

Research



Cite this article: Linden NJ, Tabuena DR, Steinmetz NA, Moody WJ, Brunton SL, Brunton BW. 2021 Go with the FLOW: visualizing spatiotemporal dynamics in optical widefield calcium imaging. *J. R. Soc. Interface* **18**: 20210523.
<https://doi.org/10.1098/rsif.2021.0523>

Received: 23 June 2021
 Accepted: 2 August 2021

Subject Category:
 Life Sciences—Engineering interface

Subject Areas:
 computational biology

Keywords:
 widefield calcium imaging, computational neuroscience, dynamical systems, coherent structures, finite-time Lyapunov exponents

Author for correspondence:
 Bingni W. Brunton
 e-mail: bbrunton@uw.edu

Electronic supplementary material is available online at <https://doi.org/10.6084/m9.figshare.c.5557409>.

Go with the FLOW: visualizing spatiotemporal dynamics in optical widefield calcium imaging

Nathaniel J. Linden^{1,2}, Dennis R. Tabuena^{2,3}, Nicholas A. Steinmetz⁴, William J. Moody², Steven L. Brunton⁵ and Bingni W. Brunton²

¹Department of Bioengineering, ²Department of Biology, ³Graduate Program in Neuroscience, ⁴Department of Biological Structure, and ⁵Department of Mechanical Engineering, University of Washington, Seattle, WA 98195-0005, USA

SLB, 0000-0002-6565-5118; BWB, 0000-0002-4831-3466

Widefield calcium imaging has recently emerged as a powerful experimental technique to record coordinated large-scale brain activity. These measurements present a unique opportunity to characterize spatiotemporally coherent structures that underlie neural activity across many regions of the brain. In this work, we leverage analytic techniques from fluid dynamics to develop a visualization framework that highlights features of flow across the cortex, mapping wavefronts that may be correlated with behavioural events. First, we transform the time series of widefield calcium images into time-varying vector fields using optic flow. Next, we extract concise diagrams summarizing the dynamics, which we refer to as *FLOW (flow lines in optical widefield imaging) portraits*. These FLOW portraits provide an intuitive map of dynamic calcium activity, including regions of initiation and termination, as well as the direction and extent of activity spread. To extract these structures, we use the finite-time Lyapunov exponent technique developed to analyse time-varying manifolds in unsteady fluids. Importantly, our approach captures coherent structures that are poorly represented by traditional modal decomposition techniques. We demonstrate the application of FLOW portraits on three simple synthetic datasets and two widefield calcium imaging datasets, including cortical waves in the developing mouse and spontaneous cortical activity in an adult mouse.

1. Introduction

Coordinated organization of neural activity among brain regions is believed to serve many crucial roles, including performing specific computations in the cortex [1–3] and supporting brain development [4–6]; further, its disruption may lead to neurological disease [7–9]. One prominent characteristic of neural activity at the scale of brain regions is the rapid and coherent propagation of activity across the cortex, which has been widely observed in a variety of contexts, including spontaneous activity, task engagement, sleep and development [10–12]. Qualitatively similar patterns of neural activity propagation have also been observed in the retina, often referred to as retinal waves, during development [13–16]. Although such spatiotemporal dynamic features are often visually salient, it remains challenging to quantify and succinctly summarize their behaviour directly from neural recordings.

Widefield optical imaging of calcium activity provides a unique opportunity to study coordinated spatiotemporal neural activity among brain areas, because this experimental approach achieves large fields of view with high temporal and spatial resolution [17,18]. In general, widefield imaging experiments involve fluorescence imaging of the entire brain surface of animals that express optical indicator proteins in known populations of neurons [19–24]. Many

experiments choose to use genetically encoded calcium indicators from the GCaMP family to image neural calcium dynamics, which is a proxy for electrical neuronal activity [25–28]; more generally, the visualization methods we discuss here can be applied to any widefield optical imaging experiment, such as imaging with voltage-sensitive dyes [29,30]. Cortical activity has been measured using widefield calcium imaging in a variety of experiments, notably to study perceptual decision making [1,3,31–35], to extract cortical functional connectivity [8,36,37], to characterize cortical activity that organizes brain development [38] and to study the effects of disease in the cortex [7–9]. In all of these data, it is typical to observe multiple regions activating transiently or in regular succession, with distinct initiation sites and wave-like flows across the fields of view. These features can often be described as flow of activity with coherent travelling fronts; interestingly, all of these patterns are well studied as nonlinear features of spatiotemporal dynamical systems [39].

The most widely applied approaches to analyse time-varying recordings of high-dimensional neural activity are dimensionality reduction techniques, which extract *modes* that correspond to dominant, low-dimensional features of the high-dimensional data [40–42]. These low-dimensional features are useful as representations of the neural activity that facilitate further analysis and modelling. A dynamical model is needed when the analytic goals include prediction in time or real-time control. Furthermore, the observation that the dynamics of neuronal populations can be reduced to a relatively small number of features may be a clue about the mechanisms that underlie coordinated neural activity [43–46]. Common modal decomposition algorithms used in neuroscience [40,47] include singular value decomposition (SVD), which is closely related to principal component analysis (PCA), independent component analysis (ICA) and non-negative matrix factorization (NNMF). These techniques all solve for combinations of relatively few modes in space and time that reconstruct an estimate of the original high-dimensional data; their solutions differ by making different assumptions about the statistical structure of the modes.

There are many exciting recent innovations in modal decomposition for analysing large-scale neural data, some of which are extensions and derivatives of SVD, ICA and NNMF. Interestingly, while some of these methods have explicit representations of the temporal dynamics (for instance, jPCA [44], dynamic mode decomposition (DMD) [48–50] and NNMF with temporal constraints [51–54]), they largely set out to achieve space–time separation. Applying PCA and NNMF to segments of synthetic and experimental data (figure 1*a*) yields a set of spatial modes (figure 1*b*; temporal modes not shown) that provide a representation of the activity. However, these representations are static modes and may not adequately summarize spatiotemporal data. As an illustrative example, the synthetic data in figure 1 are a spatial Gaussian that grows, translates, then shrinks with time, and such spatiotemporal coherent features are poorly captured by PCA and NNMF decompositions of the data.

A complementary set of methods have been developed to describe spatiotemporal patterns in widefield neural activity by explicitly extracting propagating waves. Travelling waves are often characterized by their propagation speed and their direction (see [55,56] for examples), and these measures are then aggregated for all of the waves observed in a recording to quantify the trends in wave dynamics. While this

information has proven useful in studying the roles of waves, the approach is limited because waves need to be identified individually. Several related methods have used the computation of optical flow to convert widefield activity to time-varying vector fields [57,58]. This velocity field can then be analysed using tools from vector calculus to identify fixed points, including classifying each fixed point as a source or a sink of activity, and to quantify individual propagating waves. Nevertheless, these methods are constrained to wave-by-wave analyses and cannot summarize complex global activity.

Our visualization approach is inspired by the similarity of spatial flows observed in widefield optical imaging to flows of physical fluids. Humans have a deep intuition about fluid flows from our everyday experiences (e.g. the patterns of milk mixing in coffee, a river flowing). Representing brain data as a flow allows us to leverage this intuition and decades of methods from flow analysis and visualization. Propagation of neural activity has many commonalities with and differences from physical fluid flows. In both, there exist coherent structures whose boundaries may be invariant even as the activity changes with time. In fluid physics, these invariant manifolds are known as Lagrangian coherent structures (LCSs) [59–61], which act as transport barriers in the flow, either repelling or attracting material. LCSs are often visualized by computing ridges in the finite-time Lyapunov exponent (FTLE) field [62–65], although there are other computational approaches based on variational theory [66]. Some noteworthy biological applications include the use of LCSs to study the physics of jellyfish feeding [67] and understanding cardiovascular haemodynamics [68,69]. Unlike physical flows, neural activity is not governed by fundamental conservation laws; nevertheless, these dynamics are well described by time-varying vector fields [57,58,70,71].

In this work, we develop a visualization framework to capture the spatiotemporal dynamics of neural activity by extracting field lines in optical widefield imaging, which we call *FLOW* (*flow lines in optical widefield imaging*) portraits. FLOW portraits are generated by considering frame-by-frame dynamics as time-varying optical flow vector fields, from which we compute and integrate the ridges in its FTLE. To validate and develop intuition for our approach, we show that FLOW portraits give accurate and interpretable visual summaries of simple synthetic datasets. Next, we apply our methods to analyse bouts of activity from two widefield calcium imaging datasets in mice, both of which exhibit spontaneous, widespread activity across the cortex. The first data are recordings of spontaneous cortical activity of GCaMP6s-expressing mouse pups during their first 8 postnatal days [38]. The second example is a recording of spontaneous cortical activity in a GCaMP6s-expressing adult mouse [35]. In both examples, we demonstrate that FLOW portraits extract meaningful and interpretable outlines of the dominant patterns in the cortical activity that contribute to our understanding of the animals' developmental and behavioural states.

2. FLOW portraits

This work introduces FLOW portraits, which are visualizations that provide a concise and intuitive summary of the spatiotemporal dynamics, highlighting coherent structures in widefield

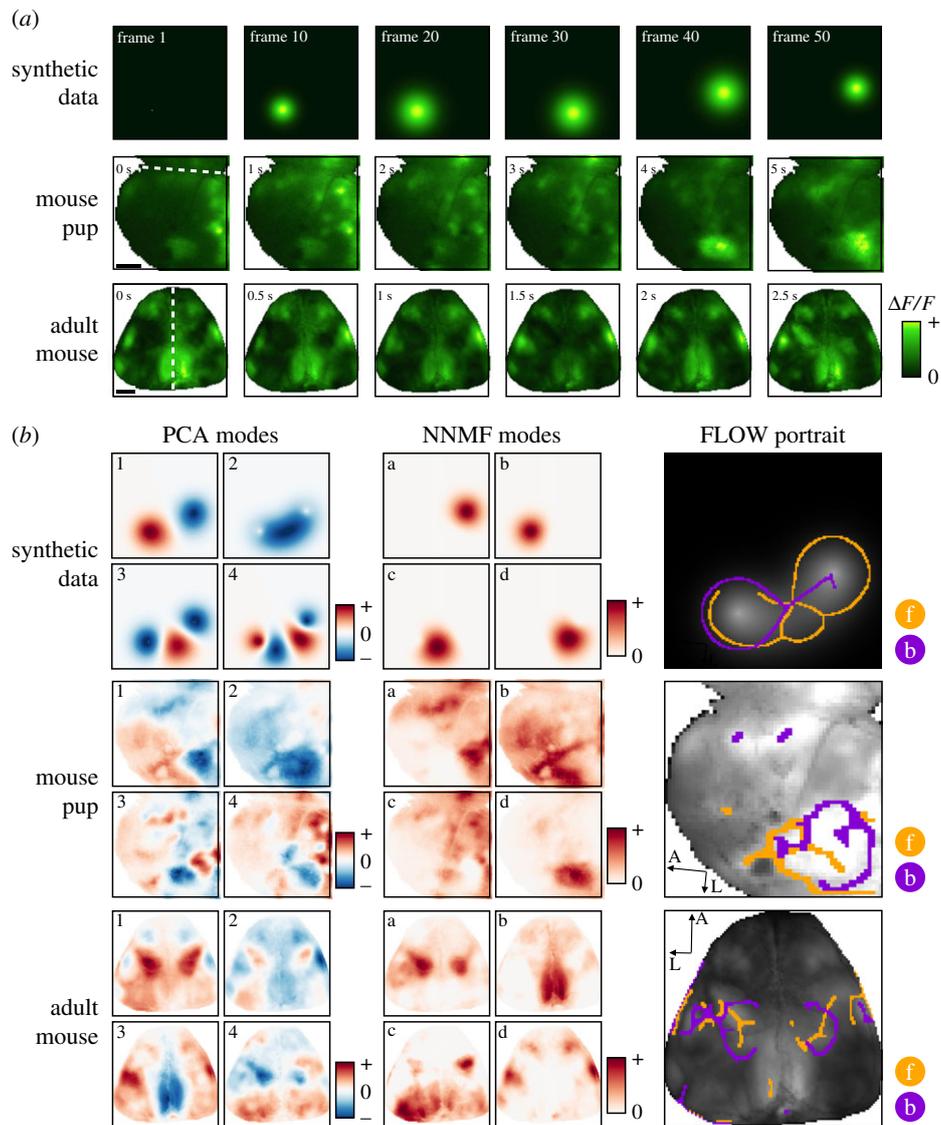


Figure 1. FLOW portraits capture coherent propagation of structures that are poorly represented by common modal decompositions that aim to achieve space–time factorization. (a) Three examples of spatiotemporal data for which we compare PCA, NNMF and our FLOW portraits. One synthetic example is a two-dimensional Gaussian that grows, translates to the right, then shrinks. Two further *in vivo* examples are widefield calcium imaging data from a developing pup and an adult mouse. The dashed white lines at 0 s indicate the midline of the brain. The mouse pup data include a pan-cortical wave from a postnatal day 7 (P7) animal; scale bar is 1 mm. The adult mouse data show spontaneous widefield calcium activity recorded in the dark; scale bar is 2 mm. (b) FLOW portraits show a succinct summary of the spatiotemporal flow in each example dataset, while spatial PCA and NNMF modes do not. The PCA modes are the first four spatial components; the NNMF modes are from a four-mode solution to the factorization and are not ordered. Both sets of modes decompose the growth and translation of activity into static spatial images, from which the flow of the activity cannot be easily appreciated. By contrast, our FLOW portraits highlight regions of activity initiation and termination, as well as the direction and extent of activity spread. Orange structures ('f'; forward-time FTLE) capture regions where activity propagates from, and purple structures ('b'; backward-time FTLE) capture regions where activity propagates towards. Videos illustrating all datasets are available as electronic supplementary material, videos 1–5. FLOW portraits were computed with integration lengths of 10 frames, 40 frames and 15 frames and the threshold percentile was set to the 85th, 93rd and 93rd percentiles for the synthetic, mouse pup and adult mouse datasets, respectively.

