without additional blood autofluorescence.

5.3 Filtering and Compression Residuals

The full ICA decomposition has as many components as the number of frames in the original video (12000 for a 20-minute 10Hz video). However, many of these components are noise-associated and represent an insignificant amount of information in the original dataset.

To be conservative, the signal and artifact components are all saved, as well as an additional portion of noise components, such that 25% of the components saved are noise associated. This results in a limited loss of information, mainly in non-cortical tissue surrounding the neocortex where there was little signal originally. The spatial residuals of lost signal is stored by default with every ICA decomposition, so that signal quality can be monitored (fig. 5.3).

Saving this reduced dataset results in an output file size $\sim$ 10x smaller than the original video, with very little loss of information (spatial signal lost is $3.4 - 05$ % of the total signal). An additional benefit of saving in this file format, is its portability and usability–Rather than locally storing a 64GB data file, a 6.4GB file can be opened and stored locally on most personal computers. While the full 64GB of RAM is still required for opening the full video, localized time series or video segments can be rebuilt with minimal RAM requirements.

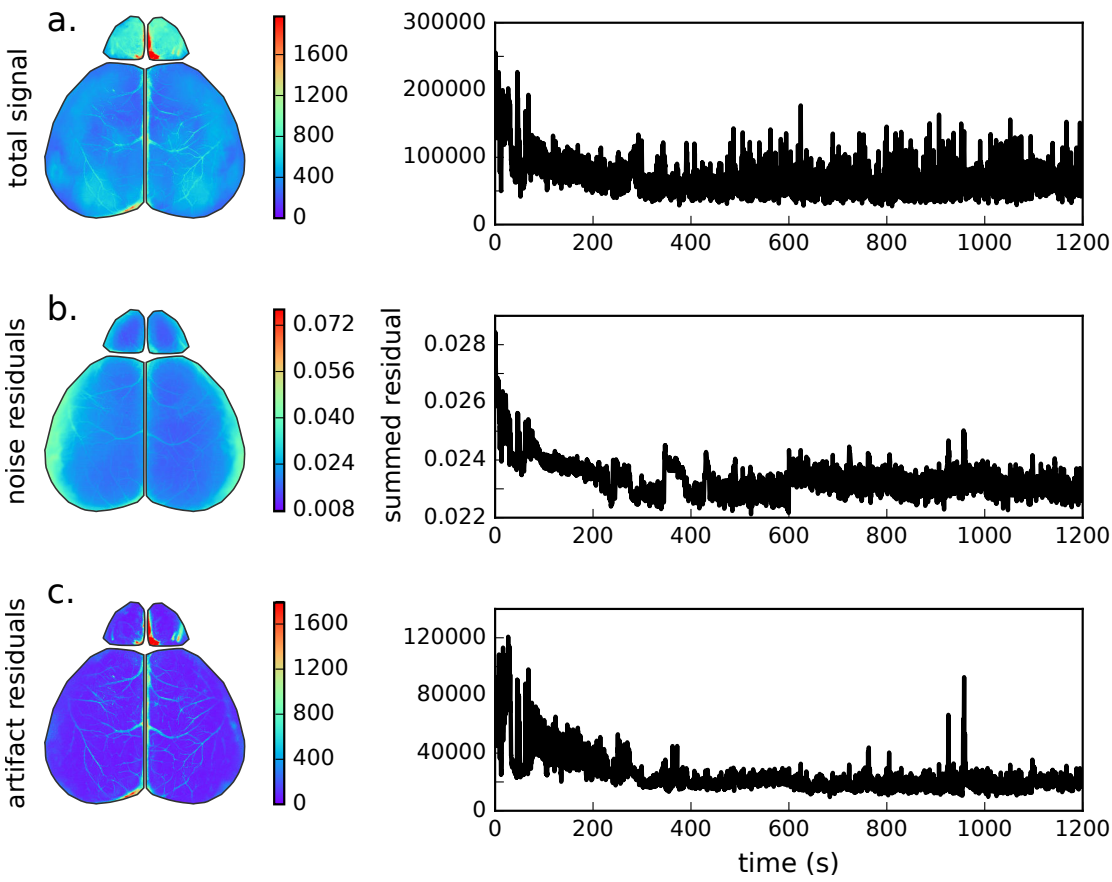

Figure 5.3: Comparison of spatial and temporal information content through compression and filtering. a) The original spatial information captured as quantified by a mean subtracted absolute value projected spatially (left) or temporally (right). b) The difference in information between the original input data and the rebuilt ICA projection after only saving a cropped portion of the noise components. The difference movie is projected spatially or temporally to visualize where information was lost in compression. c) Information removed by artifact filter. The artifact movie is rebuilt and projected spatially or temporally to visualize where information was modified by the ICA-based artifact filter.

5.4 Component Sorting

Signal and artifact components can be easily manually differentiated by their spatiotemporal properties. We originally did this by manual examination and classification in a custom built GUI (sec. 5.1). To reduce the amount of human interaction necessary for interacting with the data, we trained a machine learning classifier to sort these components instead.

We first extracted a series of spatial and temporal metrics from each component (fig. 5.4). Neural calcium signal components have distinct localized spatial domains, and often have smoother time courses, dominated by frequencies in the 2-8Hz range. Artifact components can a variety of spatial structures, including dispersed branching structures, indicative of a blood vessel artifact, or arcs on the border of the cortex, representing surgery artifacts. Artifact components can have varying temporal properties, such as movement associated spikes, or sudden jumps in level, where reflectance on the surgical surface was altered. These additional surgical artifacts can also be removed with the blood artifacts, cleaning up previously unusable data. We performed spatial binarizations based on a single tailed gaussian threshold cutoff, and applied morphometric analyses to extract spatial components. For temporal components, a wavelet decomposition was applied and various frequency metrics were extracted. No single metric was enough to differentiate signal and artifact components, since distributions are largely overlapping. When visualized as the first two components of a SVD decomposition (fig. 5.5, bottom), it is visually apparent that the neural components cluster more tightly together, while the artifact components occupy a wider spread of metric space. For this reason, we applied a machine learning classifier for component sorting.