recordings. Importantly, FLOW portraits differ from modal decomposition techniques in that they do not provide a basis in which to approximate the data and cannot quantitatively explain variance in the recordings. Instead, FLOW portraits explicitly convert the image stack into time-varying vector fields to extract patterns of activity propagation in the data (figure 1). As our approach leverages and adapts analytic techniques from fluid dynamics [61] that are unfamiliar to most neuroscientists, this section describes how to compute the FTLE from time-varying vector fields. We also build intuition for how the ridges of the FTLE field can be interpreted in the context of widefield calcium imaging, using several simple synthetic examples.

The steps of our approach to compute FLOW portraits are illustrated in figures 2 and 3, and in electronic supplementary material, video S1. The input data are a video (i.e. image stack) of the relative change in fluorescence of the imaged optical protein indicator, $\Delta F/F$, as it changes in time over many frames. The raw fluorescence may drift over the course of an experiment, so $\Delta F/F$ is considered to be a robust proxy for the magnitude of neural activation, normalizing the change in fluorescence over a moving-window baseline [72]. FLOW portraits are well suited to summarizing data where optical activity is seen to diffuse or flow across the field of view, with varied patterns throughout the recording. To characterize the propagation of recorded neural activity

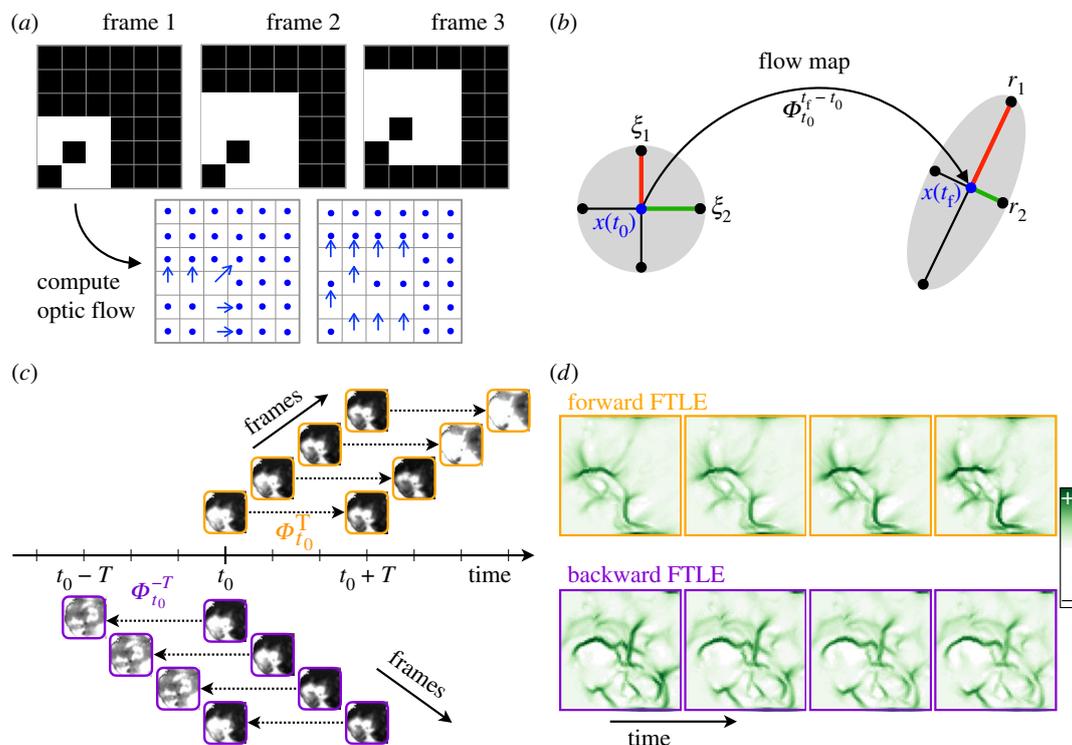


Figure 2. FTLE fields are computed from spatiotemporal data. (a) An illustration of how optic flow is computed from successive frames of images by correlating the relative movement of pixel intensities. This procedure converts widefield imaging data into a vector field of velocities. (b) The flow map at every pixel location is a virtual particle at x integrated through the vector field for a duration of T , from t_0 to $t_0 + T$; in reverse time, particles are integrated from t_0 to $t_0 - T$. This integration stretches neighbouring particles in some directions (r_1) and compresses them in others (r_2). (c) The flow map computation is repeated starting at each frame of the video, at base time $t_0 + k\Delta t$, where Δt is the separation between frames; forward maps are orange and backward maps are purple. (d) The FTLE fields are computed from the Jacobians of these flow maps; example FTLE fields are illustrated for successive frames of widefield imaging data.

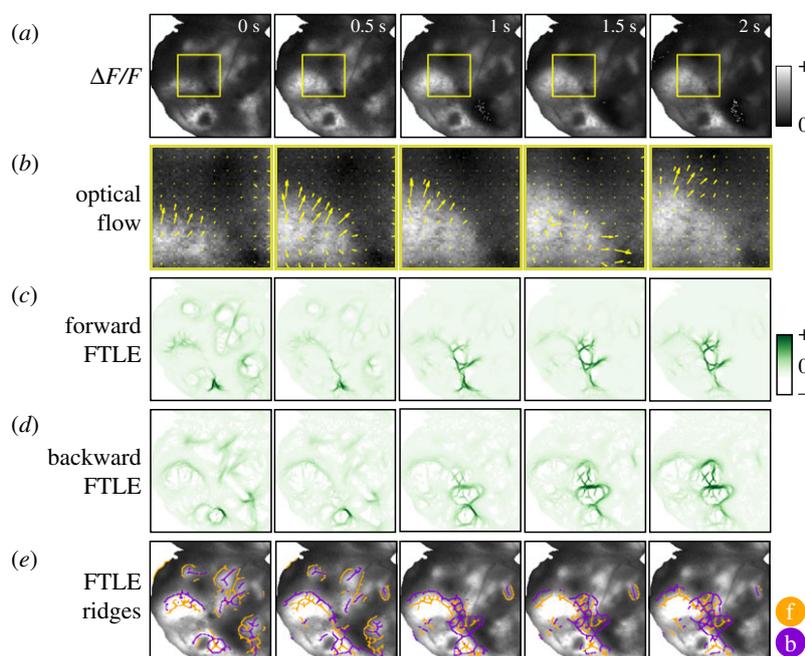


Figure 3. Steps to compute a FLOW portrait. Starting with widefield data preprocessed as $\Delta F/F$ (a), optical flow is used to convert the frame-by-frame changes in pixel intensity to a vector field, shown zoomed in for the smaller area outlined with the yellow box and at 1/6 spatial resolution for clarity (b). Next, the FTLE fields are computed in forwards (c) and backwards (d) time using an integration length of 2 s (40 frames); here we show only the non-negative Lyapunov exponents. Ridges of these fields highlight coherent structures of the flow (e), and these ridges are used to compute the final FLOW portrait (see figure 4). The threshold percentile was set to 93%. The forward-time FTLE ridges (orange) highlight regions that repel flow, while the backward-time ridges (purple) show regions that attract activity. Note that ridges in neighbouring frames are similar but do vary in time.

across brain areas through space, we first compute the flow vector field using optic flow. Next, the FTLE is computed from the time-varying vector field using the standard

integration method, as outlined by Onu *et al.* [73]. Last, the FTLE field is post-processed to visualize ridge-like features that highlight the coherent features of a spatial flow [61,62].

It is important to note that we refer to the processed FTLE ridges as *FLOW portraits* to avoid misinterpretation with traditional LCS analysis in fluid dynamics [61]. Details of data collection, preprocessing and computation are described in the Methods (§5).

2.1. Optical flow of widefield imaging data

We describe the frame-by-frame spread of neural activity as time-varying vector fields, computed by optical flow. Specifically, as regions of high pixel intensity in $\Delta F/F$ move and diffuse across the field of view, these coherent motions can be converted into a vector field of velocities, dx/dt and dy/dt , at every pixel in the recording (figure 2a). We denote this vector field as $\mathbf{v}(\mathbf{x}, t)$, defined at every point in space \mathbf{x} at time t . Motion velocities are commonly estimated from video data in computer vision using optical flow algorithms [74], and biological visual systems of vertebrates and invertebrates also perceive moving scenes with computations akin to optical flow [75,76]. In addition, some prior work has explored optical flow computations in widefield calcium imaging data [57,58]. Here, we use the Horn–Schunck (HS) [77] method because of its simplicity and its observed strong performance on our sample data.

Figure 3a,b shows an example of snapshots of $\Delta F/F$ data and the extracted optical flow vector fields. The magnitude and direction of the vector at each pixel is computed by solving for the optimal vector field that describes the change from each frame to the subsequent frame (see schematic in figure 2a). In order to minimize the effects of noise and numerical differentiation on the optical flow field, we apply temporal scaling and smoothing to the computed vector fields. Briefly, the magnitude of each optical flow vector is scaled proportionally to the relative change in the raw pixel intensity for the corresponding pixel over a prescribed time delay. This scaling attenuates the magnitudes of vectors that do not represent corresponding changes in the widefield imaging data. To mitigate the effects of pixel noise, we also apply temporal Gaussian smoothing to the scaled vector fields. The scaled and smoothed HS optical flow vector fields are used throughout the rest of the FLOW portrait algorithm where velocity data are required. This process of computing the optical flow vector field from widefield imaging data is analogous to the process of extracting the motion vector field from particle image velocimetry data [78,79] in experimental fluid dynamics. Both approaches approximate the velocity field from experimental data of *material* transported through the studied flow.

2.2. The finite-time Lyapunov exponent

Once a flow velocity field, $\mathbf{v}(\mathbf{x}, t)$, is computed, there are numerous computational approaches that can be performed to study and characterize the flow. These methods include instantaneous metrics from vector calculus, such as the divergence and the curl of the vector field; modal decomposition techniques [80,81], such as POD and DMD; and Lagrangian metrics, such as the FTLE [59,61,62]. Although instantaneous metrics have the potential to extract relevant features from widefield imaging optical flow fields (electronic supplementary material, figure S1; [57]), the unsteady nature of these data suggests that Lagrangian metrics may provide a more useful summary of the activity. Here, we compute

the FTLE fields [62] to extract time-invariant features of flow-like widefield activity.

The FTLE field is a scalar field $\sigma(\mathbf{x}, t_0, T)$ defined at every point in space \mathbf{x} and time t_0 , with respect to some relevant time scale of integration, T . The FTLE field is a measure of how much neighbouring initial conditions separate when integrated through the velocity field \mathbf{v} for a duration T . Thus, regions of high stretching for positive T (forward time) or negative T (backward time) provide time-varying analogues of stable and unstable manifolds, respectively [39,61,62]. The FTLE field is typically approximated numerically from flow-field snapshots at discrete instants in time [62,65]. First, the flow map $\Phi_{t_0}^T$ is approximated on a discretized set of spatial points, typically the same discretized domain where the velocity field is defined. The flow map $\Phi_{t_0}^T$ describes the position of an initial condition $\mathbf{x}(t_0)$ after it is integrated along the vector field \mathbf{v} for a duration T (figure 2b) and is defined as

$$\mathbf{x}(t_0 + T) = \Phi_{t_0}^T(\mathbf{x}(t_0)) = \mathbf{x}(t_0) + \int_{t_0}^{t_0+T} \mathbf{v}(\mathbf{x}(\tau), \tau) d\tau. \quad (2.1)$$

Next, the flow map Jacobian $\mathbf{D}\Phi_{t_0}^T$ is approximated via finite-difference derivatives with neighbouring points in the flow. In two dimensions, the flow map Jacobian at a point \mathbf{x} is

$$\mathbf{D}\Phi_{t_0}^T(\mathbf{x}) \approx \begin{bmatrix} \frac{\Phi_{x,t_0}^T(\mathbf{x}+\Delta x) - \Phi_{x,t_0}^T(\mathbf{x}-\Delta x)}{2\Delta x} & \frac{\Phi_{x,t_0}^T(\mathbf{x}+\Delta y) - \Phi_{x,t_0}^T(\mathbf{x}-\Delta y)}{2\Delta y} \\ \frac{\Phi_{y,t_0}^T(\mathbf{x}+\Delta x) - \Phi_{y,t_0}^T(\mathbf{x}-\Delta x)}{2\Delta x} & \frac{\Phi_{y,t_0}^T(\mathbf{x}+\Delta y) - \Phi_{y,t_0}^T(\mathbf{x}-\Delta y)}{2\Delta y} \end{bmatrix}, \quad (2.2)$$

where Φ_{x,t_0}^T denotes the x component of $\Phi_{t_0}^T$, and Φ_{y,t_0}^T denotes the y component. The finite-time Lyapunov exponent σ is finally computed from the largest eigenvalue λ_{\max} of the Cauchy–Green deformation tensor $\Delta = (\mathbf{D}\Phi_{t_0}^T)^T \Phi_{t_0}^T$, which is the maximum singular value of the flow map Jacobian

$$\sigma(\mathbf{x}_0, t_0, T) = \frac{1}{|T|} \ln\left(\sqrt{\lambda_{\max}[\Delta(\mathbf{x}_0, t_0, T)]}\right). \quad (2.3)$$

The FTLE value at a point \mathbf{x}_0 determines the maximum stretching that may occur between \mathbf{x}_0 and a perturbed location $\mathbf{x}_0 + \boldsymbol{\epsilon}$ after time T

$$\Phi_{t_0}^T(\mathbf{x}_0 + \boldsymbol{\epsilon}) \approx \Phi_{t_0}^T(\mathbf{x}_0) + \mathbf{D}\Phi_{t_0}^T(\mathbf{x}_0) \cdot \boldsymbol{\epsilon}, \quad (2.4)$$

where the amplification of the perturbation $\boldsymbol{\epsilon}$ is bounded by

$$\|\mathbf{D}\Phi_{t_0}^T(\mathbf{x}_0) \cdot \boldsymbol{\epsilon}\|_2 \leq \exp(\sigma|T|) \|\boldsymbol{\epsilon}\|_2. \quad (2.5)$$

The σ term is understood to depend on \mathbf{x}_0 , t_0 and T .

The FTLE field is quite robust to noisy measurements of the vector field $\mathbf{v}(\mathbf{x}, t)$ [59], since the computation involves integration in time, which tends to average out noise. This robustness was a major factor in its wide adoption in fluid mechanics, where experimentally acquired velocity fields often contain noise and outliers. The same robustness is appealing for optical widefield imaging.

Figures 2 and 3 illustrate the intuition behind this FTLE computation, and additional implementation details are provided in the Methods. The key insight in the FTLE computation is that virtual particles at every pixel location flow according to the vector field from t_0 to $t_0 + T$, and these integrated optical flow fields form a *flow map* $\Phi_{t_0}^T$ (figure 2b). This flow stretches neighbouring virtual particles, so that equidistant particles have stretched in some directions and

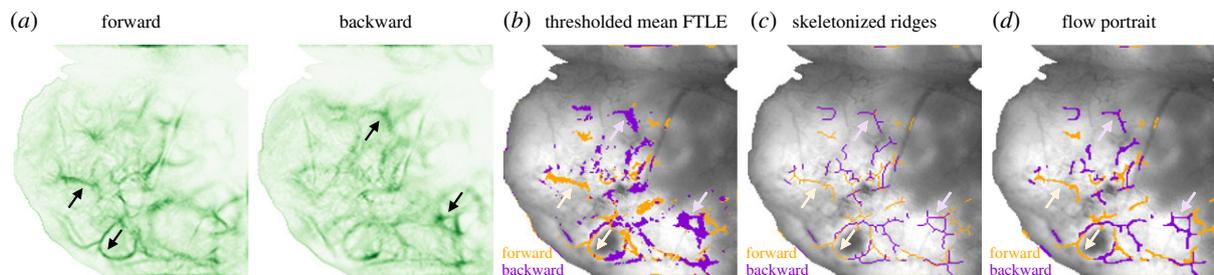


Figure 4. Ridges in the FTLE field are extracted to form the FLOW portrait. (a) The forward and backward FTLE fields are separately averaged to aggregate flow structures over time. Black arrows indicate examples of FTLE ridges which are extracted in the following analysis. (b) Next, the mean FTLE fields are binarized using a threshold which is chosen by the user at a specified percentile (denoted the threshold percentile; here this is chosen to be the 95th percentile). The binary forward and backward FTLE fields are shown overlaid on the mean $\Delta F/F$ image. Pale orange and purple arrows show the same forward and backward ridges as in (a); pale orange corresponds to the forward-time ridges and pale purple to those in backward time. (c) Ridges in the FTLE are approximated by performing a skeletonization procedure on the binarized FTLE fields. (d) Lastly, FLOW portraits are produced by further morphological image processing to smooth the approximate FTLE ridges. The final FLOW portrait highlights the example ridges observed in the original mean FTLE fields.

compressed in others (see also electronic supplementary material, video S1). Relative deformations are described by the Cauchy–Green strain tensor at every pixel, and the FTLE corresponds to the log-normalized leading eigenvalue of this tensor. The same procedure is repeated by reversing the ordering of frames to compute flow maps in backwards time. The forward and backward FTLE fields computed for each example time snapshot are shown in figure 3*c,d*.

Drawing again on our analogy to physical fluid flows, ridges in the FTLE field correspond to time-varying analogues of invariant manifolds, and they approximate LCSs [60,61]. In forward time, these features repel fluid material, similar to a *stable* manifold in a dynamical system. The opposite is true for backward time ridges, where material is attracted in forward time, as with the *unstable* manifold. A similar interpretation can be extended to the FTLE of optical activity flows, where forward-time structures repel activity, while backward-time structures attract activity. However, additional care must be taken when interpreting the intensity of FTLE ridges for brain activity, since the induced velocity field is not divergence free, as is typically the case when analysing incompressible fluid systems. When the velocity field is incompressible, then the determinant of the flow map Jacobian is equal to 1, so the largest eigenvalue is greater than or equal to 1. However, for compressible vector fields (as is the case for widefield imaging of neural activity), the divergence is non-zero and the product of the eigenvalues of the flow map Jacobian may not equal 1. In this case, we may locally have two positive or two negative Lyapunov exponents. Here, we consider only the non-negative Lyapunov exponents, which correspond to repelling ridges in forward time and attractive ridges in backward time (figure 3*c,d*).

2.3. Ridge extraction for FLOW portrait visualization

By aggregating the forward and backward FTLE ridges within a window in time, we summarize the coherent structures of propagating activity within that window with a single FLOW portrait. Ridges of an FTLE field have been shown to approximate LCSs, and several mathematical definitions are suggested to extract them from data [60,62,82,83]. We found that implementing existing strategies for ridge extraction on FTLE fields of widefield calcium imaging data did not adequately extract ridge-like features. Therefore, we developed a post-processing approach to visualize ridges from the forward and backward mean FTLE fields.

Ridges lie along local extremes in a field, thus we can approximate their locations by extracting maximal regions and computing the skeleton structure. To compute the dominant features over the entire recording, we first threshold the mean of all non-negative FTLE values (figure 4*a*) to isolate local maxima in the field (figure 4*b*). Next, we approximate ridges from the local FTLE maxima by performing a morphological skeletonization operation (figure 4*c*). Lastly, these ridges are smoothed by applying morphological image processing (figure 4*d*). Thus, the resulting visualization depicts the average approximate FTLE ridges in a recording window to summarize the time-invariant patterns of activity. We refer to this visualization as FLOW portraits because it is designed for compressible vector fields typical of widefield imaging of calcium activity.

There are two parameters that the user must choose: the integration time T for the flow map Φ_0^T and the threshold percentile for FTLE values to include in the visualization. The choice of these parameters depends on knowledge of the time scales of relevant coherent activity propagation in each dataset. Larger integration time windows filter out shorter time-scale waves; lower percentile thresholds admit more ridges with less intense coherence, which can also admit more spurious ridges if the data are noisy. Electronic supplementary material, figure S4 shows how a range of these parameters yields different FLOW portraits. As a practical recommendation to users, we recommend repeating the computation for a range of parameter values so that the visually salient features in a dataset are reflected in the FLOW portraits.

2.4. How to interpret a FLOW portrait

To build intuition and illustrate how spatiotemporal patterns are visualized by FLOW portraits, let us examine them for several simple synthetic datasets (inspired by those used in [57]), each capturing the types of coherent activity commonly observed in widefield calcium imaging. The first example is a plane wave that starts in the middle of the field of view and travels to the right (figure 5*a*). In the corresponding FLOW portrait, the forward-time FTLE structures delineate where the wave originates in the middle of the field of view, while the backward-time FTLE structures outline where the wave terminates (figure 5*b*). This type of travelling plane wave closely resembles the spread of neural activity observed by widefield imaging (for instance, data from mouse pup in figure 5). The second synthetic dataset is a circular wave that

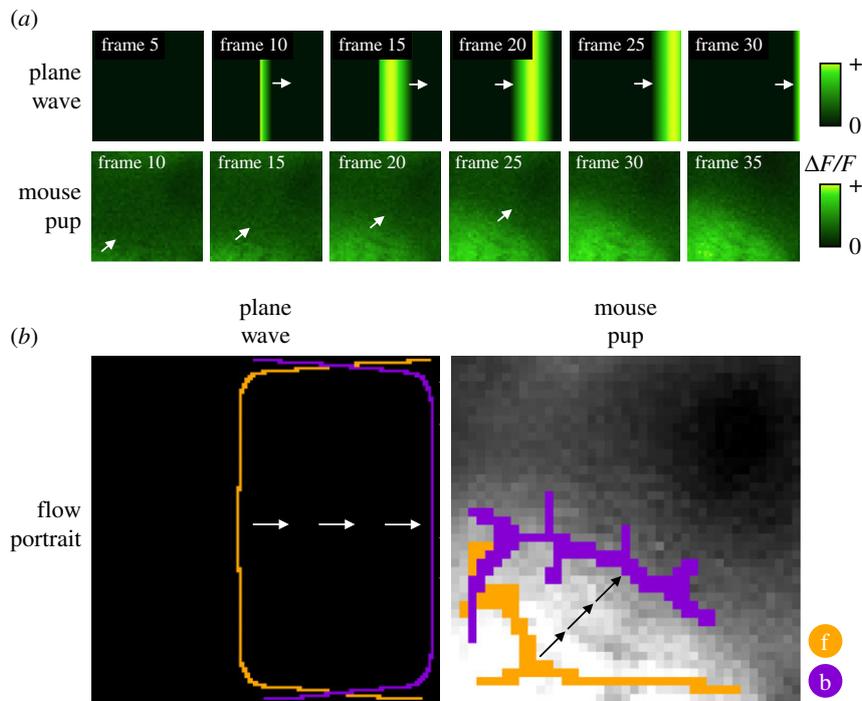


Figure 5. FLOW portraits summarize the activity within a segment of data by highlighting where activity begins and ends. (a) Simple travelling waves for which we illustrate FLOW portraits. The ‘plane-wave’ example shows a synthetic travelling wave which begins in the centre of the frame and travels to the right. The ‘mouse pup’ example shows a short travelling wave within a segment of the widefield calcium imaging experiment of a P7 mouse pup. White arrows indicate the direction of activity propagation. (b) The FLOW portraits for both waves highlight the regions where the wave begins and where the wave ends. Arrows (black and white) show the general direction of wave propagation. FLOW portraits were computed with integration lengths of 15 frames and 5 frames and the threshold percentile was set to the 91st and the 90th percentiles for the plane-wave and mouse pup datasets, respectively.

initiates in the middle, then grows larger towards the edges (electronic supplementary material, figure S2). Here, the forward-time FTLE structures mark the site of initiation, while the backward-time FTLE structures outline the maximal spatial extent of the circle’s spread. Our third synthetic example combines both travelling and growing/shrinking wavefronts. As shown in figure 1 and electronic supplementary material, video S1 and figure S2, the two-dimensional (2D) Gaussian dataset includes a Gaussian blob that appears in the field of view, grows in diameter, translates to the right, then shrinks. Forward-time FTLE structures capture where the activity originates, including the back edge of the Gaussian as it starts to translate and the outside perimeter of the blob as it shrinks. Similarly, backward-time FTLE structures capture where the activity terminates, including the outside perimeter of the blob as it grows and the centre of the shrinking blob.