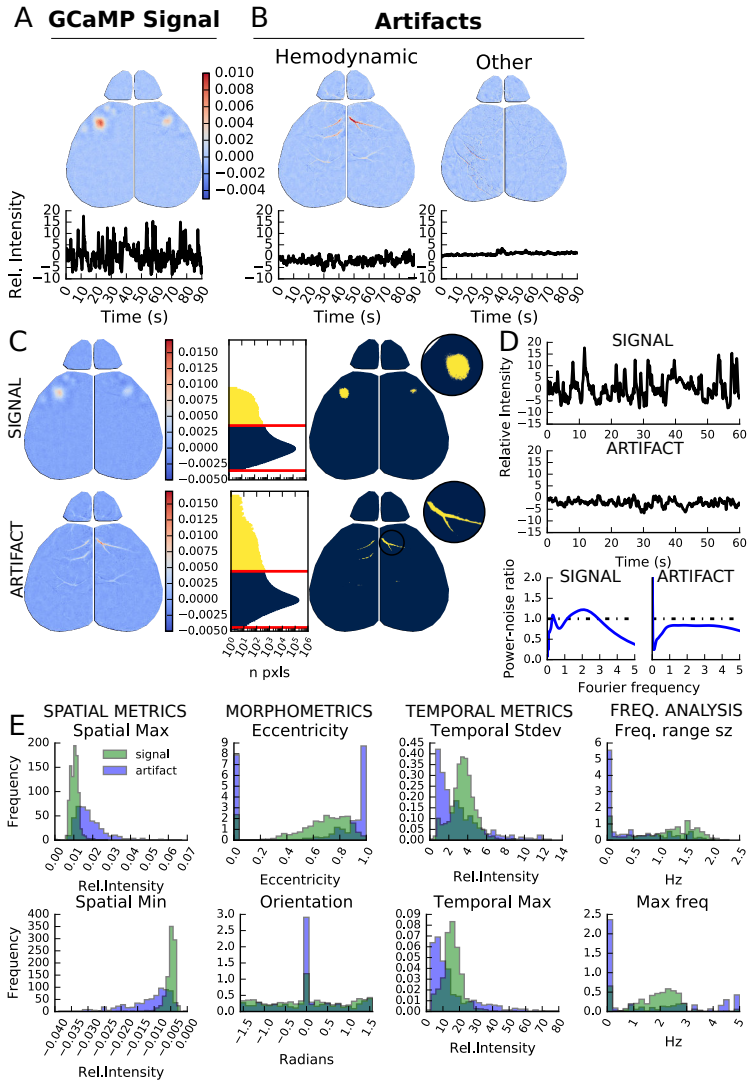

Figure 5.4: a) An example of a spatial map of a signal component and its associated time series. b) Extracted artifacts can be either hemodynamic or noise associated. Images displayed represent an example hemodynamic and movement artifact. These components have timecourses that are different from each other, and from the signal component. c) Example signal and artifact components are binarized by thresholding the pixel intensity histogram. The histogram is displayed vertically, highlighting the single tail detected in yellow. The spatial result of this threshold binarization is displayed on the right, with circular panels highlighting the difference in detected binarized structure detected between signal and artifact components. d) Example time series from signal and artifact components are displayed in terms of relative intensity. The power-noise ration is displayed below; the signal component had significant power between 0-3Hz, while the artifact component only had significant power below 0.25Hz. e) Comparison of distributions from example extracted spatial, spatial morphometric, temporal, and frequency metrics.

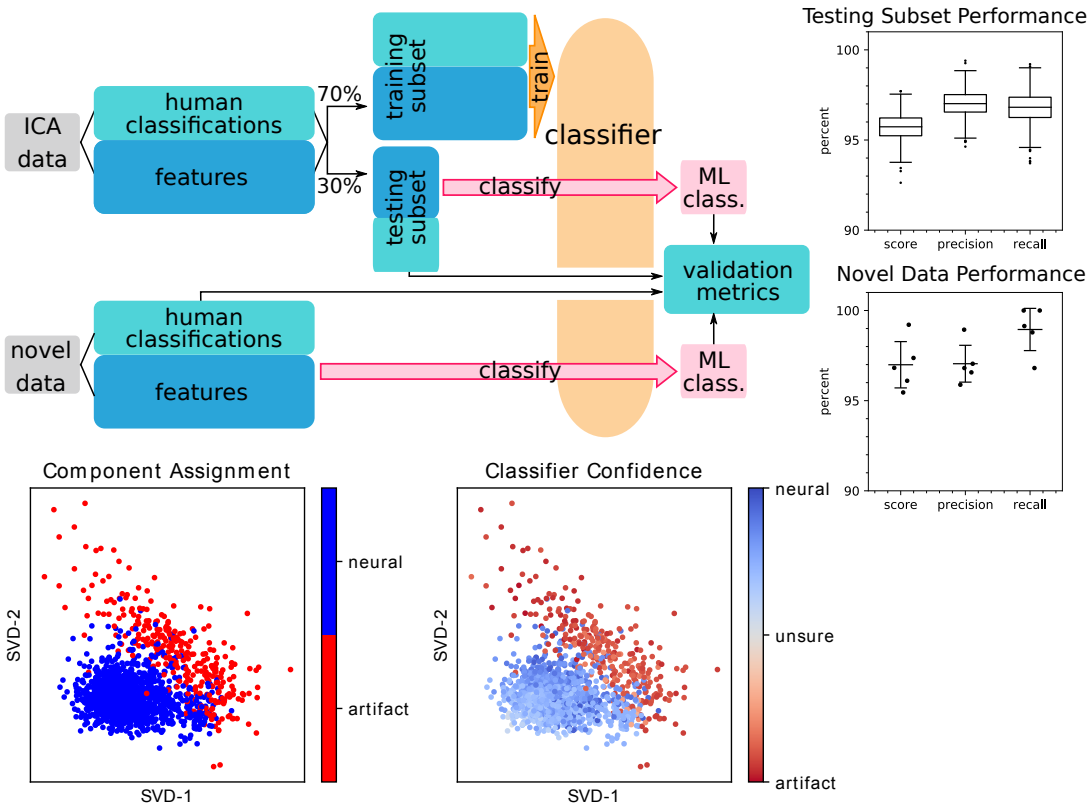

Figure 5.5: Schematic of machine learning approach for ICA component sorting. The P21 ICA dataset was split into testing and training groups. The training group was used to train the classifier, and the testing subset with manual component assignment was used for validation of the classifier's performance. The performance of novel data was then compared against a lab member's independent assessment of the component assignment to generate validation metrics. On the top right panel, score, precision, and recall metrics are calculated for the testing subset of the classifier trained on 1000 different random divisions of data. In the bottom right panel, score, precision, and recall are displayed for a set of 5 novel datasets. The bottom two panels represent the first two dimensions of a SVD decomposition of our metric data. In the left panel, the distribution is colored by manual human selection. In the right panel, the distribution is colored by the classifier's confidence in each classification.

We then trained a machine learning voting classifier to distinguish signal and artifact components based on their spatial and temporal extracted metrics (fig. 5.5). Our ICA dataset of manually assigned components and their associated metrics was split into a training and testing subset. The classifier was trained on the training subset, then tested on the testing subset. This resulted in a classifier with a score, precision, and recall of $>90\%$. When tested against 5 novel datasets it had not been previously exposed to, the classifier did even better, with score, precision, and recall all $>95\%$. This confidence was roughly on the order of human precision for component sorting.

After implementing this classifier, data could fully processed on a computing cluster without any need for a graphical login. Thus, we could process, sort, extract time series, and generate metrics without the need to process or store data on a local computer.

5.5 Wavelet Analysis

The wavelet decomposition [68] was used for frequency power analysis, mean signal filtration, and will be applied in sec. 3.1 as a measure of signal correlation. Wavelet decomposition takes a signal (fig. 5.6a) and transforms it into a continuous time-frequency domain, similar to a spectrogram (fig. 5.6c). In addition to this time-frequency domain, phase information is stored in an additional matrix (not shown). A wavelet distribution must be chosen to represent the shape of the 'wave' that we are convoluting with our signal. A sine-like wave would provide extremely accurate information about the frequency of the signal at any point at time, but sacrifice precision on the temporal accuracy of the signal. A gaussian distribution, or derivative-of-gaussian distribution would provide high accuracy on the location of the signal, but less information on the frequency. We used a $\omega = 4$ Morlet wavelet wave (fig. 5.6b) to slightly bias toward temporal accuracy.

The wavelet power spectra (fig. 5.6c, top) represents the power of the wavelet convolved with the source signal at any given time point, for any possible frequency. When this signal is integrated over time, a power spectrum (fig. 5.6c, top right) is produced.
We used a red noise model for significance for our data, which is fit based on the autocorrelation of each input signal. The noise floor is shown on the power spectrum as a dotted line.

The wavelet power spectra can also be normalized to this noise floor (fig. 5.6c, bottom). This view highlights where significant signal was present in our dataset, and is more useful for visual interpretation. The power spectra for the regular and normalized power spectra are displayed again in fig. 5.6d, with a more regular frequency axis.

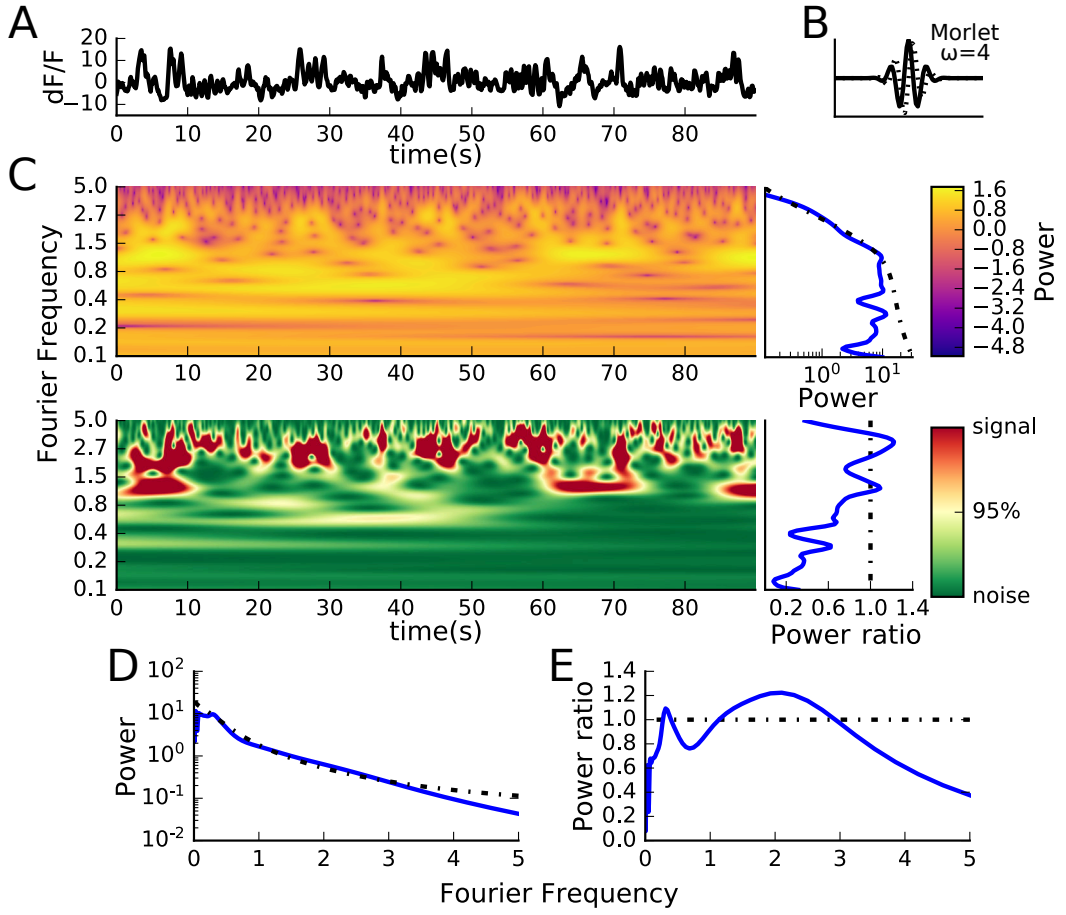

Figure 5.6: A visual summary of wavelet power methods. a) An example time dF/F series extracted from a component time series. b) The $\omega = 4$ Morlet wave used for wavelet decomposition. c) The power of the wavelet decomposition at each time point across the possible fourier-converted frequency range (top) along with it's integrated power spectrum, compared to the noise distribution.

In the bottom panel, the power ratio is displayed with an evenly spaced frequency axis. The power of the wavelet decomposition at each time point across the possible fourier-converted frequency range (top) along with it's integrated power spectrum, compared to the noise distribution.

Similar to eigendecompositions, the original data can be reconstructed from the full power and phase matrices with little to no loss of information. We can use this property to selectively filter ranges of frequencies, which we apply for mean signal filtering in sec. 5.2.

We use the normalized power spectra to generate temporal properties of mean signals in sec. 5.2, for sorting independent components based on temporal properties in component sorting sec. 5.4, and to generate time series properties in sec. 3.1.

Two signals can be compared in wavelet-space by convolving their power and phase matrices to calculate coherence. This value varies between 0 and 1, and can be used to represent the similarity of two signals. We do this to calculate network interactions in sec. 3.1.

5.6 Region Assignment

Regions were assigned to domains by a combination of data-driven and manual selection methods. First, a hierarchical clustering was computed using time series generated from each domain, giving us a starting map (fig. 5.7a). We also looked at images generated by visualizing time course metrics plotted over each cortical domain (fig. 5.7b). We included metrics such as the percent of the cortical area covered by each domain, the spatial extent of the domain, the standard deviation of the domain time course, and the significant frequency range of the domain time series, as quantified through its wavelet power spetrum. These metrics gave us some metrics for dividing higher order from primary sensory regions. We also took into account the expected functional regions as delineated by the Allen Brain Institute's functional map [56] (fig. 5.7c). Together, we used all this information to manually assign region flags to each domain. fig. 5.7d shows this final assignment. A machine-learning classifier could likely be trained for automated domain sorting.