In all of these examples, FLOW portraits represent succinct summaries of spatiotemporal coherent activity, highlighting regions of activity initiation and termination, as well as the direction and spatial extent of how activity spreads. Specifically, activity originates *from* the forward-time FLOW ridges (orange lines, analogous to stable manifolds) and goes *to* the backward-time FLOW ridges (purple lines, analogous to unstable manifolds). This visualization caricaturizes features of coherent activity not accessible by established methods, including modal decomposition (figure 1), instantaneous metrics like divergence and curl (electronic supplementary material, figure S1) and source/sink classification of fixed points (electronic supplementary material, figures S2 and S3). The forward- and backward-time FTLE structures carry more information than sources and sinks because they are not constrained to be fixed points; thus, these structures are able to delineate travelling fronts. The intersection of two or

more FLOW structures, such as where the orange and purple ridges intersect in figure 1b, can occur for several reasons. First, intersections of the forward and backward FTLE ridges are reflected as intersections in the FLOW portraits. Points where these FTLE ridges intersect correspond to time-dependent saddle points, as the forward and backward FTLE ridges are time-dependent analogues of the stable and unstable manifolds of the vector field. Second, two different spatiotemporal structures may occur at the same spatial location at different times during the recording, as in the case of the 2D Gaussian synthetic dataset.

3. FLOW portraits of widefield calcium imaging data

We demonstrate the application of our approach on several optical widefield datasets, all recordings of spontaneous calcium activation imaged from the cortical surface of transgenic mice. In each example, we have chosen to focus on windows in time when bouts of activity are observed across large portions of cortex. We show that FLOW portraits extracted from these windows summarize the extent and direction of calcium flow, highlighting cortical areas whose neural activations can be interpreted in the context of the behavioural and developmental context of the animals.

3.1. Example 1: Pan-cortical waves

Pan-cortical waves are bouts of activity that propagate across large areas of the cortex [5,84–87] and are suggested to play a critical role in cortical development [38]. These events are defined heuristically as activity that propagates to include a

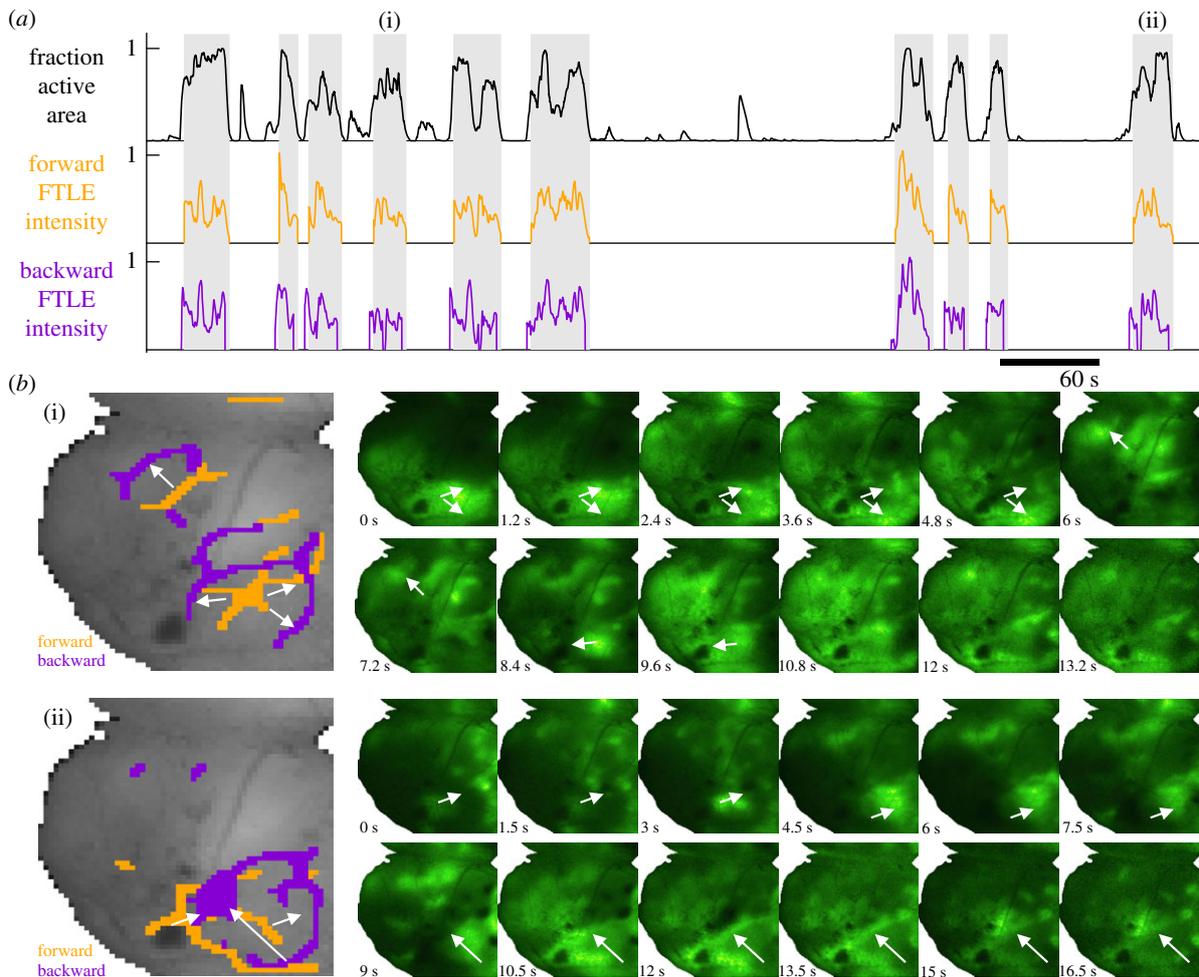


Figure 6. Pan-cortical wave events in a P7 mouse pup are summarized as FLOW portraits. (a) Pan-cortical waves are defined as events where the fraction of active cortex (black trace) exceeds 50%. Briefly, the fraction of active cortex is defined as the fraction of pixels whose intensity is greater than one standard deviation above the mean (in time) for that pixel. FTLE intensity is defined as the sum of the FTLE values for each frame, normalized by the maximum value in time; this intensity is computed for both the forward and backward FTLE time series. (b) FLOW portraits are shown for two example waves, indicated by (i) and (ii) in (a). Orange indicates forward-time FTLE ridges where calcium activity originates. Purple indicates backward-time FTLE ridges where calcium activity propagates towards. White arrows highlight the general direction of activity propagation during the cortical wave. The FLOW portraits are computed using an integration length of 2 s (40 frames) and a threshold percentile of 93%.

large fraction of the imaged cortical surface. In figure 6a, the grey bars highlight individual cortical wave events, defined as when the fraction of active cortex rises to above 1/2 and falls back to the baseline ($\sim 1/10$). To contribute to our understanding of pan-cortical waves in development, we use FLOW portraits to summarize the activity during each wave event, thus facilitating direct comparisons across individual waves and developmental time points.

We construct FLOW portraits to summarize the flow of activity during each pan-cortical wave. Spatial integration of the FTLE fields yields the FTLE intensity (figure 6a, orange and purple traces), which indicates the relative amount of time-averaged *flow* throughout the recording. The resulting FLOW portraits for two pan-cortical waves can be seen in figure 6b, alongside 12 frames of the $\Delta F/F$ data from each wave (see also electronic supplementary material, videos S2 and S3). The portraits of every pan-cortical wave are shown in electronic supplementary material, figure S5.

Each FLOW portrait provides a summary of the prominent activity observed during each wave event, highlighting the regions that repel (forward FTLE, orange) and attract (backward FTLE, purple) activity. Indeed, both waves shown in

figure 6b exhibit two stages of propagation, where activity spreads and pauses briefly at the sensorimotor cortex (outlined by the purple rings) before spreading towards the frontal cortex. This concise visualization allows us to easily compare such qualitative features of wave propagation without having to parse through the raw data manually.

3.2. Example 2: Sleep-state cortical activity changes in development

To further investigate the role of spontaneous cortical activity during development, we analysed optical recordings of spontaneous calcium activity in mouse pups during the first 8 postnatal days of development. We computed FLOW portraits on bouts of spontaneous cortical activity during sleep in 12 animals of ages P1, P2, P3, P5, P7 and P8 (figure 7). Briefly, the sleep state was determined by binning time points into three categories (sleep, wake and moving-wake) using the power of the nuchal electromyography (EMG) spectrum [38,88,89]. We chose to focus on sleep-state cortical activity for its proposed developmental roles and observed changes during development [38]. For each animal, we computed FLOW portraits for up to the 10 longest bouts of sleep