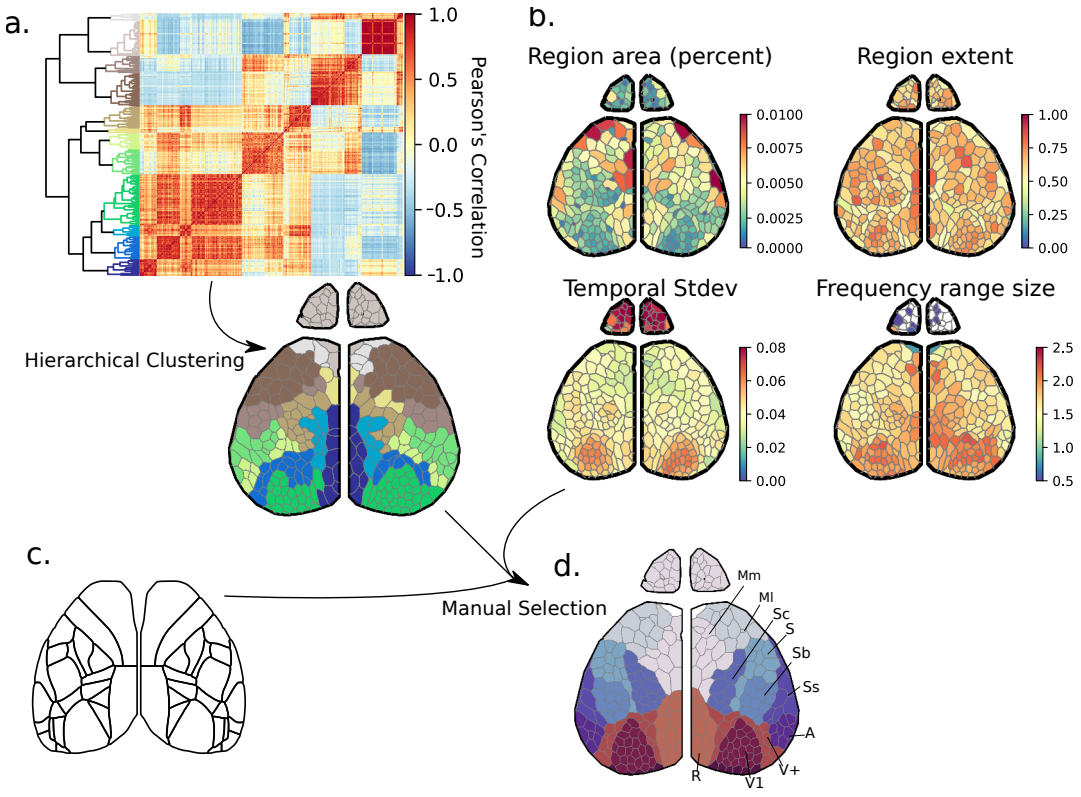

Figure 5.7: Using data-guided methods to assign domains to cortical regions. a) Hierarchical clustering based on Pearson's correlation produces a set of $\sim$ 13 regions across the cortical surface. b) Domains colored by various calculated spatial and temporal metrics to aid region assignment. Region area is calculated as a percent of the total cortical surface. Region extend ranges from 0 to 1 and calculates the relative area of a domain to its bounding box. Temporal standard deviation is calculated from the extracted time series, and frequency range size is calculated from wavelet significance. c) The Allen Brain atlas map [56] is additionally used for reference. d) The final manually assigned region, with associated labels.

Chapter 6

Abbreviations and Terms Defined

P21 - postnatal day 21

ΔF/F (dFoF) - change in fluorescence over mean fluorescence

ICA - Independent Component Analysis

PCA - Principal Component Analysis

BTBR - Black and Tan Brachyury (Mouse Line)

Domain Map - maximum projection map of ICA components

Domain - A single contiguous unit from a domain map, represents an ICA component's maximal region of influence

Mosiac Movie - a video representation of the time series extracted under each domain in the domain map

Region Abbreviations (Used in Region Maps in fig. 2.3 and fig. 3.3.

Key	Region
R	Retrosplenial
V+	Higher Order Visual
A	Auditory
Ss	Somatosensory Secondary
Sc	Somatosensory Core
Sb	Somatosensory Barrel
S	Somatosensory (Other)
Mm	Motor Medial
Ml	Motor Lateral
O	Olfactory

Chapter 7

Methods

Mice

All animal studies were conducted in accordance with the UCSC Office of Animal Research Oversight andInstitutional Animal Care and Use Committee protocols. Snap25 GCaMP6s transgenic mice (JAX: 025111), microglial GFP (JAX: 005582), and astrocyte GFP (MGI: 3843271) were maintained in UCSCs mouse facilities. Transgene presence was determined through genotyping before any procedure was performed. The expression of this transgene resulted in pan-neuronal expression of GCaMP6s throughout the nervous system.

Surgical procedure

All mice were anesthetized with isoflurane (2.5% in pure oxygen) for the procedure. General anesthesia was first induced in an isoflurane chamber then maintained with a nose cone. Body temperature was maintained at 37C for the duration of the surgery and recovery using a

feedback-regulated heading pad. Lidocaine (1%) was applied subcutaneous on the scalp, followed by careful removal of skin above the skull. If the eyes were open, opthalmic ointment was used protect the eyes. The head was glued using cyanoacrylate to two head bars, one across the back of the skull and the other on the lateral parietal bone. For young animals (P2 - P9), they were left on a heated platform for the duration of the recording. For animals that could crawl/walk (P10 - P21), the mice were transfer to a rotating disk for the duration of the recording. At the end of the recording session, the animal was either euthanized or perfused and the brain dissected.

Recording calcium dynamics

In-vivo wide-field fluorescence recordings were collected in a minimally invasive manner. Imaging through the skull by single-photon excitation light from two blue LED light produces a green fluorescent signal that is collected through coupled 50mm Nikon lenses (f5.6 / f1.2, optical magnification 1x) into a scientific cMOS camera (PCO Edge 5.5MP, 6.5μm pixel resolution). Excitation light was filtered with a 480/30 nm bandpass and the emission signal was filtered with 520/36 nm bandpass. Data collection was performed in a dark, quiet room with minimal changes in ambient light or sound. Each video segment consisted of a set of continuously collected images at 10 frames per second for 10 minutes. Raw data was written directly as a set of 16 bit multi-image TIFF files. Generally the total amount of recorded data for each animal was 40 min and the amount of time in between video segments was less than 1 minute. When video length over 20 minutes was used, multiple videos were concatenated together.

Recordings used

For the quantifications in ICA-based Filtering and Segmentation (sec. 2), the following animals were used. All of the animals were of the genotype Snap25-GCaMP6s, and between 21-22 days of age. 40 minutes of data were recorded from each animal. Animals recorded on the same date were littermates.

Date	Animal Number	Condition
190506	1	Control
	2	Control
190508	3	Control
	4	Control
190408	5	PPE
	6	PPE
190423	7	PPE
	8	PPE

Example components, videos, and video stills for filtering and segmentation were all generated from data recorded from animal 4.

Spatial resolution analyses were performed on a single 20 minute segment from a different animal (191019). Temporal resolution analyses were performed on a 10 minute recording at 20Hz from this same animal. Recording was performed as specified in “Recording calcium dynamics”, apart from the temporal resolution.

4 P21 control animals were used to determine the frequency ranges of hemodynamic artifacts for mean filtering were recorded from the following mouse lines: microglia GFP, astrocyte GFP, and wildtype mice on the same Bl/6 that lacked the Snap25 GCaMP6s transgene.

Mean filtration

Mean filtration was performed with a wavelet filter. First, time series signals were fitted with a $\omega = 4$ morlet wave for a wavelet decomposition using code adapted from C. Torrence and G. Compo [68], available at URL: http://paos.colorado.edu/research/wavelets/ Significance was determined using the red-noise model fit to the time series autocorrelation.

Frequency distributions are all displayed as the ratio of the wavelet power integrated over the time series, relative to the noise cutoff. For wavelet filtering, the original signal was rebuilt excluding all frequency signals in a certain range. In our case, we omitted frequencies between 0-0.5Hz.

ICA decomposition and saving

ICA was performed using FastICA [32], implemented through python's sklearn decomposition [1]. The ICA decomposition was applied to the flattened 2-D representation of the video data under the cortical ROI mask. The mean time series is pre-subtracted from the array before SVD decomposition or ICA decomposition, since ICA cannot separate sources with a mean signal effect. The filtered, unfiltered mean, and associated metadata are all saved. Data is stored and saved in this flattened format for storage optimization. Components are locally spatially

reconstructed for visualization in the GUI.

Requesting the full number of components resulted in extremely lengthy processing times. To reduce the processing time, the data was preprocessed through Singular Value Decomposition (SVD) whitening, and noise components were cropped. To ensure that no signal was lost, and there were ample dimensions left for ICA separation, the cutoff between SVD signal and noise was identified, and components were cropped to 5 times the SVD signal to noise cutoff value. This cutoff can be adjusted through the `svd_multiplier` option while ICA projecting.

After calculating and sorting the ICA results, excessive noise components are removed from the dataset for compression. The cutoff was determined by identifying the inflection point in the lag-1 autocorrelation distribution with a two-peaked KDE fit, as detailed in sec. 2.2. Components were saved such that 75% of the components saved were signal and artifact, and 25% of the components saved were noise associated. If not enough noise components were returned by the ICA decomposition, there is a risk that signals were not sufficiently isolation, so the ICA decomposition was repeated with a higher SVD cutoff until enough additional noise components were included.

ICA returns unsorted components that are unsorted and often flipped. Components were first sorted by their time series variance. Signal components were characterized by a single tailed gaussian distribution (detailed in methods: binarization). If the tail was negative, the component was flipped spatially and temporally. In this way, components were all identified as positive affectors for visualization, and movie rebuilding was not affected.

All code used in this paper will be available on github under www.github.com/ackmanlab/pyseas

or as a package–pySEAS in the python package index (`pip install seas`).

Data processing

ICA decompositions were run on single node on a computing cluster. 10 minute videos could be processed on a node with 512 GB of RAM, while 20 minute videos often required larger nodes, depending on the number of components requested. All data could be processed on a node with 1024 GB of RAM. After ICA processing, most map creation and time series processing could be computed on local computers with 16-32GB of RAM.

Compression and filtering residuals

Compression residuals are calculated while saving the ICA decomposition results. The original movie is rebuilt from the reduced ICA results, and residuals are calculated by taking the absolute value of the difference between the two videos. The spatial and temporal projection of this absolute difference movie is saved as the spatial and temporal residuals of the decomposition, and is stored as metadata with each ICA decomposition.

Statistical significance

Statistical significance was calculated using a paired t-tests from `statsmodel.formula.api` with Holm-Sidak multiple testing correction. When comparing two small, related samples, `ttest_rel` from `scipy.stats` was used instead.

The following symbols are used to indicate statistical significance in figures:

Symbol	Meaning
ns	$p > 0.05$
	$p \leq 0.05$
	$p \leq 0.01$
*	$p \leq 0.001$

Component classification

We used an ensemble random forest classifier from the scikit-learn [1] package to train and classify between signal and artifact activity. We first manually scored our components as signal or artifact associated. These were later used for training as well as validation. The data was randomly split into two subsets: the training set contained 70% of the dataset, and the testing set other contained 30%. We trained an ensemble of decision tree classifiers using the training set, including extracted features of the cortical dynamics and human classification scores. After training, the machine classifier predicted classification scores on the testing subset. We then validated the machine classification with a set of human classifications to generate accuracy (percent correct classifications), precision (correctly classified artifacts / all selected artifacts), and recall (correctly classified artifacts / missed artifacts) scores for various randomly generated data subsets.

Component binarization

Spatial metrics for component classification were generated from the binarized representation of each component. The spatial histogram of pixel intensity across each component can be visualized as a single tailed gaussian distribution centered around 0, where the tail represents the spatial domain of each component.

The two edges of the distribution are first identified. The boundary closer to 0 is taken as the edge of the central noise distribution, and the pixels between the noise distribution and the wider tail are included in the binarized version of the component.

Spatial metric generation

The following spatial metrics were generated for each independent component for component sorting.

The spatial minima and maxima were identified from the smallest and largest pixel value present in the components.

The center of mass is also detected for each component. If a binarized domain was detected, the thesholded area is additionally calculated.

Wavelet temporal metric generation

The following temporal metrics were generated for each independent component for component sorting. The time series maxima, minima, lag-1 autocorrelation, and standard deviation were calculated for each component.

Using wavelet analysis, the following frequency metrics were additionally generated. The highest and lowest values detected for significant frequencies were identified, as well as the size of the significant span of frequencies. The most significant dominant frequency frequency was identified, and the signal to noise ratio at this value was also stored.

Domain map creation

Domain maps were created by separating the cortex into regions represented by different ICA components. Each component was blurred by a 51-pixel kernel, then the maximum projection was taken through the component layers.

The resulting data is a cortical map that denotes the component with maximum influence over any given pixel.

This map was then further processed to get rid of domains smaller than 1/10th the mean domain. Any domain smaller than this size is checked to see if the second most significant component would produce a larger continuous structure. If after a few loops of this, pixels cannot be assigned into a larger structure, the points are excluded from the final map. Indices are then adjusted such that any non-continuous regions represented by the same domain are assigned to different units.

Olfactory bulbs were included in map generation, but domains were highly variable, and were excluded from map quantifications.

Jaccard map comparisons

For every domain or region in the original map, the nearest neighbor was identified in the comparison map with a KNN tree. To quantify the spatial similarity of each identified domain or region, the Jaccard index (spatial overlap / union) was then calculated. For each comparison, n Jaccard indices were calculated, where n is the number of identified units in the original map.

When comparing maps generated from different animals, the optimal alignment was calculated by shifting the second map up to 100 pixels in any direction. The optimal direction was determined by maximizing the Jaccard overlap. Each generated map was compared to one map from the same animal, one littermate, and two non-littermates, as well as one randomly generated voronoi map.

Voronoi map creation

Voronoi maps were created to match the same number of domains or regions as the original map, n, and shares the same cortical mask as the original map.

To create this map, n points were distributed randomly across the cortical mask. To turn these points into regions, the voronoi diagram was created using the scipy spatial package [62] and was applied as a voronoi map.

Grid map creation

Grid maps were created to match or exceed the same number of units as the original map, and share the same cortical mask as the original map. A uniformly spaced 2D grid was placed

over the original map, and resulting units were counted. If the number of resulting spatial units exceeded that of the original map by < 15, the map was accepted as a valid comparison. Otherwise, the map was rejected and a new grid map was calculated.

Region assignment

[[assigning regions]]

Key	Region
R	Retrosplenial
V+	Higher Order Visual
A	Auditory
Ss	Somatosensory Secondary
Sc	Somatosensory Core
Sb	Somatosensory Barrel
S	Somatosensory (Other)
Mm	Motor Medial
Ml	Motor Lateral
O	Olfactory

Residuals and total signal analyses

To quantify the amount of signal present in the original movie that was not included in the domain map, residuals were calculated by subtracting the 'mosaic movie', representing time series from each spatial domain from the original movie. The absolute value was then applied so that all numbers represented a positive difference, and residuals were summed to create a single value. The time series was not re-added to either the original movie or mosaic movie, since this can be easily summarized as a different temporal metric. To represent the amount of relative variation to the original dataset, this number was compared to the summed absolute value of the mean-subtracted original movie.

Variation analysis

To quantify the variation between time series, the variance at each time point was calculated by calculating the variance between that time point at each extracted time series, then summing across all time points.

Isoflurane recordings

Isoflurane recordings were performed on 3 different Snap25-GCaMP6s mice:

Date	Animal Number
200108	1
200109	2

Date	Animal Number
200110	3

20 minutes were recorded from each animal with the isoflurane nose cone in place, with no oxygen or isoflurane applied. Next, 20 minutes were recorded from each animal with oxygen flow but no isoflurane. This is the 0% isoflurane condition. Maps generated from animal 3 were used for example figures.

For the 0.25%, 0.5%, and 1% isoflurane recordings, the isoflurane dose was increased, the animal was allowed to adjust to the higher isoflurane level for 20 minutes, then an additional 20 minutes of data were recorded for each condition.