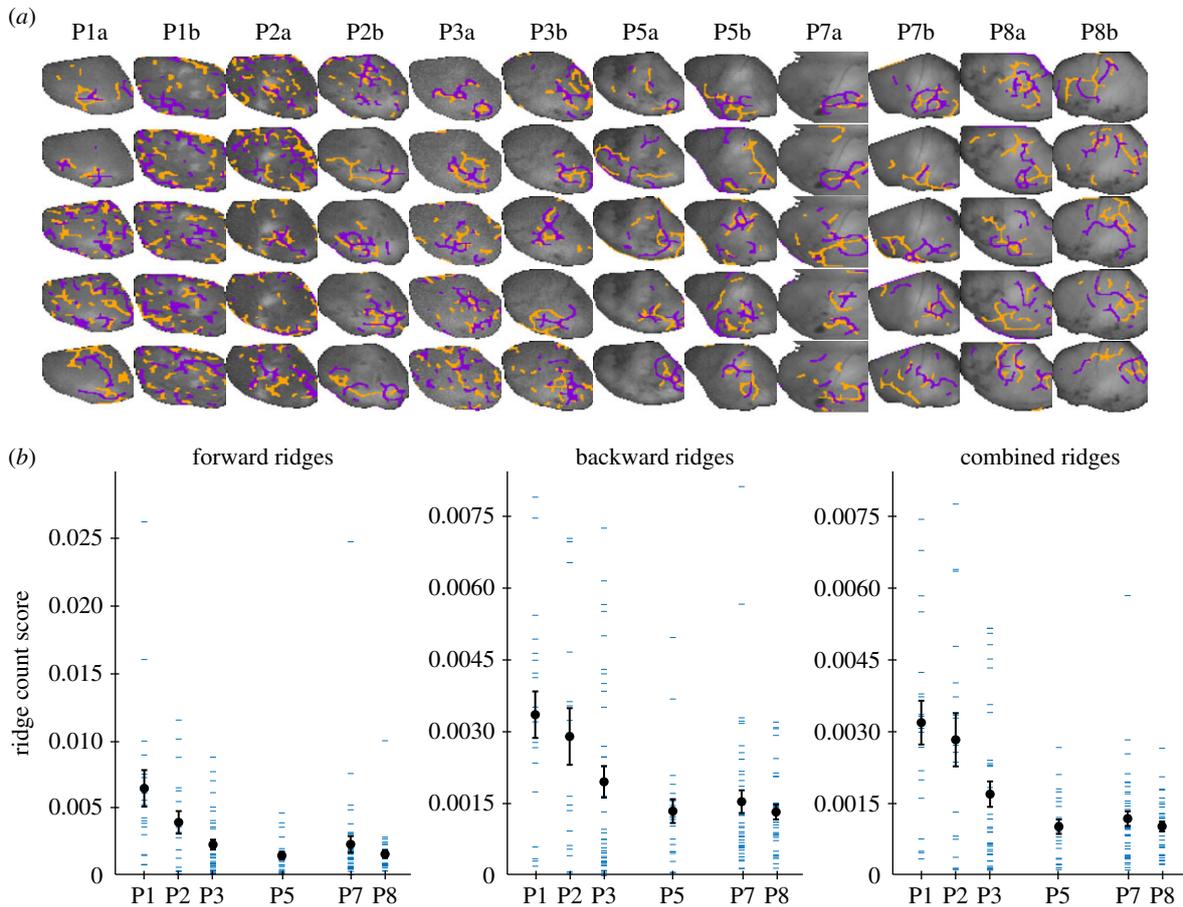


Figure 7. FLOW portraits highlight developmental changes of sleep-state cortical activity. (a) FLOW portraits for five sleep bouts from 12 P1–P8 mouse pups are shown. During the first 3 postnatal days, activity is diffuse, as indicated by many short-length structures in the FLOW portrait. As animals grow older (postnatal days 5–8), sleep-state cortical activity becomes more structured, as indicated by a consolidation of features in the FLOW portraits. Orange indicates repelling structures, and purple indicates attracting structures. All images are of the left hemisphere, such that the mid-line and anterior directions are oriented towards the bottom and left of the images, respectively. FLOW portraits were computed using an integration length of 2 s (40 frames) and a threshold percentile of 93% for all sleep bouts shown. (b) Quantification of the number of FLOW ridges seen versus developmental day. The ridge count score for a FLOW portrait is computed by counting the number of FLOW ridges, either forward, backward or both combined, and dividing by the total area of FLOW ridges in that portrait. When the ridge count score is high there are many smaller ridges in the image, whereas when the score is low there are fewer ridges with a larger ridge area. Here, the mean ridge count score (black point) decreases between developmental day 1 and day 5 and then remains constant. Furthermore, the mean ridge count score over days P1–P3 (0.0038, 0.0027 and 0.0022 for forward, backward and combined, respectively) is significantly different from the mean ridge count score over day P5–P8 (0.0019, 0.0015 and 0.0010 for forward, backward and combined, respectively; paired *t*-test, *p*-values 8.70×10^{-4} , 7.16×10^{-5} and 1.62×10^{-7} for forward, backward and combined, respectively). Blue dashes show individual data points; the ridge count score for an individual FLOW portrait. Error bars show ± 1 s.e. measure.

(fewer portraits were computed for short recordings when there were fewer than 10 sleep bouts).

Five example FLOW portraits for each animal are shown in figure 7a, with a complete set in electronic supplementary material, figure S6. This organization allows us to leverage FLOW portraits to examine developmental changes in cortical activity across long recordings from different animals. We observe a qualitative change between the portraits from the early postnatal days (P1–P3) to the later days (P5–P8).

The FLOW portraits from the early days show more diffuse activity, with fewer consolidated FTLE ridges. After P5, the FLOW portraits show cortical activity during sleep becoming more consolidated and following more defined flow patterns. We quantify this transition to more consolidated FTLE ridges after P5 by defining a ridge count score. Briefly, this metric is computed by counting the total number of disconnected ridges in a FLOW portrait and dividing this sum by the total area of the FLOW portrait. We computed ridge count scores for all FLOW portraits seen in electronic supplementary material, figure S6 and

summarized the mean over each developmental day (figure 7b). We found that the ridge count score for forward FTLE structures, backward FTLE structures and both combined all decreased between P1 and P5. Furthermore, we found that the mean ridge count score over the early developmental days (P1–P3) was significantly different (*p*-values of 8.70×10^{-4} , 7.16×10^{-5} and 1.62×10^{-7} for forward, backward and combined, respectively; paired *t*-test) from that over the later developmental days (P5–P8).

3.3. Example 3: Cortical activity during spontaneous movement

Lastly, we analyse the FLOW portraits of spontaneous cortical activity in a head-fixed, behaving adult mouse [35]. To investigate how FLOW portraits align with an animal's behaviour, we analyse infrared videos of spontaneous facial and limb movements alongside cortical calcium activity. A movement score was assigned to each recording time point by using the total pixel-wise difference between the current and next

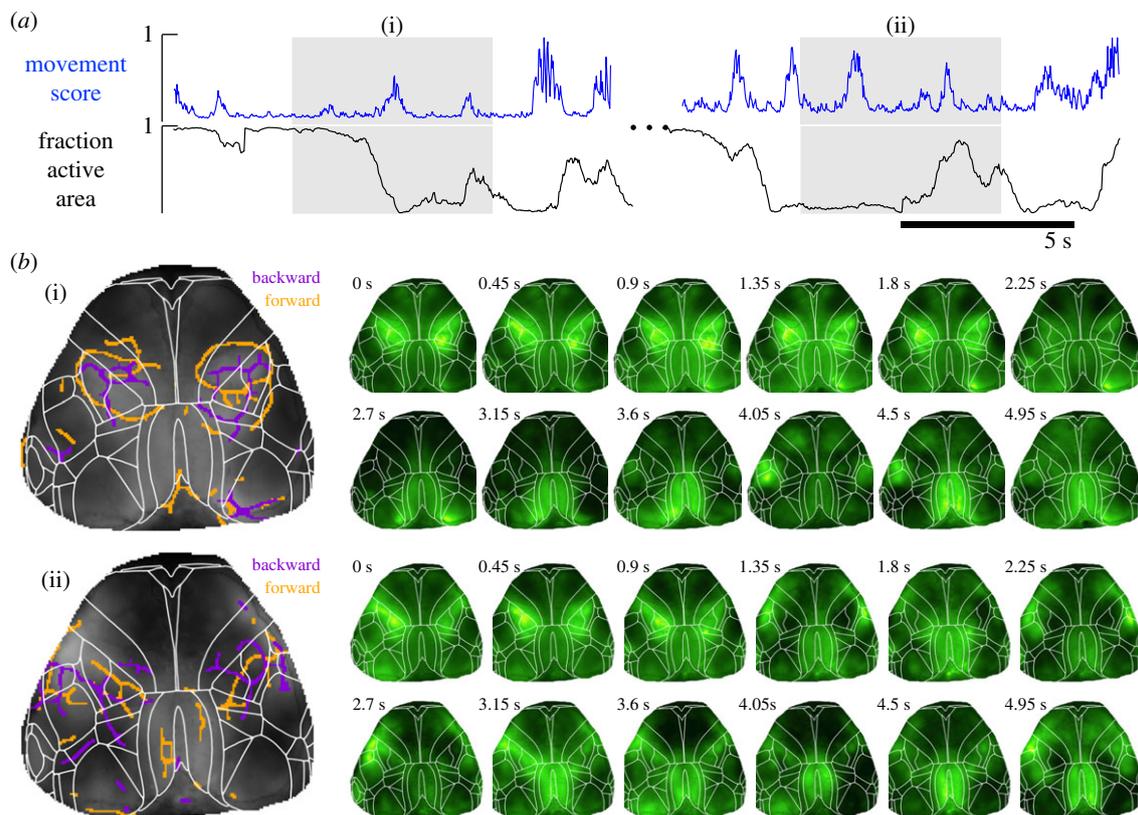


Figure 8. Examples of spontaneous cortical calcium activity associated with movements of an adult mouse summarized as FLOW portraits. (a) A movement score extracted from an infrared video of the mouse moving spontaneously in the dark shows bouts of large movements among more quiescent periods. These bouts of movements do not correspond necessarily to when a large fraction of the cortical surface is active (see Methods for threshold criteria). (b) FLOW portraits for two bouts involving spontaneous movements labelled (i) and (ii) show that coherent structures that highlight activity appear in sensorimotor regions and are then attracted to the centres of these regions bilaterally. Boundaries aligned to the Allen Mouse Brain CCF [90] are overlaid in white. FLOW portraits were computed with a 15-frame integration length and a threshold percentile of 93%.

frames (the forward difference) and normalizing this to the maximum observed difference. During bouts of limb movement or whisking the movement score was greater, approaching the maximum score of 1, than during periods of rest, when the score approached the minimum score of 0.

We chose two bouts of spontaneous movement (grey shading in figure 8a highlights the two bouts, (i) and (ii)) to compute the corresponding FLOW portraits (see also electronic supplementary material, videos S4 and S5). Large variations in the movement score (figure 8a, blue trace) can be observed throughout these bouts, indicating that the animal is continuously switching from a resting to a moving state.

We see signatures of these movement behaviours in the calcium activity, when we expect sensorimotor cortical regions to be more active than during periods of rest. Indeed, the FLOW portrait for each activity bout provides a clear summary of calcium activity surrounding the sensorimotor cortex (figure 8b). During both bouts, a ring-like repelling (forward, orange) field line outlines the sensorimotor region, while attracting (backward, purple) field lines fill in the centres of the rings. We note that these patterns are different from our analysis of example 1 of pan-cortical waves. Specifically, these features suggest a dominant pattern of cortical calcium activity as diffusion of activity from the entirety (or outer edges) of sensorimotor regions towards the centre. In other words, our FLOW portraits point to sensorimotor cortex as a terminus of cortical activity during spontaneous movement behaviours. Interestingly, compared

with the overlaid Allen Mouse Brain Common Coordinate Framework (CCF) (white lines in figure 8b), the attracting (backward, purple) field lines are close to the boundary between somatosensory and primary motor cortices. We note that the integration length and the threshold percentile parameters chosen for these examples determine which ridges are highlighted in the FLOW portraits.

4. Discussion

This paper introduces FLOW portraits as a novel approach to visualize spatiotemporal flow of coherent features in optical widefield calcium imaging data. Viewed at this meso-scale of temporal and spatial resolution, neural activity at the cortical surface is typified by multiple brain regions activating transiently and sometimes in spatial succession. Motivated by an analogy between this flow of neural activity over cortex and physical fluid flows, we leverage techniques well established to study physical fluid flows, in particular the FTLE. Here, we convert videos of $\Delta F/F$ over the cortical surface into vector fields, and the FTLE ridges in these vector fields form an intuitive map of dynamic calcium activity. Importantly, our FLOW portraits do not decompose the data into modes and are not models of the data. Instead, they capture succinct portraits of diverse, variable and non-stationary spatiotemporal patterns, such as those often observed in spontaneous or task-driven widefield calcium imaging experiments.

The FLOW portrait analysis makes several assumptions that are usually true of physical fluid systems but often not met by neural data. Coherent propagation of neural activity on the cortex does not obey mass or energy conservation, so the extraction of FTLE ridges are only approximate 'material' accumulation lines. This assumption is particularly invalid for long bouts of data and over long integration windows, so caution must be exercised in choosing these parameters in the analysis (the same is true of FTLE analysis in fluid flows). Because the integration window effectively low-pass filters the dynamics of the data, activity that is on a faster time scale may be attenuated, and local activity may integrate to appear more coherent. The optimal choice of FTLE parameters for visualization widefield activity and how these depend on spatiotemporal statistics will be important to understand in future applications. Furthermore, although widefield imaging offers much larger fields of view at a higher temporal resolution than many other imaging methods, there remains much unobservable neural activity. Brain areas outside the imaging window and underneath the cortical surface contribute to the imaged activity, yet the flow of neural activity among these regions cannot be captured by our analysis and may bias the extracted flow lines. This limitation is more severe in considering brains with sulci and gyri, as our analysis fundamentally assumes that neighbouring pixels are also neighbours on the cortical sheet.

The quality and interpretability of FLOW portraits requires the imaging data to have been acquired with sufficient temporal and spatial resolution to support the analysis. To be specific, we require that the sampling in time be fast enough that successive frames of the video are very similar. If the frame rate is too slow and neighbouring frames differ substantially, then the optical flow computation infers inaccurate vector fields and can no longer disambiguate between a gradual flow of activity and sudden jumps in activation. Despite the relatively slow dynamics of GCaMP6s compared with single-neuron activity [27], the temporal dynamics of lasting neural synchrony at this meso-scale is adequately matched to the kinetics of the indicator protein in all the data we highlight here. The choice of calcium or voltage indicator also introduces filtering in time, so that our analysis relies on the dynamics of the indicator to be faster than the dynamics of the underlying flow across the brain. Similarly, the spatial resolution of the data need not support disambiguation of single neurons, but it is important that spatial averaging in the field of view does not obscure coherent features of interest.

We suggest that our approach expands the toolbox of techniques to analyse and understand widefield imaging data, especially facilitating direct comparison of multiple bouts of spatiotemporal activity that are interpretable in the context of behaviour and development. This visualization framework can be developed to explicitly quantify features of the flow (for example, the ridge count score analysis in figure 7). Such quantification may be of value in further work that connects features of FLOW portraits with states of relevance to behaviour, development or disease. The transformation of widefield calcium imaging data into a vector field representation suggests multiple avenues for the development of analytic tools. For instance, where multiple coherent waves are present and propagate locally, future work may develop visualizations of the direction of activity propagation, from individual forward FLOW ridges to backward FLOW ridges. Intriguingly, it may be possible to discover partial differential

equations that govern the flow of activity through these vector fields using data-driven techniques [91,92].

5. Methods

5.1. Widefield calcium imaging and data preprocessing

5.1.1. Developing mouse datasets

These experimental procedures were conducted at the University of Washington, and all protocols were reviewed and approved by the University of Washington IACUC. Neonatal mice expressing GCaMP6s in cortical neurons were bred by crossing mice heterozygously expressing an *Emx1*-driven *Cre* (*Emx1-Cre+/-*; Jackson Labs ID 005628) with mice homozygously expressing GCaMP6s under the control of a *cre* promoter (*Ai162+/-*; donated by the Allen Institute; Jackson Labs ID 031562). This cross resulted in mice expressing GCaMP6s primarily in glutamatergic cortical neurons early in development. On the day of recording, mice were placed on a heating pad and anaesthetized using 1–2% isoflurane carried by 100% O₂, while local anaesthetic bupivacaine was delivered subcutaneously at the scalp. The skin over the cortex was removed over a window spanning between the ears to just above the eyes of the pup, to reveal the skull. The periosteum was then removed with fine-tip forceps and cotton swabs. At this developmental stage, the skull is uncalcified and largely transparent, so thinning or cutting a window was unnecessary. A stainless steel U-shaped bracket was then attached to the skull with cyanoacrylate glue. The bracket was clamped in place to the heating pad and stage to stabilize the head. To prevent the skull from drying and to preserve clarity, the exposed skull was also sealed with a thin layer of cyanoacrylate. Silver wire hook leads were implanted into the nuchal muscle through the same incision to monitor neck EMG.

Once the glue had dried, isoflurane anaesthesia was removed and the pup along with the heating pad and stage was positioned for imaging on a Nikon AZ100 with 2× objective and 0.6× reducer. Nuchal EMG activity was amplified with an AM Systems model 1700 amplifier (10 Hz high pass, 60 Hz notch, 10 kHz low pass) and was sampled at 10 kHz using a Powerlab 4/26 and Labchart v. 8 (AD Instruments). GCaMP6s activity was excited using an Intensilight mercury lamp (Nikon), captured using a CCD camera (ORCA Flash 2.8) and recorded using the HCImage application (Hamamatsu). Frame capture rates varied from 10 Hz to 50 Hz with maximum exposure times (100–20 ms, respectively). To further increase the signal-to-noise ratio, the camera was set to perform online hardware-based pixel binning, reducing a 1920 × 1440 pixel image to 960 × 720 pixels. Individual recordings began when the animal began cycling regularly between sleep and wake, and recordings typically lasted between 40 and 60 min, after which the pup was euthanized.

Ca²⁺ records were processed using Matlab (Mathworks) to create $\Delta F/F$ image stacks for FLOW portrait analysis. Briefly, imaging runs were further downsampled by pixel binning the 960 × 720 pixel image down to 480 × 360 pixels. To compensate for slow drift, a moving window of 40 s was used to calculate baseline F for each frame; each pixel in F was set to the minimum value for that pixel across the 40 s window. ΔF was calculated as the difference between the raw pixel intensity and this calculated moving minimum. The difference was then normalized to relative change by dividing ($\Delta F/F$). A small Gaussian spatial blur was used to attenuate 'speckled' noise. Region of interest (ROI) masks of the visible cortical surface were generated by excluding any pixel whose mean-to-variance ratio was greater than 400 : 1. This value was determined heuristically to optimize exclusion of any pixels that displayed minimal change in fluorescence over time, such as those that lie outside the cortical window.

5.1.2. Adult mouse dataset

These experimental procedures were conducted at University College London, UK, according to the UK Animals Scientific Procedures Act (1986) and under personal and project licences granted by the Home Office following appropriate ethics review. The dataset and associated procedures were described previously [35]. In brief, the data were from an adult (30 weeks) male mouse expressing GCaMP6s in excitatory neurons (tetO-GCaMP6s; CaMK2a-tTa genotype [11]). The mouse was implanted with a metal headplate, plastic light isolation chamber and transparent covering over the dorsal skull. On the day of recording, the mouse was head-fixed under the microscope on a stable seat with a rubber wheel underneath the forelimbs. Video cameras captured the frontal aspect of the mouse as well as its eye. Imaging was conducted at 70 Hz with alternating blue and violet illumination, and imaging data were corrected for haemodynamic components. The data were processed by SVD compression.

The images were aligned to the Allen CCF [90] by manually identifying the bregma point and the orientation of the midline in the images. Bregma was taken to be located at the coordinate 5.7 mm anteroposterior in the CCF. Since the pixel size in the camera was known (21.7 $\mu\text{m}/\text{pixel}$), the CCF region boundaries could then be overlaid on the images.

5.2. Imaging analysis

5.2.1. Pan-cortical wave segmentation

Pan-cortical waves, as defined by [38], are cortical activity events where recorded activity spreads over a large area of the imaged cortex. We defined the large cortical area to be when 50% of the cortical pixels (pixels which show the cortex) are active. At any time point, a pixel is active if its intensity is more than 1 s.d. above the temporal mean for that pixel. To extract pan-cortical wave events, we computed the fraction of active cortical pixels throughout the recording, and noted the time points where the active area exceeded the 50% threshold. Each wave event was then defined by the time points when the active area crossed 10% active prior to the time of crossing the 50% threshold and the time when the active area crossed this 10% lower bound following the peak. Overlapping wave events were merged into a signal pan-cortical wave to avoid redundancy. Furthermore, events that lasted less than the FTLE integration length (T) plus the optical flow scaling delay (3.5 s or 70 frames for the mouse pup data) were not analysed because the FTLE and optical flow computations require longer bouts of data.

5.2.2. Sleep bouts during development

Sleep-state cortical activity was segmented using the nuchal EMG as an indicator of state (sleep or awake). Time points were clustered into three groups based on the nuchal EMG power spectrum as in [38,88,89], where the lowest power group is known to represent the sleep state. We defined a sleep bout as a period of continuous classification in the sleep state, and extracted the 10 longest bouts from each recording over the developmental time span. Any bout that did not meet the FTLE and optical flow length requirement (3.5 s or 70 frames for the mouse pup data; 0.8 s or 30 frames for the adult mouse data) was not analysed further. In cases when there were fewer than 10 bouts that met the length requirement, we chose to include fewer sleep bouts for that recording.

5.2.3. Movement event extraction

We extracted movement events from the video of the face and front arms of the adult mouse during the widefield imaging experiment. We defined a movement score for each time point in the video based on the difference between the

current time point and the previous time point. Each video frame was assigned a movement score given by the sum (over all pixels in the frame) of the difference between the current and previous frame. For time point t , the score is given by $\text{MovementScore}_t = \sum_{\text{pixels}} (I_t - I_{t-1})$, where I is the pixel intensity for each of the pixels in the frame. The time series of movement scores was normalized to the maximum observed value for ease of interpretation and visualization. Time stamps of video frames were determined by recording transistor-transistor logic pulses emitted by the camera on each exposure, for both calcium imaging and behavioural videos. We then compared cortical activity across varying movement regimes.

5.3. Optical flow computation

5.3.1. Horn–Schunck optical flow

We computed optical flow vector fields using the HS optical flow algorithm [77] implemented in Matlab [93]. Two parameters must be supplied to the optical flow algorithm: the maximum number of iterations and the α smoothness parameter. Values for both parameters were selected such that the errors in the HS minimization problem (see [77] for details) were simultaneously minimized. We set the maximum number of iterations to 100 and α to 1 for all computations.

5.3.2. Optical flow scaling and smoothing

To minimize the effects of noise on optical flow fields, we applied an activity-based scaling to the magnitudes of the optical flow vectors. First, we created a time series of weights for each pixel by normalizing change in raw pixel intensity between the current time and the intensity of that pixel 1.5 s in the past to the maximum observed change. We chose a time delay of 1.5 and 0.5 s, for the developmental and adult mouse datasets, respectively, to empirically match the time scale of large changes observed in the raw data. Next, we took the sliding windowed average, over a window of 0.25 s, of the weights in order to further reduce the effects of recording noise. We then scaled the magnitude of the optical flow vectors by applying the weights to the corresponding vector. Lastly, we temporally smoothed the optical flow fields using a five-point Gaussian window created with Matlab's `gausswin()` function. The `gausswin` function takes an additional parameter, α , which is proportional to the inverse of the standard deviation of the Gaussian smoothing kernel. We set this parameter to 1.25 for all smoothing operations because of its observed ability to reduce noise in the processed vector fields.

5.4. Finite-time Lyapunov exponent fields

We computed the FTLE of all vector fields using the LCS Tool [73] (<https://github.com/jeixav/LCS-Tool>) in the Matlab software package. We computed the FTLE using an integration length of 2.0 s (40 frames) for the developing mouse data and an integration length of approximately 0.5 s (15 frames) for the adult mouse data. Additionally, we used integration lengths of 15 frames, 12 frames and 10 frames for the plane wave, the circular wave and the travelling Gaussian examples, respectively. To choose the integration length T , we followed the criteria outlined in [62] of choosing a value such that the FTLE ridges are sufficiently resolved. Using a sample of each dataset, we computed the FTLE for a range of integration lengths (0–100 frames) and visualized the resulting FTLE fields. We then chose the smallest integration length where the corresponding FTLE field had well-resolved, sharp ridges. Electronic supplementary material, figure 4B illustrates the effects of computing FLOW portraits with a range of integration lengths.

5.5. FLOW portrait construction

FLOW portraits are constructed through several image-processing steps that aim to extract ridges from an FTLE field (see figure 4 for a visualization of the intermediate processing steps). It is important to note that we process the forward and backward FTLE fields separately and overlay them on the mean $\Delta F/F$ image to create the final FLOW portrait.

We begin by averaging the FTLE time series to aggregate the flow features into mean forward and backward FTLE fields. Next, we isolate possible ridge-like features by thresholding the mean FTLE field at a chosen percentile to form a binarized image. This thresholding step is motivated by recognizing that a ridge can be thought of as a continuous path along a local maximum in the field [60,62]. Therefore, the binarized mean FTLE fields are thought to contain the ridges whose values are above the chosen threshold value. Throughout this work, we denote the specified threshold value as a parameter named the *threshold percentile*. For each FLOW portrait analysis, we chose the threshold percentile to extract the FTLE ridges (see black arrows in figure 4a for example ridges). Figure 4a,b shows the correspondence between the mean FTLE field and the binary versions (the threshold percentile was set to 95%). We used threshold percentiles between 90% and 93% for the mouse pup dataset and a threshold percentile of 93% for the adult mouse dataset. Additionally, we thresholded the plane wave, the circular wave and the travelling Gaussian examples to the 91st, the 93rd and the 85th percentiles, respectively. Electronic supplementary material, figure S4C illustrates that the FLOW portraits compute with a range of threshold percentiles.

Next, we perform two sets of morphological image processing operations on each binarized mean FTLE field to produce the final FLOW portrait. The first set of operations aims to denoise approximate ridges from the FTLE fields, while the second set smoothes the ridges to produce the FLOW portrait. We found that these two series of operations provide strong approximations to the ridge features that we observe in the FTLE fields. We use the `bwmorph()` function in Matlab for all morphological image-processing operations (see <https://www.mathworks.com/help/images/ref/bwmorph.html> for details). This function applies a specified morphological operation iteratively, with the number of iterations specified by the n parameter, or until the input image remains unchanged, $n = \text{Inf}$. Unless otherwise specified, we performed morphological operations until the image no longer changed, with $n = \text{Inf}$. We refer the reader to the Matlab documentation, Gonzales *et al.* [94] and Haralick & Shapiro [95] for the mathematical details of each morphological processing operation used.

The first set of operations aims to transform the noisy, disconnected ridges in the binarized images to connected ridges that reassemble those observed in the raw data. First, we perform the 'close' operation (morphological dilation followed by erosion) to close any gaps within the binary image. Next, we use the 'thin' operation to thin the blob-like structures seen in figure 4b to a series of lines. Lastly, we skeletonize the image by applying the 'skel' (performed with $n=4$). Together, these operations convert the disconnected, blob-like structures seen in figure 4b to the connected, single-pixel structures in

figure 4c. These skeletonized structures can be thought to approximate the centrelines of the FTLE ridges.

The second set of operations aims to smooth the skeletonized image to produce the FLOW portrait. Here, we perform the 'diag' operation to connect regions where two pixels lie corner to corner with an additional pixel. We then apply the 'spur' operation to remove any remaining single-pixel spurs from the ridges. Lastly, we close any gaps introduced with the 'close' operation.