Anesthesia network analysis

For network analyses, time series and region classifications were extracted from the 0% domain map generated from each animal. Time series were grouped by the 10 regions as defined in methods: region assignment. Within each group, time series were compared to all other extracted time series through wavelet coherence. A percentage metric was generated to determine what percent of time series were highly correlated (coherence > 0.4) to time series extracted in the compared region. Coherence was calculated by applying the algorithms described in [22] to the wavelet package by Torrence and Compo [68].

This percentage of highly associated time series was then combined across all animals under the given condition, then subtracted to identify group differences.

Bibliography

[1] Alexandre Abraham, Fabian Pedregosa, Michael Eickenberg, Philippe Gervais, Andreas Mueller, Jean Kossaifi, Alexandre Gramfort, Bertrand Thirion, and Ga?l Varoquaux. Machine learning for neuroimaging with scikit-learn. *Frontiers in Neuroinformatics*, 8, 2014.

[2] James B. Ackman, Timothy J. Burbridge, and Michael C. Crair. Retinal waves coordinate patterned activity throughout the developing visual system. *Nature*, 490(7419):219–225, October 2012.

[3] James B Ackman and Michael C Crair. Role of emergent neural activity in visual map development. *Current Opinion in Neurobiology*, 24:166–175, February 2014.

[4] James B. Ackman, Hongkui Zeng, and Michael C. Crair. Structured dynamics of neural activity across developing neocortex. *bioRxiv*, 2014.

[5] William E. Allen, Isaac V. Kauvar, Michael Z. Chen, Ethan B. Richman, Samuel J. Yang, Ken Chan, Viviana Gradinaru, Benjamin E. Deverman, Liqun Luo, and Karl Deisseroth. Global Representations of Goal-Directed Behavior in Distinct Cell Types of Mouse Neocortex. *Neuron*, 94(4):891–907.e6, May 2017.

[6] Penn Anna A., Riquelme Patricio A., Feller Marla B., and Shatz Carla J. Competition in Retinogeniculate Patterning Driven by Spontaneous Activity. *Science*, 279, March 1998.

[7] D. S. Bassett, E. Bullmore, B. A. Verchinski, V. S. Mattay, D. R. Weinberger, and A. Meyer-Lindenberg. Hierarchical Organization of Human Cortical Networks in Health and Schizophrenia. *Journal of Neuroscience*, 28(37):9239–9248, September 2008.

[8] Vera Brust, Philipp M. Schindler, and Lars Lewejohann. Lifetime development of behavioural phenotype in the house mouse (Mus musculus). *Frontiers in zoology*, 12(1):S17, 2015.

[9] G. Y. Cederquist, E. Azim, S. J. Shnider, H. Padmanabhan, and J. D. Macklis. Lmo4 Establishes Rostral Motor Cortex Projection Neuron Subtype Diversity. *Journal of Neuroscience*, 33(15):6321–6332, April 2013.

[10] Heng Chen, Lucina Q. Uddin, Xujun Duan, Junjie Zheng, Zhiliang Long, Youxue Zhang, Xiaonan Guo, Yan Zhang, Jingping Zhao, and Huafu Chen. Shared atypical default mode and salience network functional connectivity between autism and schizophrenia. *Autism Research*, 10(11):1776–1786, November 2017.

[11] Tsai-Wen Chen, Trevor J. Wardill, Yi Sun, Stefan R. Pulver, Sabine L. Renninger, Amy Baohan, Eric R. Schreiter, Rex A. Kerr, Michael B. Orger, Vivek Jayaraman, Loren L. Looger, Karel Svoboda, and Douglas S. Kim. Ultrasensitive fluorescent proteins for imaging neuronal activity. *Nature*, 499(7458):295–300, July 2013.

[12] Jakob Christensen and Therese Koops Gr. Prenatal Valproate Exposure and Risk of Autism Spectrum Disorders and Childhood Autism. *JAMA*, page 8, 2013.

[13] Matthew T. Colonnese, Anna Kaminska, Marat Minlebaev, Mathieu Milh, Bernard Bloem, Sandra Lescure, Guy Moriette, Catherine Chiron, Yehezkel Ben-Ari, and Rustem Khazipov. A Conserved Switch in Sensory Processing Prepares Developing Neocortex for Vision. *Neuron*, 67(3):480–498, August 2010.

[14] Julia V. Cramer, Benno Gesierich, Stefan Roth, Martin Dichgans, Marco Dring, and Arthur Liesz. In vivo widefield calcium imaging of the mouse cortex for analysis of network connectivity in health and brain disease. *NeuroImage*, 199:570–584, October 2019.

[15] Luca Dodero, Mario Damiano, Alberto Galbusera, Angelo Bifone, Sotirios A. Tsaftsaris, Maria Luisa Scattoni, and Alessandro Gozzi. Neuroimaging Evidence of Major Morpho-Anatomical and Functional Abnormalities in the BTBR T+TF/J Mouse Model of Autism. *PLoS ONE*, 8(10):e76655, October 2013.

[16] Gaelle E. Doucet, Maxwell J. Luber, Priti Balchandani, Iris E. Sommer, and Sophia Frangou. Abnormal auditory tonotopy in patients with schizophrenia. *npj Schizophrenia*, 5(1):16, December 2019.

[17] Fabrizio Esposito, Tommaso Scarabino, Aapo Hyvarinen, Johan Himberg, Elia Formisano, Silvia Comani, Gioacchino Tedeschi, Rainer Goebel, Erich Seifritz, and Francesco Di Salle. Independent component analysis of fMRI group studies by self-organizing clustering. *NeuroImage*, 25(1):193–205, March 2005.

[18] David A. Feldheim, Young-In Kim, Andrew D. Bergemann, Jonas Frisn, Mariano Barbacid, and John G. Flanagan. Genetic analysis of ephrin-A2 and ephrin-A5 shows their requirement in multiple aspects of retinocollicular mapping. *Neuron*, 25(3):563–574, 2000.

[19] Matthew F. Glasser, Timothy S. Coalson, Emma C. Robinson, Carl D. Hacker, John Harwell, Essa Yacoub, Kamil Ugurbil, Jesper Andersson, Christian F. Beckmann, Mark Jenkinson, Stephen M. Smith, and David C. Van Essen. A multi-modal parcellation of human cerebral cortex. *Nature*, 536(7615):171–178, July 2016.

[20] Luciano Custo Greig, Mollie B. Woodworth, Maria J. Galazo, Hari Padmanabhan, and Jeffrey D. Macklis. Molecular logic of neocortical projection neuron specification, development and diversity. *Nature Reviews Neuroscience*, 14(11):755–769, October 2013.

[21] Ludovica Griffanti, Gholamreza Salimi-Khorshidi, Christian F. Beckmann, Edward J. Auerbach, Gwenalle Douaud, Claire E. Sexton, Enik Zsoldos, Klaus P. Ebmeier, Nicola Filippini, Clare E. Mackay, Steen Moeller, Junqian Xu, Essa Yacoub, Giuseppe Baselli,

Kamil Ugurbil, Karla L. Miller, and Stephen M. Smith. ICA-based artefact removal and accelerated fMRI acquisition for improved resting state network imaging. *NeuroImage*, 95:232–247, July 2014.

[22] A. Grinsted, J. C. Moore, and S. Jevrejeva. Application of the cross wavelet transform and wavelet coherence to geophysical time series. *Nonlinear Processes in Geophysics*, 11(5/6):561–566, November 2004.

[23] I. L. Hanganu, Y. Ben-Ari, and R. Khazipov. Retinal Waves Trigger Spindle Bursts in the Neonatal Rat Visual Cortex. *Journal of Neuroscience*, 26(25):6728–6736, June 2006.

[24] I. L. Hanganu, J. F. Staiger, Y. Ben-Ari, and R. Khazipov. Cholinergic Modulation of Spindle Bursts in the Neonatal Rat Visual Cortex In Vivo. *Journal of Neuroscience*, 27(21):5694–5705, May 2007.

[25] Donald Hebb. *The Organization of Behavior*. Wiley & Sons, New York, 1949.

[26] ClausC. Hilgetag, Gully A. P. C. Burns, Marc A. O'Neill, Jack W. Scannell, and Malcolm P. Young. Anatomical connectivity defines the organization of clusters of cortical areas in the macaque and the cat. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 355(1393):91–110, January 2000.

[27] Elizabeth M.C. Hillman. Coupling Mechanism and Significance of the BOLD Signal: A Status Report. *Annual Review of Neuroscience*, 37(1):161–181, July 2014.

[28] Elizabeth M.C. Hillman, Anna Devor, Matthew B. Bouchard, Andrew K. Dunn, G.W. Krauss, Jesse Skoch, Brian J. Bacskai, Anders M. Dale, and David A. Boas. Depth-resolved optical imaging and microscopy of vascular compartment dynamics during somatosensory stimulation. *NeuroImage*, 35(1):89–104, March 2007.

[29] Lewis B. Holmes, Elizabeth A. Harvey, Brent A. Coull, Kelly B. Huntington, Shahram Khoshbin, Ailish M. Hayes, and Louise M. Ryan. The teratogenicity of anticonvulsant drugs. *New England Journal of Medicine*, 344(15):1132–1138, 2001.

[30] Bryan M. Hooks and Chinfei Chen. Distinct Roles for Spontaneous and Visual Activity in Remodeling of the Retinogeniculate Synapse. *Neuron*, 52(2):281–291, October 2006.

[31] Andrew D. Huberman, Marla B. Feller, and Barbara Chapman. Mechanisms Underlying Development of Visual Maps and Receptive Fields. *Annual Review of Neuroscience*, 31(1):479–509, July 2008.

[32] A. Hyvarinen and E. Oja. Independent component analysis: algorithms and applications. *Neural Networks*, 13:411–430, March 2000.

[33] M V Jones, P A Brooks, and N L Harrison. Enhancement of gamma-aminobutyric acid-activated Cl- currents in cultured rat hippocampal neurones by three volatile anaesthetics. *The Journal of Physiology*, 449(1):279–293, April 1992.

[34] Mariel G. Kozberg, Ying Ma, Mohammed A. Shaik, Sharon H. Kim, and Elizabeth M. C. Hillman. Rapid Postnatal Expansion of Neural Networks Occurs in an Environment of Altered Neurovascular and Neurometabolic Coupling. *The Journal of Neuroscience*, 36(25):6704–6717, June 2016.

[35] Matthew D Krasowski and Neil L Harrison. The actions of ether, alcohol and alkane general anaesthetics on GABA $_{a}$ and glycine receptors and the effects of TM2 and TM3 mutations. *British Journal of Pharmacology*, 129(4):731–743, February 2000.

[36] Elena Kutsarova, Martin Munz, and Edward S. Ruthazer. Rules for Shaping Neural Connections in the Developing Brain. *Frontiers in Neural Circuits*, 10, January 2017.

[37] Alexandra H. Leighton and Christian Lohmann. The Wiring of Developing Sensory CircuitsFrom Patterned Spontaneous Activity to Synaptic Plasticity Mechanisms. *Frontiers in Neural Circuits*, 10, September 2016.

[38] Heiko Luhmann. Subplate cells: amplifiers of neuronal activity in the developing cerebral cortex. *Frontiers in Neuroanatomy*, 3, 2009.

[39] Heiko J. Luhmann, Anne Sinning, Jenq-Wei Yang, Vicente Reyes-Puerta, Maik C. Stttgen, Sergei Kirischuk, and Werner Kilb. Spontaneous Neuronal Activity in Developing Neocortical Networks: From Single Cells to Large-Scale Interactions. *Frontiers in Neural Circuits*, 10, May 2016.

[40] Li Ma, Wentai Liu, and Andrew Hudson. Propofol Anesthesia Increases Long- range Frontoparietal Corticocortical Interaction in the Oculomotor Circuit in Macaque Monkeys. *Anesthesiology*, 130(4):560–571, 2019.

[41] Ying Ma, Mohammed A. Shaik, Mariel G. Kozberg, Sharon H. Kim, Jacob P. Portes, Dmitriy Timerman, and Elizabeth M. C. Hillman. Resting-state hemodynamics are spatiotemporally coupled to synchronized and symmetric neural activity in excitatory neurons. *Proceedings of the National Academy of Sciences*, 113(52):E8463–E8471, December 2016.

[42] Ying Ma, Mohammed A. Shaik, Mariel G. Kozberg, Sharon H. Kim, Jacob P. Portes, Dmitriy Timerman, and Elizabeth M. C. Hillman. Resting-state hemodynamics are spatiotemporally coupled to synchronized and symmetric neural activity in excitatory neurons. *Proceedings of the National Academy of Sciences*, 113(52):E8463–E8471, December 2016.

[43] Eleonora Maggioni, Jorge Arrubla, Tracy Warbrick, Jrgen Dammers, Anna M. Bianchi, Gianluigi Reni, Michela Tosetti, Irene Neuner, and N. Jon Shah. Removal of Pulse Artefact from EEG Data Recorded in MR Environment at 3T. Setting of ICA Parameters for Marking Artefactual Components: Application to Resting-State Data. *PLoS ONE*, 9(11):e112147, November 2014.

[44] Hiroshi Makino, Chi Ren, Haixin Liu, An Na Kim, Neehar Kondapaneni, Xin Liu, Duygu Kuzum, and Takaki Komiyama. Transformation of Cortex-wide Emergent Properties during Motor Learning. *Neuron*, 94(4):880–890.e8, May 2017.

[45] M J McKeown, S Makeig, G G Brown, T-P Jung, S S Kindermann, A J Bell, and T J Sejnowski. Analysis of fMRI Data by Blind Separation into Independent Spatial Components. *Human Brain Mapping*, page 31, 1998.

[46] David A. McVea, Timothy H. Murphy, and Majid H. Mohajerani. Large Scale Cortical Functional Networks Associated with Slow-Wave and Spindle-Burst-Related Spontaneous Activity. *Frontiers in Neural Circuits*, 10, December 2016.

[47] Kimford J Meador, Gus A Baker, Nancy Browning, Morris J Cohen, Rebecca L Bromley, Jill Clayton-Smith, Laura A Kalayjian, Andres Kanner, Joyce D Liporace, Page B Pennell, Michael Privitera, and David W Loring. Fetal antiepileptic drug exposure and cognitive outcomes at age 6 years (NEAD study): a prospective observational study. *The Lancet Neurology*, 12(3):244–252, March 2013.

[48] Vinod Menon. Large-scale brain networks and psychopathology: a unifying triple network model. *Trends in Cognitive Sciences*, 15(10):483–506, October 2011.

[49] Nicholas J. Michelson and Takashi D. Y. Kozai. Isoflurane and ketamine differentially influence spontaneous and evoked laminar electrophysiology in mouse V1. *Journal of Neurophysiology*, 120(5):2232–2245, November 2018.