Lastly, we overlay the processed forward and backward images on the corresponding mean $\Delta F/F$ image to create the final FLOW portrait. An example FLOW portrait can be seen in figure 4d.

5.6. Quantification of FLOW portrait consolidation during development

In order to quantify the consolidation of FLOW portraits during development we computed a metric that we denote as the *ridge count score*. The ridge count score is defined as the number of disconnected ridges in the FLOW portrait divided by the total area of the FLOW portrait. This metric is computed by counting disconnected ridges (objects) in the FLOW portrait and dividing by the total number of pixels included in the FLOW portrait. This score was computed for forward FLOW, backward FLOW and both combined for each sleep bout. We then take the mean ridge count score of all sleep bouts from animals of the same developmental age. Lastly, we use a paired t -test to determine whether the mean mean ridge score for developmental days P1–P3 is statistically different from that for developmental days P5–P7.

Data accessibility. Our code is publicly available without restriction, other than citation, on Github at <https://github.com/natejlinden/FLOWPortrait>. The code and data in this repository can reproduce all main analyses, findings and figures from our paper.

Authors' contributions. N.J.L., S.L.B. and B.W.B. conceived of the study and designed the analyses. N.J.L. carried out the analyses. D.R.T., N.A.S. and W.J.M. collected the imaging data and helped interpret the results. N.J.L. and B.W.B. wrote the paper, and all authors contributed to editing the manuscript.

Competing interests. The authors have no competing interests.

Funding. N.J.L. acknowledges support through a Neuroengineering Undergraduate Research Fellowship from the University of Washington Institute of Neuroengineering (UWIN) and the Washington Research Foundation Funds for Innovation in Neuroengineering. D.R.T. acknowledges funding support from the UW Neuroscience Graduate Program (T32NS099578) and the UW Computational Neuroscience Center (5T90DA032436). N.A.S. was supported by the Human Frontiers Science Program (Fellowship LT001071) and the European Union's Horizon 2020 research and innovation programme (Marie Skłodowska-Curie fellowship 656528). W.J.M. acknowledges funding from a Simons Foundation Autism grant. S.L.B. acknowledges funding support from the Army Research Office (ARO W911NF-19-1-0045). B.W.B. and S.L.B. acknowledge funding from the Air Force Office of Scientific Research (AFOSR FA9550-19-1-0386). B.W.B. acknowledges funding from the Washington Research Foundation, the Alfred P. Sloan Foundation and Weill Neurohub.

Acknowledgements. We are grateful for helpful discussions with Aditya Nair, Kameron Decker Harris and Seth Hirsh.

References

1. Wekselblatt JB, Flister ED, Piscopo DM, Niell CM. 2016 Large-scale imaging of cortical dynamics during sensory perception and behavior. *J. Neurophysiol.* **115**, 2852–2866. (doi:10.1152/jn.01056.2015)
2. Muller L, Chavane F, Reynolds J, Sejnowski TJ. 2018 Cortical travelling waves: mechanisms and computational principles. *Nat. Rev. Neurosci.* **19**, 255. (doi:10.1038/nrn.2018.20)
3. Musall S, Kaufman MT, Juavinett AL, Gluf S, Churchland AK. 2019 Single-trial neural dynamics are dominated by richly varied movements. *Nat. Neurosci.* **22**, 1677–1686. (doi:10.1038/s41593-019-0502-4)

4. Corlew R, Bosma MM, Moody WJ. 2004 Spontaneous, synchronous electrical activity in neonatal mouse cortical neurones. *J. Physiol.* **560**, 377–390. (doi:10.1113/jphysiol.2004.071621)
5. Conhaim J, Cedarbaum ER, Barahimi M, Moore JG, Becker MI, Gleiss H, Kohl C, Moody WJ. 2010 Bimodal septal and cortical triggering and complex propagation patterns of spontaneous waves of activity in the developing mouse cerebral cortex. *Dev. Neurobiol.* **70**, 679–692. (doi:10.1002/dneu.20797)
6. Luhmann HJ, Sinning A, Yang J-W, Reyes-Puerta V, Stüttgen MC, Kirischuk S, Kilb W. 2016 Spontaneous neuronal activity in developing neocortical networks: from single cells to large-scale interactions. *Front. Neural Circuits* **10**, 40. (doi:10.3389/fncir.2016.00040)
7. Rossi LF, Wykes RC, Kullmann DM, Carandini M. 2017 Focal cortical seizures start as standing waves and propagate respecting homotopic connectivity. *Nat. Commun.* **8**, 1–11. (doi:10.1038/s41467-016-0009-6)
8. Cramer JV, Gesierich B, Roth S, Dichgans M, Düring M, Liesz A. 2019 In vivo widefield calcium imaging of the mouse cortex for analysis of network connectivity in health and brain disease. *Neuroimage* **199**, 570–584. (doi:10.1016/j.neuroimage.2019.06.014)
9. McGirr A, LeDue J, Chan AW, Xie Y, Murphy TH. 2017 Cortical functional hyperconnectivity in a mouse model of depression and selective network effects of ketamine. *Brain* **140**, 2210–2225. (doi:10.1093/brain/awx142)
10. Siapas AG, Wilson MA. 1998 Coordinated interactions between hippocampal ripples and cortical spindles during slow-wave sleep. *Neuron* **21**, 1123–1128. (doi:10.1016/S0896-6273(00)80629-7)
11. Weksselblatt JB, Niell CM. 2019 Distinct functional classes of excitatory neurons in mouse V1 are differentially modulated by learning and task engagement. *bioRxiv* 533463. (doi:10.1101/533463)
12. Liu M, Song C, Liang Y, Knöpfel T, Zhou C. 2019 Assessing spatiotemporal variability of brain spontaneous activity by multiscale entropy and functional connectivity. *NeuroImage* **198**, 198–220. (doi:10.1016/j.neuroimage.2019.05.022)
13. Feller MB, Wellis DP, Stellwagen D, Werblin FS, Shatz CJ. 1996 Requirement for cholinergic synaptic transmission in the propagation of spontaneous retinal waves. *Science* **272**, 1182–1187. (doi:10.1126/science.272.5265.1182)
14. Feller MB, Butts DA, Aaron HL, Rokhsar DS, Shatz CJ. 1997 Dynamic processes shape spatiotemporal properties of retinal waves. *Neuron* **19**, 293–306. (doi:10.1016/S0896-6273(00)80940-X)
15. Wong RO. 1999 Retinal waves and visual system development. *Annu. Rev. Neurosci.* **22**, 29–47. (doi:10.1146/annurev.neuro.22.1.29)
16. Tiriac A, Smith BE, Feller MB. 2018 Light prior to eye opening promotes retinal waves and eye-specific segregation. *Neuron* **100**, 1059–1065. (doi:10.1016/j.neuron.2018.10.011)
17. Ren C, Komiyama T. 2021 Characterizing cortex-wide dynamics with wide-field calcium imaging. *J. Neurosci.* **41**, 4160–4168. (doi:10.1523/JNEUROSCI.3003-20.2021)
18. Urai AE, Doiron B, Leifer AM, Churchland AK. 2021 Large-scale neural recordings call for new insights to link brain and behavior. (<http://arxiv.org/abs/2103.14662>).
19. Sato TK, Nauhaus I, Carandini M. 2012 Traveling waves in visual cortex. *Neuron* **75**, 218–229. (doi:10.1016/j.neuron.2012.06.029)
20. Dana H, Chen T-W, Hu A, Shields BC, Guo C, Looger LL, Kim DS, Svoboda K. 2014 *Thy1-GCaMP6* transgenic mice for neuronal population imaging in vivo. *PLoS ONE* **9**, e108697. (doi:10.1371/journal.pone.0108697)
21. Stirman JN, Smith IT, Kudenov MW, Smith SL. 2016 Wide field-of-view, multi-region, two-photon imaging of neuronal activity in the mammalian brain. *Nat. Biotechnol.* **34**, 857–862. (doi:10.1038/nbt.3594)
22. Silasi G, Xiao D, Vanni MP, Chen AC, Murphy TH. 2016 Intact skull chronic windows for mesoscopic wide-field imaging in awake mice. *J. Neurosci. Methods* **267**, 141–149. (doi:10.1016/j.jneumeth.2016.04.012)
23. Steinmetz NA *et al.* 2017 Aberrant cortical activity in multiple GCaMP6-expressing transgenic mouse lines. *eneuro* **4**, ENEURO.0207-17.2017. (doi:10.1523/ENEURO.0207-17.2017)
24. Couto J *et al.* 2021 Chronic, cortex-wide imaging of specific cell populations during behavior. *Nat. Protoc.* (doi:10.1038/s41596-021-00527-z)
25. Nakai J, Ohkura M, Imoto K. 2001 A high signal-to-noise Ca²⁺ probe composed of a single green fluorescent protein. *Nat. Biotechnol.* **19**, 137–141. (doi:10.1038/84397)
26. Tian L *et al.* 2009 Imaging neural activity in worms, flies and mice with improved GCaMP calcium indicators. *Nat. Methods* **6**, 875. (doi:10.1038/nmeth.1398)
27. Chen T-W *et al.* 2013 Ultrasensitive fluorescent proteins for imaging neuronal activity. *Nature* **499**, 295–300. (doi:10.1038/nature12354)
28. Vanni MP, Murphy TH. 2014 Mesoscale transcranial spontaneous activity mapping in GCaMP3 transgenic mice reveals extensive reciprocal connections between areas of somatomotor cortex. *J. Neurosci.* **34**, 15 931–15 946. (doi:10.1523/JNEUROSCI.1818-14.2014)
29. McVea DA, Mohajerani MH, Murphy TH. 2012 Voltage-sensitive dye imaging reveals dynamic spatiotemporal properties of cortical activity after spontaneous muscle twitches in the newborn rat. *J. Neurosci.* **32**, 10982–10994. (doi:10.1523/JNEUROSCI.1322-12.2012)
30. Song C, Piscopo DM, Niell CM, Knöpfel T. 2018 Cortical signatures of wakeful somatosensory processing. *Sci. Rep.* **8**, 1–12.
31. Scott BB, Thiberge SY, Guo C, Tervo DGR, Brody CD, Karpova AY, Tank DW. 2018 Imaging cortical dynamics in GCaMP transgenic rats with a head-mounted widefield microscope. *Neuron* **100**, 1045–1058. (doi:10.1016/j.neuron.2018.09.050)
32. Allen WE *et al.* 2017 Global representations of goal-directed behavior in distinct cell types of mouse neocortex. *Neuron* **94**, 891–907. (doi:10.1016/j.neuron.2017.04.017)
33. Pinto L, Rajan K, DePasquale B, Thiberge SY, Tank DW, Brody CD. 2019 Task-dependent changes in the large-scale dynamics and necessity of cortical regions. *Neuron* **104**, 810–824. (doi:10.1016/j.neuron.2019.08.025)
34. Jacobs EA, Steinmetz NA, Peters AJ, Carandini M, Harris KD. 2020 Cortical state fluctuations during sensory decision making. *Curr. Biol.* **30**, 4944–4955.e7. (doi:10.1016/j.cub.2020.09.067)
35. Zatka-Haas P, Steinmetz NA, Carandini M, Harris KD. 2021 Sensory coding and the causal impact of mouse cortex in a visual decision. *eLife* **10**, e63163. (doi:10.7554/eLife.63163)
36. Wright PW *et al.* 2017 Functional connectivity structure of cortical calcium dynamics in anesthetized and awake mice. *PLoS ONE* **12**, e0185759. (doi:10.1371/journal.pone.0185759)
37. Vanni MP, Chan AW, Balbi M, Silasi G, Murphy TH. 2017 Mesoscale mapping of mouse cortex reveals frequency-dependent cycling between distinct macroscale functional modules. *J. Neurosci.* **37**, 7513–7533. (doi:10.1523/JNEUROSCI.3560-16.2017)
38. Tabuena DR, Huynh R, Metcalf J, Richner T, Stroth A, Brunton BW, Moody WJ, Easton CR. Submitted. Pan-cortical waves in the neonatal rodent brain *in vivo*: a precursor of adult sleep waves?
39. Holmes P, Guckenheimer J 1983 *Nonlinear oscillations, dynamical systems, and bifurcations of vector fields*, vol. 42. Applied Mathematical Sciences. Berlin, Germany: Springer.
40. Pang R, Lansdell BJ, Fairhall AL. 2016 Dimensionality reduction in neuroscience. *Curr. Biol.* **26**, R656–R660. (doi:10.1016/j.cub.2016.05.029)
41. Cunningham JP, Byron MY. 2014 Dimensionality reduction for large-scale neural recordings. *Nat. Neurosci.* **17**, 1500–1509. (doi:10.1038/nn.3776)
42. Dyer EL, Azar MG, Perich MG, Fernandes HL, Naufel S, Miller LE, Körding KP. 2017 A cryptography-based approach for movement decoding. *Nat. Biomed. Eng.* **1**, 967–976. (doi:10.1038/s41551-017-0169-7)
43. Ganguli S, Sompolinsky H. 2012 Compressed sensing, sparsity, and dimensionality in neuronal information processing and data analysis. *Annu. Rev. Neurosci.* **35**, 485–508. (doi:10.1146/annurev-neuro-062111-150410)
44. Churchland MM, Cunningham JP, Kaufman MT, Foster JD, Nuyujukian P, Ryu SI, Shenoy KV. 2012 Neural population dynamics during reaching. *Nature* **487**, 51–56. (doi:10.1038/nature11129)
45. Gallego JA, Perich MG, Miller LE, Solla SA. 2017 Neural manifolds for the control of movement. *Neuron* **94**, 978–984. (doi:10.1016/j.neuron.2017.05.025)
46. Gallego JA, Perich MG, Chowdhury RH, Solla SA, Miller LE. 2020 Long-term stability of cortical population dynamics underlying consistent behavior. *Nat. Neurosci.* **23**, 260–270. (doi:10.1038/s41593-019-0555-4)
47. Cunningham JP, Ghahramani Z. 2015 Linear dimensionality reduction: survey, insights, and generalizations. *J. Mach. Learn. Res.* **16**, 2859–2900.