[50] M. Minlebaev, Y. Ben-Ari, and R. Khazipov. Network Mechanisms of Spindle-Burst Oscillations in the Neonatal Rat Barrel Cortex In Vivo. *Journal of Neurophysiology*, 97(1):692–700, January 2007.

[51] V. Mountcastle. The columnar organization of the neocortex. *Brain*, 120(4):701–722, April 1997.

[52] Vernon B. Mountcastle. Modality and topographic properties of single neurons of cat's somatic sensory cortex. *Journal of Neurophysiology*, 20(4):408–434, July 1957.

[53] Vernon B. Mountcastle, Philip W. Davies, and Alvin L. Berman. Response properties of neurons of cat's somatic sensory cortex to peripheral stimuli. *Journal of Neurophysiology*, 20(4):374–407, July 1957.

[54] Timothy H. Murphy, Jamie D. Boyd, Federico Bolaos, Matthieu P. Vanni, Gergely Silasi, Dirk Haupt, and Jeff M. LeDue. High-throughput automated home-cage mesoscopic functional imaging of mouse cortex. *Nature Communications*, 7(1):11611, September 2016.

[55] Jan Mlter, Lilach Avitan, and Geoffrey J. Goodhill. Detecting neural assemblies in calcium imaging data. *BMC Biology*, 16(1):143, December 2018.

[56] Seung Wook Oh, Julie A. Harris, Lydia Ng, Brent Winslow, Nicholas Cain, Stefan Mihalas, Quanxin Wang, Chris Lau, Leonard Kuan, Alex M. Henry, Marty T. Mortrud, Benjamin Ouellette, Thuc Nghi Nguyen, Staci A. Sorensen, Clifford R. Slaughterbeck, Wayne Wakeman, Yang Li, David Feng, Anh Ho, Eric Nicholas, Karla E. Hirokawa, Phillip Bohn, Kevin M. Joines, Hanchuan Peng, Michael J. Hawrylycz, John W. Phillips, John G. Hohmann, Paul Wohnoutka, Charles R. Gerfen, Christof Koch, Amy Bernard, Chinh Dang, Allan R. Jones, and Hongkui Zeng. A mesoscale connectome of the mouse brain. *Nature*, 508(7495):207–214, April 2014.

[57] Jaime F. Olavarria and Ryoko Hiroi. Retinal influences specify cortico-cortical maps by postnatal day six in rats and mice. *The Journal of Comparative Neurology*, 459(2):156–172, April 2003.

[58] O Paulsen. Natural patterns of activity and long-term synaptic plasticity. *Current Opinion in Neurobiology*, 10(2):172–180, April 2000.

[59] Ann Petro, Hannah G. Sexton, Caroline Miranda, Anit Rastogi, Jonathan H. Freedman, and Edward D. Levin. Persisting neurobehavioral effects of developmental copper exposure in wildtype and metallothionein 1 and 2 knockout mice. *BMC Pharmacology and Toxicology*, 17(1), December 2016.

[60] Gerald R Prescott and Luke H Chamberlain. Regional and developmental brain expression patterns of SNAP25 splice variants. *BMC Neuroscience*, 12(1):35, December 2011.

[61] Mikail Rubinov and Ed Bullmore. Fledgling pathoconnectomics of psychiatric disorders. *Trends in Cognitive Sciences*, 17(12):641–647, December 2013.

[62] SciPy 1.0 Contributors, Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, lhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antnio H. Ribeiro, Fabian Pedregosa, and Paul van Mulbregt. SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nature Methods*, 17(3):261–272, March 2020.

[63] Kristin K. Sellers, Davis V. Bennett, Axel Hutt, and Flavio Frhlich. Anesthesia differentially modulates spontaneous network dynamics by cortical area and layer. *Journal of Neurophysiology*, 110(12):2739–2751, December 2013.

[64] Bridgette D. Semple, Klas Blomgren, Kayleen Gimlin, Donna M. Ferriero, and Linda J. Noble-Haeusslein. Brain development in rodents and humans: Identifying benchmarks of maturation and vulnerability to injury across species. *Progress in Neurobiology*, 106-107:1–16, July 2013.

[65] Francesco Sforazzini, Alice Bertero, Luca Dodero, Gergely David, Alberto Galbusera, Maria Luisa Scattoni, Massimo Pasqualetti, and Alessandro Gozzi. Altered functional connectivity networks in acallosal and socially impaired BTBR mice. *Brain Structure and Function*, 221(2):941–954, March 2016.

[66] Nicholas C. Spitzer. Electrical activity in early neuronal development. *Nature*, 444(7120):707–712, December 2006.

[67] Andrew Y. Y. Tan. Spatial diversity of spontaneous activity in the cortex. *Frontiers in Neural Circuits*, 9, September 2015.

[68] Christopher Torrence and Gilbert P. Compo. A practical guide to wavelet analysis. *Bulletin of the American Meteorological society*, 79(1):61–78, 1998.

[69] M. T. Valley, M. G. Moore, J. Zhuang, N. Mesa, D. Castelli, D. Sullivan, M. Reimers, and J. Waters. Separation of hemodynamic signals from GCaMP fluorescence measured with wide-field imaging. *Journal of Neurophysiology*, 123(1):356–366, January 2020.

[70] Matthieu P. Vanni, Allen W. Chan, Matilde Balbi, Gergely Silasi, and Timothy H. Murphy. Mesoscale Mapping of Mouse Cortex Reveals Frequency-Dependent Cycling between Distinct Macroscale Functional Modules. *The Journal of Neuroscience*, 37(31):7513–7533, August 2017.

[71] Patricia Wollstadt, Kristin K. Sellers, Axel Hutt, Flavio Frohlich, and Michael Wibral. Anesthesia-related changes in information transfer may be caused by reduction in local information generation. In *2015 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 4045–4048, Milan, August 2015. IEEE.

[72] Tung-Lin Wu, Arabinda Mishra, Feng Wang, Pai-Feng Yang, John C. Gore, and Li Min Chen. Effects of isoflurane anesthesia on resting-state fMRI signals and functional connectivity within primary somatosensory cortex of monkeys. *Brain and Behavior*, 6(12):e00591, December 2016.

[73] Benyi Xiong, Anan Li, Yang Lou, Shangbin Chen, Ben Long, Jie Peng, Zhongqin Yang, Tonghui Xu, Xiaoquan Yang, Xiangning Li, Tao Jiang, Qingming Luo, and Hui Gong. Precise Cerebral Vascular Atlas in Stereotaxic Coordinates of Whole Mouse Brain. *Frontiers in Neuroanatomy*, 11, December 2017.

[74] Zhijun Yao, Bin Hu, Yuanwei Xie, Fang Zheng, Guangyao Liu, Xuejiao Chen, and Weihao Zheng. Resting-State Time-Varying Analysis Reveals Aberrant Variations of Functional Connectivity in Autism. *Frontiers in Human Neuroscience*, 10, September 2016.

[75] Jun Zhuang, Williams Derric, Valley Matthew, Li Yang, Marina Garrett, and Jack Waters. An extended retinotopic map of mouse cortex. *eLife*, January 2017. PMC5218535.

[76] W. R. Zipfel, R. M. Williams, R. Christie, A. Y. Nikitin, B. T. Hyman, and W. W. Webb. Live tissue intrinsic emission microscopy using multiphoton-excited native fluorescence and second harmonic generation. *Proceedings of the National Academy of Sciences*, 100(12):7075–7080, June 2003.