48. Brunton BW, Johnson LA, Ojemann JG, Kutz JN. 2016 Extracting spatial–temporal coherent patterns in large-scale neural recordings using dynamic mode decomposition. *J. Neurosci. Methods* **258**, 1–15. (doi:10.1016/j.jneumeth.2015.10.010)
49. Tu JH, Rowley CW, Luchtenburg DM, Brunton SL, Kutz JN. 2014 On dynamic mode decomposition: theory and applications. *J. Comput. Dyn.* **1**, 391–421. (doi:10.3934/jcd.2014.1.391)
50. Kutz JN, Brunton SL, Brunton BW, Proctor JL. 2016 *Dynamic mode decomposition: data-driven modeling of complex systems*. Philadelphia, PA: SIAM.
51. MacDowell CJ, Buschman TJ. 2020 Low-dimensional spatiotemporal dynamics underlie cortex-wide neural activity. *Curr. Biol.* **30**, 2665–2680. (doi:10.1016/j.cub.2020.04.090)
52. Saxena S *et al.* 2020 Localized semi-nonnegative matrix factorization (LocaNMF) of widefield calcium imaging data. *PLoS Comput. Biol.* **16**, e1007791. (doi:10.1371/journal.pcbi.1007791)
53. Mackevicius EL, Bahle AH, Williams AH, Gu S, Denisenko NI, Goldman MS, Fee MS. 2019 Unsupervised discovery of temporal sequences in high-dimensional datasets, with applications to neuroscience. *Elife* **8**, e38471. (doi:10.7554/eLife.38471)
54. Zhou P *et al.* 2018 Efficient and accurate extraction of *in vivo* calcium signals from microendoscopic video data. *Elife* **7**, e28728. (doi:10.7554/eLife.28728)
55. Zanos TP, Mineault PJ, Nasiatou KT, Guitton D, Pack CC. 2015 A sensorimotor role for traveling waves in primate visual cortex. *Neuron* **85**, 615–627. (doi:10.1016/j.neuron.2014.12.043)
56. Blankenship AG, Hamby AM, Firl A, Vyas S, Maxeiner S, Willecke K, Feller MB. 2011 The role of neuronal connexins 36 and 45 in shaping spontaneous firing patterns in the developing retina. *J. Neurosci.* **31**, 9998–10008. (doi:10.1523/JNEUROSCI.5640-10.2011)
57. Afrashteh N, Inayat S, Mohsenzand M, Mohajerani MH. 2017 Optical-flow analysis toolbox for characterization of spatiotemporal dynamics in mesoscale optical imaging of brain activity. *Neuroimage* **153**, 58–74. (doi:10.1016/j.neuroimage.2017.03.034)
58. Townsend RG, Gong P. 2018 Detection and analysis of spatiotemporal patterns in brain activity. *PLoS Comput. Biol.* **14**, e1006643. (doi:10.1371/journal.pcbi.1006643)
59. Haller G. 2002 Lagrangian coherent structures from approximate velocity data. *Phys. Fluids* **14**, 1851–1861. (doi:10.1063/1.1477449)
60. Shadden SC. 2011 *Lagrangian coherent structures*, pp. 59–89. Hoboken, NJ: Wiley.
61. Haller G. 2015 Lagrangian coherent structures. *Annu. Rev. Fluid Mech.* **47**, 137–162. (doi:10.1146/annurev-fluid-010313-141322)
62. Shadden SC, Lekien F, Marsden JE. 2005 Definition and properties of Lagrangian coherent structures from finite-time Lyapunov exponents in two-dimensional aperiodic flows. *Physica D* **212**, 271–304. (doi:10.1016/j.physd.2005.10.007)
63. Mathur M, Haller G, Peacock T, Ruppert-Felsot JE, Swinney HL. 2007 Uncovering the Lagrangian skeleton of turbulence. *Phys. Rev. Lett.* **98**, 144502-1–144502-4. (doi:10.1103/physrevlett.98.144502)
64. Green MA, Rowley CW, Haller G. 2007 Detection of Lagrangian coherent structures in 3D turbulence. *J. Fluid Mech.* **572**, 111–120. (doi:10.1017/S0022112006003648)
65. Brunton SL, Rowley CW. 2010 Fast computation of FTLE fields for unsteady flows: a comparison of methods. *Chaos* **20**, 017503. (doi:10.1063/1.3270044)
66. Farazmand M, Haller G. 2012 Computing Lagrangian coherent structures from their variational theory. *Chaos* **22**, 013128. (doi:10.1063/1.3690153)
67. Peng J, Dabiri J. 2009 Transport of inertial particles by Lagrangian coherent structures: application to predator–prey interaction in jellyfish feeding. *J. Fluid Mech.* **623**, 75–84. (doi:10.1017/S0022112008005089)
68. Duvernois V, Marsden AL, Shadden SC. 2013 Lagrangian analysis of hemodynamics data from FSI simulation. *Int. J. Num. Methods Biomed. Eng.* **29**, 445–461. (doi:10.1002/cnm.2523)
69. Shadden SC, Arzani A. 2015 Lagrangian postprocessing of computational hemodynamics. *Ann. Biomed. Eng.* **43**, 41–58. (doi:10.1007/s10439-014-1070-0)
70. Mohajerani MH *et al.* 2013 Spontaneous cortical activity alternates between motifs defined by regional axonal projections. *Nat. Neurosci.* **16**, 1426. (doi:10.1038/nn.3499)
71. Ashby DM, LeDue J, Murphy TH, McGirr A. 2019 Peripheral nerve ligation elicits widespread alterations in cortical sensory evoked and spontaneous activity. *Sci. Rep.* **9**, 1–10. (doi:10.1038/s41598-019-51811-8)
72. Jia H, Rochefort NL, Chen X, Konnerth A. 2011 *In vivo* two-photon imaging of sensory-evoked dendritic calcium signals in cortical neurons. *Nat. Protoc.* **6**, 28–35. (doi:10.1038/nprot.2010.169)
73. Onu K, Huhn F, Haller G. 2015 LCS Tool: a computational platform for Lagrangian coherent structures. *J. Comput. Sci.* **7**, 26–36. (doi:10.1016/j.jocs.2014.12.002)
74. Paragios N, Chen Y, Faugeras OD. 2006 *Handbook of mathematical models in computer vision*. Berlin, Germany: Springer Science & Business Media.
75. Duffy CJ, Wurtz RH. 1991 Sensitivity of MST neurons to optic flow stimuli. II. Mechanisms of response selectivity revealed by small-field stimuli. *J. Neurophysiol.* **65**, 1346–1359. (doi:10.1152/jn.1991.65.6.1346)
76. Krapp HG, Hengstenberg R. 1996 Estimation of self-motion by optic flow processing in single visual interneurons. *Nature* **384**, 463–466. (doi:10.1038/384463a0)
77. Horn BK, Schunck BG. 1981 Determining optical flow. *Artif. Intell.* **17**, 185–203. (doi:10.1016/0004-3702(81)90024-2)
78. Willert CE, Gharib M. 1991 Digital particle image velocimetry. *Exp. Fluids* **10**, 181–193. (doi:10.1007/BF00190388)
79. Westerweel J. 1997 Fundamentals of digital particle image velocimetry. *Meas. Sci. Technol.* **8**, 1379. (doi:10.1088/0957-0233/8/12/002)
80. Taira K *et al.* 2017 Modal analysis of fluid flows: an overview. *AIAA J.* **55**, 4013–4041. (doi:10.2514/1.J056060)
81. Taira K, Hemati MS, Brunton SL, Sun Y, Duraisamy K, Bagheri S, Dawson S, Yeh C-A. 2020 Modal analysis of fluid flows: applications and outlook. *AIAA J.* **58**, 998–1022. (doi:10.2514/1.J058462)
82. Garth C, Gerhardt F, Tricoche X, Hans H. 2007 Efficient computation and visualization of coherent structures in fluid flow applications. *IEEE Trans. Vis. Comput. Graph.* **13**, 1464–1471. (doi:10.1109/TVCG.2007.70551)
83. Lipinski D, Mohseni K. 2010 A ridge tracking algorithm and error estimate for efficient computation of Lagrangian coherent structures. *Chaos* **20**, 017504. (doi:10.1063/1.3270049)
84. Garaschuk O, Linn J, Eilers J, Konnerth A. 2000 Large-scale oscillatory calcium waves in the immature cortex. *Nat. Neurosci.* **3**, 452–459. (doi:10.1038/74823)
85. Conhaim J *et al.* 2011 Developmental changes in propagation patterns and transmitter dependence of waves of spontaneous activity in the mouse cerebral cortex. *J. Physiol.* **589**, 2529–2541. (doi:10.1113/jphysiol.2010.202382)
86. Easton CR, Weir K, Scott A, Moen SP, Barger Z, Folch A, Hevner RF, Moody WJ. 2014 Genetic elimination of gabaergic neurotransmission reveals two distinct pacemakers for spontaneous waves of activity in the developing mouse cortex. *J. Neurosci.* **34**, 3854–3863. (doi:10.1523/JNEUROSCI.3811-13.2014)
87. Barger Z, Easton CR, Neuzil KE, Moody WJ. 2016 Early network activity propagates bidirectionally between hippocampus and cortex. *Dev. Neurobiol.* **76**, 661–672. (doi:10.1002/dneu.22351)
88. Seelke AM, Blumberg MS. 2010 Developmental appearance and disappearance of cortical events and oscillations in infant rats. *Brain Res.* **1324**, 34–42. (doi:10.1016/j.brainres.2010.01.088)
89. Blumberg MS, Gall AJ, Todd WD. 2014 The development of sleep–wake rhythms and the search for elemental circuits in the infant brain. *Behav. Neurosci.* **128**, 250. (doi:10.1037/a0035891)
90. Wang Q *et al.* 2020 The Allen Mouse Brain Common Coordinate Framework: a 3D reference atlas. *Cell* **181**, 936–953. (doi:10.1016/j.cell.2020.04.007)
91. Rudy SH, Brunton SL, Proctor JL, Kutz JN. 2017 Data-driven discovery of partial differential equations. *Sci. Adv.* **3**, e1602614. (doi:10.1126/sciadv.1602614)
92. Schaeffer H. 2017 Learning partial differential equations via data discovery and sparse optimization. *Proc. R. Soc. A* **473**, 20160446. (doi:10.1098/rspa.2016.0446)
93. Kharbat M. 2009 Horn-Schunck optical flow method. MATLAB Central File Exchange. See <https://www.mathworks.com/matlabcentral/fileexchange/22756-horn-schunck-optical-flow-method>.
94. Gonzalez C, Woods RE, Eddins LS. 2009 *Digital image processing using MATLAB*. Knoxville, TX: Gatesmark Publishing.
95. Haralick RM, Shapiro LG. 1992 *Computer and robot vision*, vol. 1. Reading, MA: Addison-Wesley.